Far-field Correlation Electromagnetic Analysis attacks against AES in real world applications

F. van Tienen

Technische Universiteit Delft

Far-field Correlation Electromagnetic Analysis attacks against AES in real world applications

by

F. van Tienen

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday July 16, 2018 at 2:30 PM.

Student number:	4094123	
Project duration:	January 4, 2018 – July 16, 2018	
Thesis committee:	Assoc. Prof. Dr. Ir. J.C.A. van der Lubbe,	TU Delft, supervisor
	Assist. Prof. Dr. C. Doerr	TU Delft
	J. Haesakkers,	Fox-IT, supervisor
	G. Stoevelaar,	Fox-IT
	R. Boix Carpi	Riscure BV

Abstract

In almost every device cryptographic functions are used to protect data and sensitive information from being intercepted. A commonly used encryption algorithm is the Advanced Encryption Standard (AES), which is a symmetric block cypher. Side-channel attacks against AES are well known and are often performed either directly on the surface of the integrated circuit or by attaching wires to the target device. These attacks are more difficult for devices with tamper protection, which can detect such an attack because the device enclosure must be removed. This limits the attack possibilities of these side-channel attacks for these real-world applications. Attacks using electromagnetic radiation from a further distance, called the far-field, can be used to prevent opening the enclosure.

For these power based side-channel attacks against AES, power traces must be recorded with the exact timing before an encryption or decryption starts. This is used to align the recorded traces and perform statistical analysis to extract the secret key. In order to achieve this often an General Purpose Input Output (GPIO) trigger is used to indicate the start of a new trace. When such an GPIO trigger is not available a smart trigger can be used, which uses a pattern to generate a trigger based on the measured power. This removes the need for making a connection with the target device.

In this thesis an approach for performing non-invasive far-field side-channel attacks against multiple target devices is evaluated. For this approach near-field analysis is performed to analyse the target leakage, with the use of Test Vector Leakage Assessment. Then for each of these targets far-field side-channel attacks are attempted using a Software Defined Radio set-up and several smart triggers are tested in real-world scenarios. The results of these attacks showed that far-field side-channel attacks without an artificial trigger are possible, and thus the enclosure of the target device can stay in tact and non-invasive attacks can be performed. For the Microsemi SF2 Basic development kit attacks up to a distance of 15 cm can be achieved in an office environment. This means that far-field side-channel attacks against AES are possible in real world applications, and when designing cryptographic devices precautions must be made to protect against these attacks.

Acknowledgements

This thesis would not have been possible without the support of my supervisors, colleagues, family and friends. I would like to thank my supervisors Jos Haesakkers Gerke Stoevelaar from Fox-IT for all their advise and motivation during my thesis, and for giving me the opportunity to perform this awesome project. Rafael Boix Carpi from Riscure has supervised and supported me with his knowledge, for which I am deeply grateful for. My thesis supervisor from the TU Delft, Jan van der Lubbe, I would like to thank for giving the opportunity to perform my own research and his advise during the meetings.

During my master I have also spend a lot of time in the MAVLab and I would like to thank them for the opportunities to work on cool research projects and participating in the Outback Challenge. I would also want to thank all the colleagues from the MAVLab for the wonderful time, their support and help during my master.

Elisabeth has always supported me during my master study and has taken care of me whenever needed, for which I am extremely grateful. I would also like to thank my family and parents for their support and the opportunity to perform this study. My friends and other close relatives I would like to thank for everything they have provided me with.

I would like to thank my colleagues at Fox-IT for providing a nice and friendly working environment, from which I got a lot of inspiration. Without Riscure I would have never been able to achieve these results, and I would like to thank them for their knowledge and tools which I used during this thesis. Both at Fox-IT and Riscure I would like to thank everyone for their unconditional support. As last I would like to thank OSPL for giving the opportunity to perform measurements inside an anechoic room, which was an awesome experience.

F. van Tienen
Delft, July 2018

Contents

Acronyms

ADC Analog to Digital Converter.

AES Advanced Encryption Standard.

ARIA Academia, Research Institute and Agency.

CEMA Correlation Electromagnetic Analysis.

CMOS Complementary Metal Oxide Semiconductor.

CPA Correlation Power Analysis.

CRT Cathode Ray Tube.

DC Direct Current.

DDR Double Data Rate.

DEMA Differential Electromagnetic Analysis.

DES Data Encryption Standard.

DPA Differential Power Analysis.

DSP Digital Signal Processor.

DVB-T Digital Video Broadcasting Terrestrial.

ECC Eliptic Curve Cryptography.

EM Electromagnetic.

EMA Electromagnetic Analysis.

FFT Fast Fourier Transformation.

FIPS Federal Information Processing Standard.

FM Frequency Modulated.

FPGA Field-programmable Gate Array.

GPIO General Purpose Input Output.

HW Hamming Weight.

IC Integrated Circuit.

IF Intermediate Frequency.

LNA Low Noise Amplifier.

PA Power Analysis.

PCB Printed Circuit Board.

PLL Phase-locked loop.

PSD Power Spectral Density.

RF Radio Frequency.

ROM Read-only Memory.

RSA Rivest–Shamir–Adleman.

SAD Sum of the Absolute Difference.

SCA Side-channel Analysis.

SDR Software Defined Radio.

SEMA Simple Electromagnetic Analysis.

SF2 SmartFusion2.

SNR Signal to Noise Ratio.

SoC System on a Chip.

SOM System on Module.

SPA Simple Power Analysis.

TVLA Test Vector Leakage Assessment.

UDP User Datagram Protocol.

1
Introduction

Almost every data processing device uses cryptographic functions to protect sensitive information, hence cryptographic devices can be found everywhere. Most of these devices make use of the Advanced Encryption Standard (AES) encryption[7], which is a Federal Information Processing Standard (FIPS)[53] since 2001. It is the successor of the Data Encryption Standard (DES)[52], which in contrary to AES is considered to be insecure[4, 24].

A lot of different implementations exist of the AES encryption, but writing a correct implementation can be a difficult task. Many applications and devices use standard implementations like the one from The OpenSSL Project [54]. This implementation is widely tested and considered to be cryptographically correct. But this doesn't mean that this implementation protects against every possible attack. Unintentionally secret data can also be leaked through so-called side-channels. These leakages depend on the actual implementation of the AES cipher and the platform, an integrated circuit, which is running the encryption.

One of these side-channels could be the electromagnetic field, which is almost directly related to the power usage of a device. Because of the commonly used Complementary Metal Oxide Semiconductor (CMOS) technology in integrated circuits, the power consumption is related to the data processed in the chip. This results in a relation between the generated electromagnetic field and the data processed in a chip. Previous research in this field focused on the interception of electromagnetic leakage on platforms such as telephones and televisions. When this research became more known other side-channels and leakages of devices were discovered. This leads to the advanced attacks against cryptographic chips we currently see in the state of the art research.

The field of analysing electromagnetic radiation started with the research from Van Eck [56]. In his paper he showed that it is possible to reconstruct a video signal from the electromagnetic radiation of a Cathode Ray Tube (CRT) tv system. He makes use of 2 different antenna's to receive both the video signal and the synchronisation signals of the CRT television. These signals are then processed and fed to a TV receiver which will display the image of the intercepted CRT television. He concludes that under optimal circumstances it may be feasible to reconstruct to the signal at distances up to 1 km. This shows that the electromagnetic leakage could lead to people eavesdropping from far away. This research is later extended by Marinov [31] for digital signals.

Before Van Eck [56] already a lot of research was done by the military in the late 1960s[23]. This research was kept secret and was known as *compromising emanation* or *TEMPEST*. The US military used the code word *TEMPEST* for the program that researched emission security and developing security standards. Later the word *TEMPEST* became a synonym for compromising emissions especially in the far-field. But even though this research was never fully published it shows that compromising emissions are used in real-world scenarios.

In near-field electromagnetic side-channels already a lot of research has been done against various symmetric cryptographic devices and implementations[1, 18, 28, 59]. The most commonly used side-channel attacks against symmetric block ciphers are Differential Electromagnetic Analysis (DEMA) and

Correlation Electromagnetic Analysis (CEMA). Both of these attacks are very successful in retrieving the secret key from various symmetric cryptographic devices and implementations in the near-field. But from further distances like the *TEMPEST* attacks from the military and Van Eck [56], these attacks aren't achieved in realistic scenarios. In Kim et al. [21] far-field attacks of 1 meter distance are shown against the lesser known ARIA, which is a Korean Standard block cipher algorithm. It makes use of similar cryptographic functions as the AES encryption, which raised the question of also AES is vulnerable to far-field electromagnetic side-channels.

Ramsay [45] filled this gap and was able to recover the full cryptographic key at a distance of 1 meter of a device running an AES cipher. But these results were only achieved in an anechoic room, where the electromagnetic interference of other signals are minimized. Also an electromagnetic trigger was added before each encryption, which is needed to recover the important parts from the recorded electromagnetic field. In real-world scenarios a lot of electromagnetic interference is introduced, and also such artificial trigger is not available.

This opens the question if it is possible to retrieve the full cryptographic key in more realistic situations of a cryptographic device running an AES cipher. In order to answer this question research is performed at Fox-IT together with the help of Riscure.

1.1. Research question

In the state of the art research it currently is possible to perform far-field DEMA and CEMA attacks against various platforms, running a symmetric block cipher. But all of the attacks make use of a predefined trigger implemented on the target platform. In real-world applications these triggers are not available without opening the devices' enclosure, which will make these attacks semi-invasive. The combining of Software Defined Radio (SDR) and far-field CEMA gives us new opportunities to implement smarter trigger systems, due to the fact that calculations can be done after the recordings have finished. In order to analyse these capabilities the following research question is constructed:

Can far-field Electromagnetic (EM) side-channel attacks be achieved against an Integrated Circuit (IC) running AES in real world scenarios?

The research question is divided into the following sub-questions:

1. Do different platforms leak secret information in similar ways in the EM near-field?
2. Can smart-triggers be used for far-field EM side-channel attacks when using SDR?
3. Can information from the near-field be used to determine the far-field leakage?

1.2. Contributions

The following contributions are made:

1. The analysis of the near-field electromagnetic leakage on the SmartFusion2 (SF2) Basic, SF2 Advanced and the PYNQ-Z1 running AES.
2. The possibility of smart-triggers in the electromagnetic far-field is verified for specific targets.
3. Near-field leakage analysis is compared to far-field leakage analysis.
4. Far-field side-channel attacks in real world scenarios without a trigger against AES are performed using SDR.

1.3. Report Outline

In Chapter 2 the background needed to understand the research is explained, which includes the AES cipher, SDR and electromagnetic attacks. Next in Chapter 3 the state of the art research for smart-triggers and DEMA attacks, in the far-field and using SDR, are described to give an introduction about the research performed in these topics. Then in Chapter 4 the measurement set-ups will be described in detail, which are used for the baseline and far-field results, and is divided into he Baseline, SDR, Analog Front-end, Analysis Software and Target Devices. In Chapter 5 the baseline results for each of the three targets will be discussed, based on the distance from the target up to 2 cm. In Chapter 6 the results from the far-field will be discussed, also based on the three different targets which are first verifying the baseline results at 2 cm and then discussing the results from the far-field. Finally all the results from the thesis are discussed and an outlook for further research will be given in Chapter 7.

2

Background

2.1. AES

The Advanced Encryption Standard (AES) is a symmetric block cipher which is also known as the Rijndael cipher[6] with a block size of 128 bits. It was proposed in 1999 as the Advanced Encryptions Standard by Daemen and Rijmen [7] and because of that it became the most commonly used symmetric cryptographic cipher. To understand the DEMA and CEMA attacks, and the different leakages of platforms, first a good understanding of the AES cipher is needed. Especially the different implementations affect the leakage of the platforms, and will described in this section.

There are three different variants of the AES cipher, namely AES-128, AES-192 and AES-256[16]. Those variants are respectively using a 128, 192 and 256 bit key size. Since the leakage of the platform itself is not affected by the different variants of AES almost always AES-128 is chosen for simplicity. We will also choose for the AES-128 variant and for the coming sections the AES-128 cipher will be referred to as AES cipher.

To understand the AES algorithm we will first define a *State* array, which is used in the different operations of the AES algorithm. This *State* array represents the state at a given time and consists of 128 bits, the block size, which are presented as 16 bytes. The input bytes *in* are copied to the *State* array as illustrated in Figure 2.1. The single bytes of the state are represented a $S_{r,c}$, where $0 \leq r \leq 4$ represents the row in the matrix and $0 \leq c \leq 4$ represents the column. The number of columns and rows in the state, Nb, are 4. In Figure 2.1 can also be seen how the *State* is copied to the output bytes *out*.

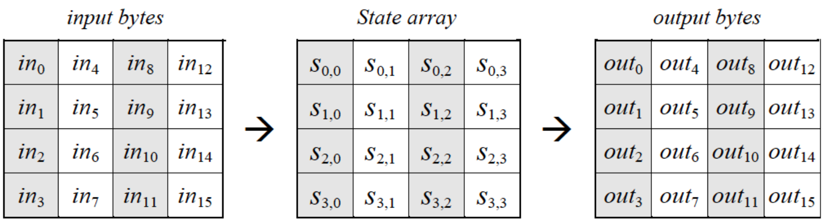

in_0	in_4	in_8	in_{12}
in_1	in_5	in_9	in_{13}
in_2	in_6	in_{10}	in_{14}
in_3	in_7	in_{11}	in_{15}

$s_{0,0}$	$s_{0,1}$	$s_{0,2}$	$s_{0,3}$
$s_{1,0}$	$s_{1,1}$	$s_{1,2}$	$s_{1,3}$
$s_{2,0}$	$s_{2,1}$	$s_{2,2}$	$s_{2,3}$
$s_{3,0}$	$s_{3,1}$	$s_{3,2}$	$s_{3,3}$

out_0	out_4	out_8	out_{12}
out_1	out_5	out_9	out_{13}
out_2	out_6	out_{10}	out_{14}
out_3	out_7	out_{11}	out_{15}

Figure 2.1: The AES state array input and output from FIPS [16]

The algorithm consists of two separate parts: Round Operations and Key Expansion. The AES-128 cipher consists of 10 rounds. At first the Key Expansion is run to generate Round Keys for each of the 10 rounds. The encryption then starts with an *AddRoundKey* operation. After that 9 rounds consisting of the following operations are executed:

1. *SubBytes*
2. *ShiftRows*
3. *MixColumns*

4. *AddRoundKey*

As last a final round will be executed containing of the following operations, omitting the *Mixcolumns* operation:

1. *SubBytes*
2. *ShiftRows*
3. *AddRoundKey*

Each of the above mentioned operations will be described in the next section and a full overview of the encryption algorithm can be seen in Figure 2.2. The input text and secret key are fed into the separate parts of the algorithm and the encrypted output will be generated.

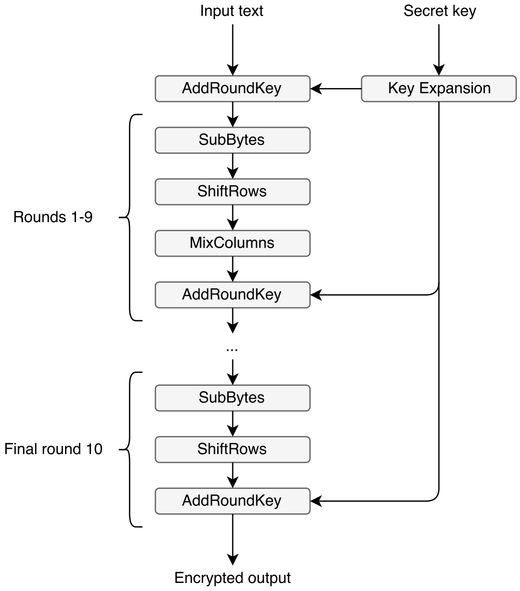

Figure 2.2: Overview of the AES-128 encryption cipher rounds

2.1.1. Round Operations

The AES cipher makes use of following round operations: *SubBytes*, *ShiftRows*, *MixColumns*, *AddRoundKey*. Each of these operations will be described briefly in this section.

SubBytes

The *SubBytes* operation is a non-linear byte substitution that operates on each individual byte of the *State*. It makes use of a substitution table, S-Box, which is invertible and can be seen in Figure 2.4. This table is constructed by composing two transformations which are described in depth in FIPS [16], and the matrix form of these transformations is shown in Equation 2.1.

$$\begin{bmatrix} b'_0 \\ b'_1 \\ b'_2 \\ b'_3 \\ b'_4 \\ b'_5 \\ b'_6 \\ b'_7 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \\ b_7 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} \tag{2.1}$$

In Figure 2.3 can be seen how the S-Box transformation is applied to the *State*. Each individual byte in the *State* matrix is substituted by the S-Box transformation.

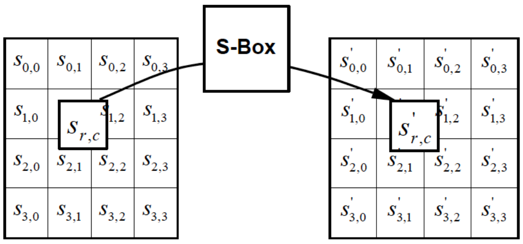
S-Box
$s_{0,0}$ $s_{0,1}$ $s_{0,2}$ $s_{0,3}$
$s_{1,0}$ $s_{r,c}$ $s_{1,2}$ $s_{1,3}$
$s_{2,0}$ $s_{2,1}$ $s_{2,2}$ $s_{2,3}$
$s_{3,0}$ $s_{3,1}$ $s_{3,2}$ $s_{3,3}$
$s'_{0,0}$ $s'_{0,1}$ $s'_{0,2}$ $s'_{0,3}$
$s'_{1,0}$ $s'_{r,c}$ $s'_{1,2}$ $s'_{1,3}$
$s'_{2,0}$ $s'_{2,1}$ $s'_{2,2}$ $s'_{2,3}$
$s'_{3,0}$ $s'_{3,1}$ $s'_{3,2}$ $s'_{3,3}$

Figure 2.3: The AES SubBytes operation from FIPS [16]

		y															
		0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
	0	63	7c	77	7b	f2	6b	6f	c5	30	01	67	2b	fe	d7	ab	76
	1	ca	82	c9	7d	fa	59	47	f0	ad	d4	a2	af	9c	a4	72	c0
	2	b7	fd	93	26	36	3f	f7	cc	34	a5	e5	f1	71	d8	31	15
	3	04	c7	23	c3	18	96	05	9a	07	12	80	e2	eb	27	b2	75
	4	09	83	2c	1a	1b	6e	5a	a0	52	3b	d6	b3	29	e3	2f	84
	5	53	d1	00	ed	20	fc	b1	5b	6a	cb	be	39	4a	4c	58	cf
	6	d0	ef	aa	fb	43	4d	33	85	45	f9	02	7f	50	3c	9f	a8
	7	51	a3	40	8f	92	9d	38	f5	bc	b6	da	21	10	ff	f3	d2
x	8	cd	0c	13	ec	5f	97	44	17	c4	a7	7e	3d	64	5d	19	73
	9	60	81	4f	dc	22	2a	90	88	46	ee	b8	14	de	5e	0b	db
	a	e0	32	3a	0a	49	06	24	5c	c2	d3	ac	62	91	95	e4	79
	b	e7	c8	37	6d	8d	d5	4e	a9	6c	56	f4	ea	65	7a	ae	08
	c	ba	78	25	2e	1c	a6	b4	c6	e8	dd	74	1f	4b	bd	8b	8a
	d	70	3e	b5	66	48	03	f6	0e	61	35	57	b9	86	c1	1d	9e
	e	e1	f8	98	11	69	d9	8e	94	9b	1e	87	e9	ce	55	28	df
	f	8c	a1	89	0d	bf	e6	42	68	41	99	2d	0f	b0	54	bb	16

Figure 2.4: The AES S-Box table for byte *xy* (in hexadecimal format) from FIPS [16]

ShiftRows

In the *ShiftRows* operation each of the rows in the *State* are cyclically shifted. The amount of shifting is determined by the row number. For example the first row is shifted by 0 and the second row is shifted by 1. This transformation can be seen in Figure 2.5. The first row remains the same, while the other rows in the *State* matrix are shifted.

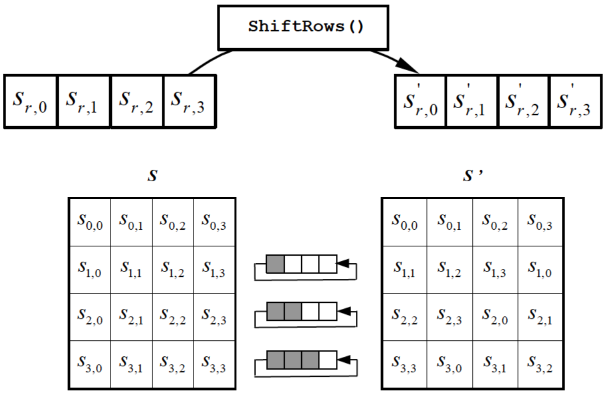
ShiftRows()
$s_{r,0}$ $s_{r,1}$ $s_{r,2}$ $s_{r,3}$
$s'_{r,0}$ $s'_{r,1}$ $s'_{r,2}$ $s'_{r,3}$
S
S'
$s_{0,0}$ $s_{0,1}$ $s_{0,2}$ $s_{0,3}$
$s_{1,0}$ $s_{1,1}$ $s_{1,2}$ $s_{1,3}$
$s_{2,0}$ $s_{2,1}$ $s_{2,2}$ $s_{2,3}$
$s_{3,0}$ $s_{3,1}$ $s_{3,2}$ $s_{3,3}$
$s_{0,0}$ $s_{0,1}$ $s_{0,2}$ $s_{0,3}$
$s_{1,1}$ $s_{1,2}$ $s_{1,3}$ $s_{1,0}$
$s_{2,2}$ $s_{2,3}$ $s_{2,0}$ $s_{2,1}$
$s_{3,3}$ $s_{3,0}$ $s_{3,1}$ $s_{3,2}$

Figure 2.5: The AES ShiftRows transformation from FIPS [16]

MixColumns

The *MixColumns* operation is transforming the *State* column-by-column. Each of the columns is treated as a four-term polynomial, which is described in FIPS [16]. This four-term polynomial transformation is described in Equation 2.2 as a matrix multiplication.

$$\begin{bmatrix} s'_{0,c} \\ s'_{1,c} \\ s'_{2,c} \\ s'_{3,c} \end{bmatrix} = \begin{bmatrix} 02 & 03 & 01 & 01 \\ 01 & 02 & 03 & 01 \\ 01 & 01 & 02 & 03 \\ 03 & 01 & 01 & 02 \end{bmatrix} \begin{bmatrix} s_{0,c} \\ s_{1,c} \\ s_{2,c} \\ s_{3,c} \end{bmatrix} \text{for } 0 \leq c \leq Nb \tag{2.2}$$

In Figure 2.6 can be seen how the MixColumns transformation is executed on the *State* matrix. On each of the columns the matrix transformation from 2.2 is executed, resulting in the new *State*.

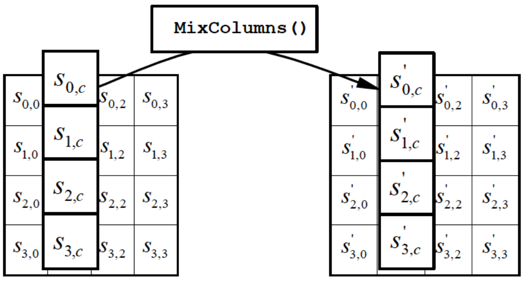

Figure 2.6: The AES MixColumns transformation

AddRoundKey

During the *AddRoundKey* operation the *RoundKey*, W_l, is generated by the Key Expansion algorithm. Each byte of the *RoundKey* will be bitwise XOR-ed with the input *State*, resulting into the new output *State*. In Figure 2.7 this operation is described as a transformation on the *State* matrix. Each input column is XOR-ed with the round key W_{l+c}, resulting in the new *State*.

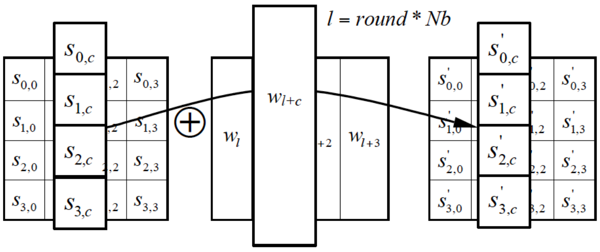

Figure 2.7: The AES AddRoundKey transformation from FIPS [16]

2.1.2. Key Expansion

The Key Expansion algorithm starts of with the input key, which in our case is 128-bits. This means that the number of AES rounds will thus be $Nr = 10$. The algorithm can be seen in Listing 2.1.2, where *Key* is the input key consisting of $Nk = 4$ words and W is the expanded output key in words. $Nb = 8$, because the number of blocks in bytes is 16 but in this algorithm is expressed as words.

```
KeyExpansion(byte Key[4*Nk], word W[Nb*(Nr+1)])
{
  for(i = 0; i < Nk; i++)
    W[i] = (Key[4*i],Key[4*i+1],Key[4*i+2],Key[4*i+3]);
  for(i = Nk; i < Nb * (Nr + 1); i++)
  {
    temp = W[i - 1];
    if (i % Nk == 0)
      temp = SubWord(RotWord(temp)) ^ Rcon[i / Nk];
    W[i] = W[i - Nk] ^ temp;
  }
}
```

Listing 2.1: AES key expansion algorithm in pseudo code from FIPS [16]

As can be seen in the algorithm, the first *Nk* words consist of the secret input key itself. Meaning that the first round during an encryption will always use the secret key. The *Rcon* variable is a round constant which is XORed with the first word of each round key. The *SubWord* preforms a *SubBytes* operation on each of the bytes in the word. The *RotWord* shifts the bytes in the word one position to the right similar to *ShiftRows*. All the other words of the expanded key are just plain XORed with the previous value.

This algorithm is reverse-able, meaning that if a single round key is to be found all the other round keys can be calculated. This is only applicable to AES-128, because AES-192 and AES-256 use a slightly different key scheduling algorithm. For AES-256 and AES-192 you need 2 consecutive round keys in order to calculate all the other round keys. Since the algorithm puts the secret input key as first round, this means that for AES-128 this key can be retrieved by retrieving any of the round-keys. This is particularly useful for side-channel attacks such as DEMA and CEMA, which is explained in the next sections.

2.1.3. Implementations

Since side-channel attacks against AES are implementation dependant, and to be able to understand these differences we will describe two different implementations of the AES encryption. These main differences exist between the implementation of the S-Box lookup, which is also the most common attack vector in side-channels. This is due to the fact that the S-Box lookup is the only non-linear function in AES. While there are many different implementations, we chose to only describe three commonly used implementations in this section.

S-Box

The S-Box lookup table implementation is straightforward. Instead of calculating the S-Box value, the whole S-Box table is saved to memory and a simple lookup is performed. In Figure 2.4 this table can be seen, which is 16 × 16 = 256 bytes. This lookup is extremely efficient for systems with fast memory accesses.

The electromagnetic leakage of this implementation is caused by the memory addresses accessing the individual table values. Depending on the target platform these memory address connections can be contained inside the chip or extend to the Printed Circuit Board (PCB). Whenever external memory is used and these lines extend to the PCB, leakage is to be expected between the memory and microcontroller chip.

T-Tables

The T-Tables implementation is an optimisation for 32-bit processors, but will increase the amount of memory needed. It was first introduced by the original standard from FIPS [16]. The OpenSSL Project [54] uses the T-Tables implementation, which is widely used in many projects. The T-Tables implementation combines the *SubBytes*, *ShiftRows* and the *MixColumns* operation into the *TTable* lookup operation. The resulting new AES encryption can be seen in Figure 2.8. Which also shows that the last round makes use of the *TTableLast* operation, because the *MixColumns* operation is omitted from that round.

The *TTable* operation combines the *SubBytes*, *ShiftRows* and the *MixColumns* operation into for different table lookups per row. Every calculation is done per row of 4 bytes, corresponding to 32 bits, to

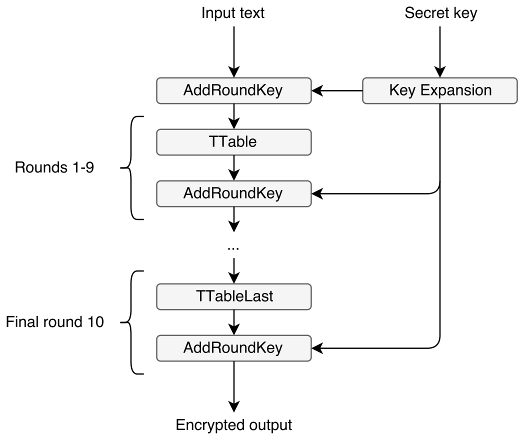

Figure 2.8: Overview of the AES-128 encryption cipher rounds for the T-Table implementation

optimise for 32-bit processor architectures. The *TTableLast* operations combines only the *SubBytes* and *ShiftRows* operations. To calculate the lookup tables for *TTable* and *TTableLast* a substitution of the combined operations is calculated.

Combinatorial
For FPGA's specific implementations exist for the *SubBytes* function. Also in FPGA's it is possible to store the full S-Box table in memory, but in Mui et al. [36] a combinatorial implementation is presented. This implementation calculates the S-Box values based on the combinatorial logic from FIPS [16].
The advantage of this implementation is that no memory is needed to store the S-Box values. But on the other side the combinatorial implementation makes use of more logical gates, compared to the S-Box implementation. Also the electromagnetic most likely will differ, because different electrical current will flow through the Field-programmable Gate Array (FPGA).

2.2. SDR

For electromagnetic side-channel attacks a recording device is needed. The most commonly used device is an oscilloscope, which can record signals with a high bandwidth and fast sample rate. In real-world situations it is more difficult to make recordings with an oscilloscope, due to the limited amount of recording memory and cost. But research showed that even with lower sample rates and a down converted signal a Software Defined Radio (SDR) can be used for electromagnetic attacks against cryptographic devices[35, 45]. These devices are more affordable and can record in real-time, which makes them more useful for real-world scenarios. In this section a short description is given how such SDR is capable of recording electromagnetic signals.

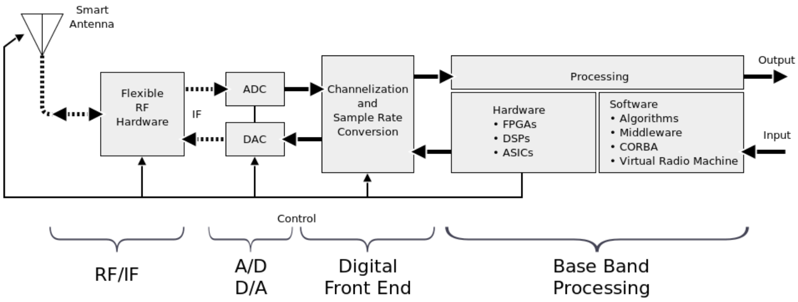

Figure 2.9: A block diagram of an typical SDR device from Wikipedia, the free encyclopedia [58].

In Figure 2.9 we can see how a typical SDR functions. On the left side the electromagnetic signals are picked up by an antenna and passed through to flexible Radio Frequency (RF) Hardware. In this flexible RF Hardware the signal is down converted through a mixer with a variable Baseband Frequency. Because of this down conversion, signals with a higher frequency than the sample rate of the Analog to Digital Converter (ADC) can be recorded. The SDR's split the baseband frequency into a 0° and a 90° phase shifted signal, respectively the *I* and *Q* samples. An example is given in Figure 2.10.

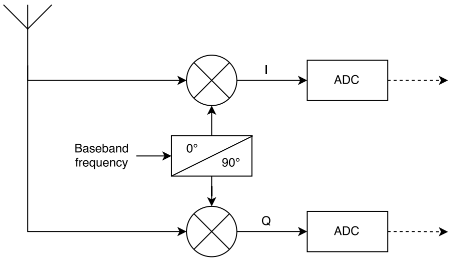

Figure 2.10: Overview of the SDR mixing and phase shifting done in the flexible RF hardware.

After the signals are processed by the RF hardware they are sampled by an ADC and forwarded to the digital front-end. This digital front-end performs channelisation and sample rate conversion and will prepare the IQ-signals for transmission to the computer software. All the base-band processing and configuring of the RF hardware is done through FPGA's and a Digital Signal Processor (DSP) in order to achieve high sample rates.

After data is transferred to the computer it will be processed by software. This software can then perform the signal analysis needed for side-channel attacks. Most SDR's output both the *I* and the *Q* signal, which are 90° phase shifted. But because we are only interested in the power consumption of the processor we will only use $P = \sqrt{I^2 + Q^2}$ during this master thesis to calculate the power. This removes the need of up-converting the signal back to the original baseband frequency.

2.3. Electromagnetic Attacks

Electromagnetic attacks against cryptographic devices fall under the category of side-channel attacks. These side-channel attacks are usually classified into three different categories: invasive, semi-invasive and non-invasive[30, 57]. The invasive side-channel attacks require physical access to the device and will result in tamper evidence. The semi-invasive side-channel attacks also requires physical access to the device, but does not provide tamper evidence of the device. The last category is the non-invasive, where no physical access to the device is required and thus also tamper evidence of the device is avoided.

Between side-channel attacks also a second categorisation can be made: active and passive attacks[30, 57]. The active attacks will tamper with the devices' functionality, for example introducing faults to change the expected behaviour during computation. While the passive attacks will only observe and/or measure the devices' behaviour, such as measuring the computing time of a specific operation.

This thesis will focus on the semi-invasive passive and non-invasive passive attacks in the electromagnetic field. First the Electromagnetic Radiation of the device will be described, then a relation to the different Leakage Models will be given. After that the different methodologies used in Electromagnetic Attack will be explained.

2.3.1. Electromagnetic Radiation

Because of the dependency between electromagnetic radiation and the intermediate state of the AES algorithm, it is possible to detect the AES key using the electromagnetic radiation. To understand the electromagnetic radiation of an IC running cryptographic algorithms a description of CMOS technology will be given. Next the electromagnetic field generated by an IC will be explained. In the last section the difference and definitions of the near-field and far-field will be defined.

Complementary CMOS

Almost every IC used in cryptographic devices, makes use of CMOS technology[43]. A CMOS cell consists of a p-type and n-type MOSFET arranged in a complementary structure as can be seen. Where the p-type MOSFET forms a pull-up network and the n-type MOSFET forms a pull-down network, which are constructed in such a way that the networks are never conducting at the same time.

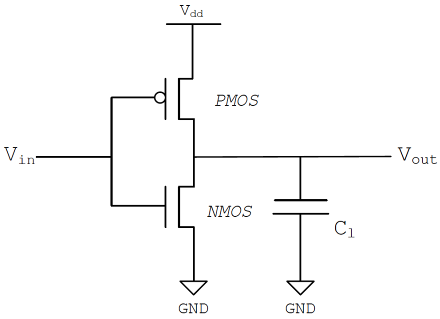

Figure 2.11: CMOS inverter

Every logic cell, such as a NAND or OR gate, in an IC is constructed using this complementary CMOS technology. An example of a CMOS inverter can be seen in Figure 2.11. During static operation, when fixed inputs are applied at the input, the CMOS transistor will have almost no current flowing between V_{dd} and GND. This current is called I_{leak}, representing the leakage current going through the transistor. The static power can be expressed as $P_{stat} = I_{leak} \cdot V_{dd}$.

The dynamic power of the CMOS transistor, P_{dyn}, is based on the data being processed by the logic cell. Because of the complementary construction in CMOS, this dynamic power in a CMOS transistor is much higher than the static power. Considering the different state changes of a CMOS transistor, we can define the following power consumptions:

- $0 \rightarrow 0 : P_{stat}$
- $1 \rightarrow 0 : P_{stat} + P_{dyn}$
- $0 \rightarrow 1 : P_{stat} + P_{dyn}$
- $1 \rightarrow 1 : P_{stat}$

This dynamic power is the part that is used in Power Analysis (PA) attacks and also Electromagnetic Analysis (EMA) attacks. In an IC this dynamic power of multiple transistors will consists of P_{op}, the operational power, and P_{data}, based on the operation data. Where the operational power consists of the instruction to execute. Resulting in $\sum P_{dyn} = P_{op} + P_{data}$. Both the operational power and data power can be used in side-channel attacks, but when using the AES algorithm only the P_{data} will be used. This is due to the fact that the AES operations do not depend on the cryptographic key.

Electromagnetic Field

The electromagnetic field differs in several ways from the power analysis side-channel. The electromagnetic side-channel is a three dimensional vector field that changes over time. The power analysis side-channel is only a simple amplitude waveform representing the current at a specific time. The electromagnetic emanations in an IC come from the different current flows within the IC, as can be explained by Maxwell's Equations[33]. The most interesting current flow is the P_{dyn} of the CMOS transistors, as these contain both the operational and data power.

There are broadly two type of electromagnetic emanations: Direct Emanations and Unintentional Emanations[1]. Direct emanations are generated by a current carrying element, while unintentional emanations are origination from modulated effects. The direct emanations mostly come from sharp rising edges in current flows, which will result in emanations which are observable in a wide frequency band. For the unintentional emanation, the most commonly used carrier signal is the square-wave clock-signal.

Near-Field and Far-Field

When dealing with the electromagnetic field a distinction can be made between the near-field and the far-field. This definition is not clearly defined in literature, as several different similar definitions exist. But for this thesis we use the definition used in De Mulder [8]. She defines the near-field as $r < \lambda/(2\pi)$, where r is the distance from the target and λ the wavelength of the signal. The definition of the far-field is respectively $r > \lambda/(2\pi)$.

2.3.2. Leakage Models

In order to analyse the data leakage of a given device a model must be made which describes the power usage of the given data. There are two common leakage models used for EMA attacks, these are the Hamming weight model and the Hamming distance model[43].

Hamming Weight Model

The Hamming weight model is a very simple model assuming that the electromagnetic emanation is linear to the number of bits that are high, '1', during an operation. The formula is is given in Equation 2.3, where the number of bits is given by n and d_j represent bit j in D. Where D represents the input data (byte or word).

$$HW(D) = \sum_{j=0}^{n-1} d_j \tag{2.3}$$

The assumptions of this model are that each bit uses the same amount of power, and that there is a linear relation between the number of bits that are '1' and the power. This power will then generate the electromagnetic emanations. But as explained before most IC's make use of CMOS technology, which don't have a linear relation between the number of bits that are '1' and the power.

Hamming Distance Model

Because the Hamming weight model doesn't make use of the dynamic leakage of CMOS transmitters, the Hamming distance model was introduced by Brier et al. [5]. The Hamming distance model, models the switching of the bits in CMOS transistors and assumes that there is no difference in power usage between a 0 → 1 and 1 → 0. Using Equation 2.3 we define the Hamming distance Equation 2.4 between D and D'. Where D and D' are two input bytes or words.

$$HD(D, D') = HW(D \oplus D') \tag{2.4}$$

Also the Hamming distance model assumes that each bit uses the same amount of power. But differentiates from the Hamming weight model in that it assumes that the power usage is linear to the $HW(D \oplus D')$. This leakage model will represent the P_{data} going through multiple CMOS transistors. This model doesn't account for the entire power consumption, but because the bus lines are often the most power consuming elements of an IC it can model those parts very well.

2.3.3. Side-channel Attacks

In the electromagnetic field there are mainly two different attacks: Simple Electromagnetic Analysis (SEMA) and Differential Electromagnetic Analysis (DEMA). Correlation Electromagnetic Analysis (CEMA) is a slight variation of the DEMA attack, but the naming for these attacks are used to describe both attacks. We will describe the different attack techniques in more detail in this section. Also a leakage assessment method, Test Vector Leakage Assessment (TVLA), will be described in order to define the leakage of a cryptographic device and to analyse the leakage behaviour of the different targets.

SEMA

Simple Electromagnetic Analysis (SEMA) comes from the field of Power Analysis. In Kocher et al. [22] Simple Power Analysis (SPA) is introduced. With Simple Power Analysis (SPA) a single trace could retrieve the secret key. There are two different SPA attacks: visual and template. Which can both also be performed in the electromagnetic field, where instead of power, electromagnetic emanation is analysed.

With visual SPA attacks an attacker will visually analyse the Power trace over time and tries to determine the different operations performed by the cryptographic device. This means that there must be a dependency between the secret key and the operations performed by the device, which will result in a leakage in P_{op}. But since most AES implementations currently run a high frequencies and due to the noise generated by other parts in the chip, this is currently very hard to perform. This technique is most commonly used against asymmetric encryption algorithms such as Rivest–Shamir–Adleman (RSA), which perform different branching operations based on the secret key.

Template attacks make use of a multivariate distribution to characterise the power usage over time. At the first stage templates are made of the cryptographic device, while performing operations with known data and secret key. These templates are then used to determine the probability of a certain secret key on the real cryptographic device where the secret key is not known. These templates have the advantage that instead of only including P_{op}, it could also include P_{data}. A disadvantage of the template attack is that you need to characterise these templates on a similar device. Also the search space for generating these characterisation could grow immensely when modelling P_{op} and P_{data}.

DEMA and CEMA

Another attack using power analysis or electromagnetic emanation are respectively Differential Power Analysis (DPA) and Differential Electromagnetic Analysis (DEMA) [22, 28]. These attacks make use of the varying amount of power used when processing data[44]. Differential Power Analysis (DPA) and DEMA are also used in literature as a name for Correlation Power Analysis (CPA) and Correlation Electromagnetic Analysis (CEMA). Correlation Power Analysis (CPA) and CEMA compared to DPA and DEMA only differ in the way the statistical analysis of the data is performed, but the general process is similar. An overview for DEMA and CEMA attacks can be seen in Figure 2.12.

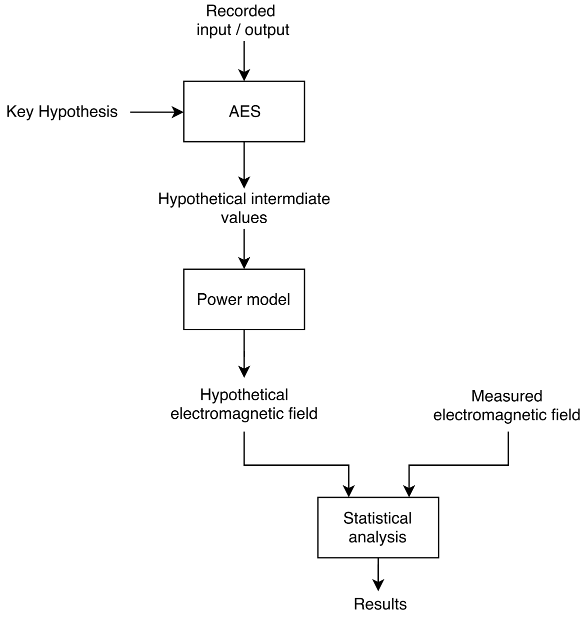
Recorded input / output
Key Hypothesis
AES
Hypothetical intermdiate values
Power model
Hypothetical electromagnetic field
Measured electromagnetic field
Statistical analysis
Results

Figure 2.12: Overview of DEMA and CEMA attacks.

The attack starts of with the recording of the electromagnetic field and input or output of the cryptographic device. Depending on if the input or the output is recorded, either the first or last round of the AES device can be attacked. Often in literature both are recorded, to verify the correctness of the device. Then a guess of a single expanded key byte, key hypothesis, will be made. This is repeated for every expanded key byte of a single round, which will give 16 × 256 = 4096 guesses for a full AES-128 round key. Both the input/output and key hypotheses are fed into the AES algorithm to calculate an hypothetical intermediate value after one or more AES round operations.

This hypothetical intermediate value will be input into the power model, which will generate a hypothetical electromagnetic field. This power model is explained in Section 2.3.2. Since the power is related to the electromagnetic field, this will also give an hypothetical electromagnetic field.

The last task is to compare the hypothetical electromagnetic field and the measured electromagnetic field with statistical analysis. Because of the key hypothesis there will be 256 hypothetical electromagnetic fields for each key byte. For each of these guesses a comparison need to be made with the measured electromagnetic signal. But since the measured electromagnetic signal isn't measured at a single time but over the whole AES operation, the results will be per key-byte. This will result in 256 key-byte guesses over time for a single key-byte.

Since the *SubBytes* operation of the AES algorithm is a non-linear function, the incorrect key-bytes will generate a normal distribution. If there is EM leakage from the IC there will be a correlation between the measured EM-radiation and the correct key-byte. Meaning that the correct key-byte will generate a peak in the statistical analysis at the point in time where the operation is executed. This is caused by P_{data}, which will be containing the intermediate result at that time. When enough traces are analysed this peak will be higher than the normal distribution of the incorrect key-bytes. This makes it possible to retrieve a key-byte, based on the measured EM-radiation and can be repeated for all other key-bytes until the full key is retrieved.

TVLA

In Gilbert Goodwill et al. [19], Test Vector Leakage Assessment (TVLA) is introduced. DEMA and CEMA attacks require multiple traces and do not give a good quantitative measurement if for example the power leakage model used is incorrect. The EM spectral intensity only gives information about the presence of certain frequencies which might not have a direct relation to the leakage. Therefore TVLA is introduced to give a fast quantitative measurement of the leakage of an IC, which requires less traces than DEMA and CEMA attacks[2].

TVLA is based around a t-test where two different groups of traces will be acquired and will be compared. For this thesis the first group is always based on random input. This input should always generate an uniform distribution for the measured EM radiation. The first group input is random to identify with the null-hypothesis if the two groups are statistically different. The input for the second group of the t-test is based on two intermediate values which will be called *MixColumns* and *SubBytes* in the next sections.

The *MixColumns* intermediate is defined as the output of the round 5 *MixColumns* operation of the AES-encryption. During the generation of the input data for the TVLA test, the Hamming Weight (HW) of the *MixColumns* intermediate will be controlled. For the second group of traces the *MixColumns* intermediate HW will be less than 8. The generated *MixColumns* intermediate value is generated randomly and based on the intermediate and the AES key the AES input will be calculated.

The *SubBytes* intermediate is defined as the output of the round 5 *SubBytes* operation of the AES-encryption. Similar to the *MixColumns* input data will be generated with a HW less than 8. In total this generates three different groups: *Random*, *MixColumns*, *SubBytes*. For the TVLA analysis only two groups are used, either *Random* and *MixColumns* or *Random* and *SubBytes*. During measurements the three different groups are randomly interleaved in order to reduce measurement errors due to external conditions.

The results of the TVLA test are computed using a Welch t-test, which is shown in Equation 2.5. Where $\bar{X_1}$, s_1 and N_1 are the mean, variance and number of traces of the *Random* group. The $\bar{X_2}$, s_2 and N_2 are respectively the mean, variance and number of traces of the second group. This second group is either the *MixColumns* or *SubBytes* intermediate. The t-test is calculated over time based on the measured EM radiation. But it can also be performed in the frequency domain after a Fast Fourier Transformation (FFT) transformation on the EM radiation.

$$t = \frac{\bar{X_1} - \bar{X_2}}{\sqrt{\frac{s_1^2}{N_1} + \frac{s_2^2}{N_2}}} \tag{2.5}$$

The pass fail criteria introduced by Gilbert Goodwill et al. [19] is $|t| > 4.5$. If $|t| > 4.5$ the null hypothesis does not hold with a confidence > 99.99%. Based on the selected second intermediate can thus be concluded that either the *MixColumns* or *SubBytes* intermediate is leaking. Based on this results analysis can be done on the leakage position and frequencies of the IC.

3

State of the art

The state of the art research is separated in two parts: DEMA attacks and Smart triggering. We will first describe the interesting DEMA attacks related to far-field DEMA attacks against AES using SDR. We will mainly focus on the works related to either far-field DEMA or DEMA using SDR. Next the research related to smart triggering is described related to DEMA attacks, and as last we will conclude by explaining the interesting research areas.

3.1. DEMA attacks

There has already been done a lot of research in the field of DEMA or DPA attacks against FPGA implementations of the AES cipher[34, 41]. But the field of far-field DEMA attacks is very small and only a small amount of researchers have approached this field[21, 45]. The use of Software Defined Radio, SDR, for side-channel attacks against AES is researched by Montminy et al. [35], but was discovered by Jun and Kenworthy [20]. In this section different attacks related to the far-field DEMA attacks will be discussed and also DEMA attacks using SDR.

3.1.1. DEMA attacks against ARIA in the far-field

The first paper describing far-field DEMA attacks is from Kim et al. [21]. In his paper the differential side-channel resistance of the Academia, Research Institute and Agency (ARIA) cipher is analysed in various ways. The ARIA cipher is a symmetric block cipher used as the national Korean standard algorithm. The ARIA cipher uses similar round operations, but differs in the S-Box operation which uses two different S-Box tables instead of a single S-Box table. These S-Box tables are alternated in the different rounds of the algorithm, where one of these S-Boxes is the Rijndael S-Box. The attacks are performed against different hardware implementation of the ARIA cipher running on an ARIA, Altera EP20K300EQC240-3. In the paper also masked FPGA implementations are tested, but only the non-masked implementations will be described. The measurement set-up is shown in Figure 3.1.

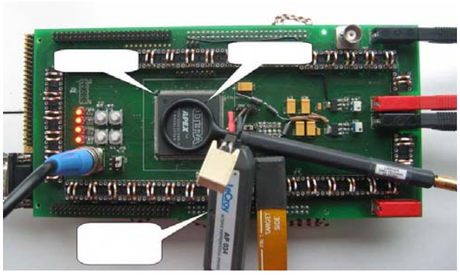

Figure 3.1: Measurement setup of Kim et al. [21], using an oscilloscope and Langer probe to measure both the Power and EM-radiation at the same time

Two different implementations are constructed for the FPGA, where the first implementation makes use of a table lookup in Read-only Memory (ROM). The second S-Box table of this implementation isn't stored in ROM, but calculated by inverting the first S-Box table. The second implementation is based on a multiplicative inverter and calculates each S-Box transition. Four S-Box calculations are computed in a single cycle, resulting in 4 cycles for the substitution layer per round.

To be able to correctly align the traces during the test a separate trigger pin is connected to the device's I/O, which can be seen in Figure 3.1 on the left. They state that in practice it is difficult to obtain a trigger signal without connecting to the FPGA board. But it should still be possible to trigger with the use of a smart trigger, which triggers at specific patterns of the far-field electromagnetic signal. If this method doesn't provide good enough results an attacker could also apply alignment techniques and discard outliers to perform far-field DEMA attacks.

$$N = 3 + 8\left(\frac{Z_a}{\ln\frac{1+\rho_{max}}{1-\rho_{max}}}\right)^2 \tag{3.1}$$

The results of the DPA and DEMA attacks can be seen in Table 3.1 for the two different FPGA implementations. The ρ_{max} is the maximum correlation of the correct key byte using a Hamming Distance model. The minimum amount of traces needed to be able to retrieve the key is t_{min} and is calculated by Equation 3.1 from Mangard [29]. The probability a determines the confidence level and for the results a confidence level of $a = 0.9999$. The confidence level is determined by a rule of thumb from Mangard et al. [30]. This will result in $Z_a = 3.719$.

	DPA		DEMA	
	Inverter	Table	Inverter	Table
ρ_{max}	0.1521	0.2222	0.2289	0.3068
t_{min}	1180	545	512	278

Table 3.1: DPA and Near-Field DEMA results

The results show that for both DPA and DEMA the table implementation is easier to attack. This can be explained by the fact that the inverter is composed of more complicated combinatorial logic which will result in more noise during the measurements. This means that the signal-to-noise, SNR, of the table implementation is significantly higher and resulting in a higher ρ_{max} and lower t_{min}. Also can be seen that the DEMA attack performs better than the DPA attack.

In the far-field instead of four S-Boxes only a single S-Box is calculated per clock cycle in order to improve the Signal to Noise Ratio (SNR). The equivalent ρ_{max} of DPA and DEMA with this adapted implementation are 0.6 and 0.8 respectively. The measurements are performed in a non-shielded environment with a directional antenna, ranging from 200 MHz to 1 GHz. The signal of the antenna is amplified by 30 dB and connected to an oscilloscope. Note that the signal isn't filtered, and will contain a wide range of frequencies mainly from the antenna range. Measurements were made at 0.5 m and 1 m distance and can be found in Table 3.2.

	Far-field DEMA	
	0.5 m	1 m
ρ_{max}	0.0461	0.0324
t_{min}	13000	26335

Table 3.2: Far-Field DEMA results

The far-field results are very low compared to the DEMA and DPA attacks. This is due to the high amount of noise, which will cause a lower SNR. Even though the results are worse it is still possible to retrieve the secret key from a distance of 1 meter. The SNR can be improved by either using a spectrum analyser to remove frequencies with a low data dependency or to perform the attacks in a shielded environment.

The work from Kim et al. [21] attacks the ARIA cipher, which is mainly used in Korea. While this cipher has many similarities to the more commonly used AES encryption the results of these attacks could be totally different. In order to determine how the leakage compares to other ciphers, recordings should have been made with different ciphers on the same platform. Due to the different behaviours of power within microprocessors, it is impossible to conclude how the side-channel leakage of ARIA compares to AES from this work.

The calculations of Mangard [29] are used to calculate the minimum amount of traces needed in order to retrieve the key. The paper doesn't describe that they performed attacks which recovered the full secret key. There is only a theoretical assumption that the key could be fully retrieved.

Recordings were made in a non-shielded environment which would make these attacks more plausible in real-world scenarios. This means that no access to the target device or close by environment is needed in order to shield the device from the noise. But in the attacks a trigger is used from a GPIO pin in order to identify and align the traces. This means that, although the tests are performed in a non-shielded environment, access to the device is still needed to perform the attack.

3.1.2. Near-Field DEMA against AES using SDR

In Montminy et al. [35], near-field DEMA attacks are shown against an AES-128 cipher running on an ARM Cortex-M4F at 50 MHz. Although they have made no approach to attack the far-field emanations of the cryptographic device, this research is still interesting since they analyse the use of Software Defined Radio (SDR) for DEMA attacks against AES.

In Jun and Kenworthy [20] attacks against implementations of RSA and Eliptic Curve Cryptography (ECC) are shown on smart phones. They made use of SDR as digitizer and a receiver demodulated the signal. It was shown that AES operations can be observed in the demodulated signal, but an attack was never performed against AES. Next to Montminy et al. [35] no other research related to using SDR for side-channel analysis have been found.

In Montminy et al. [35] two different set-ups are used for comparison, the baseline set-up uses an oscilloscope and the main set-up uses an SDR for recording the traces. The analog front-end is similar for both set-ups, containing of a low-sensitive probe from Riscure[48] with a differential amplifier of 60 dB gain and a low-pass filter of 1000 MHz. The SDR set-up uses an extra 20 dB attenuator, because the amplifier is build in the Riscure probe. The set-up makes use of a different collection and control PC in order to fully imitate a real attack. The SDR set-up can be seen in Figure 3.2.

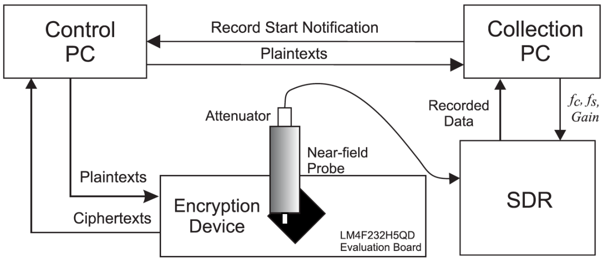

Figure 3.2: Measurement setup of Montminy et al. [35], using SDR and a remote control PC

Two different SDRs are used and analysed for CEMA side-channel. The USRP2[15] can sample at f_s = 100 MSa/s with 14-bit ADC resolution and the LRFX daughter board is used which is designed to receive signal between 0-30 MHz. It can receive signals up to 50 MHz, but the LRFX uses a third-order low-pass filter with a cut-off at 30 MHz to prevent aliasing. With the USRP2, the frequencies between 18 an 40 MHz are recorded with an output sampling rate of f_s^D = 2 MSa/s and f_s^D = 4 MSa/s.

The second SDR is a Low-cost Digital Video Broadcasting Terrestrial (DVB-T) USB dongle which can be used as SDR. The DVB-T is based on the Realtek RTL2832U[37], which has an 8-bit ADC.The Elonics E4000 is chosen because it is capable of receiving signal between 50 MHz - 2.2 GHz, outside its specification[49]. Center frequencies between 53.5 and 73 MHz and a fixed gain of 1.5 is used during all collections. As sample rate 2.8 MHz is chosen to avoid sample loss when the data is transferred of USB 2.0 to the PC.

The oscilloscope recordings are made at a frequency of f_s = 2.5 GSa/s. To be able to determine the best location for evaluating the attacks an XY-scan is performed on top of the target using the oscilloscope set-up. In Figure 3.3a the Power Spectral Density (PSD) across the chip at 625 locations is shown. But in order to verify that the PSD relates to the actual leakage also the mean magnitude of the maximum correlation coefficient, based on the hamming weight, for the 16 different key bytes are calculated at the same 625 positions. For each of the positions 1000 traces were made and the results can be seen in Figure 3.3b.

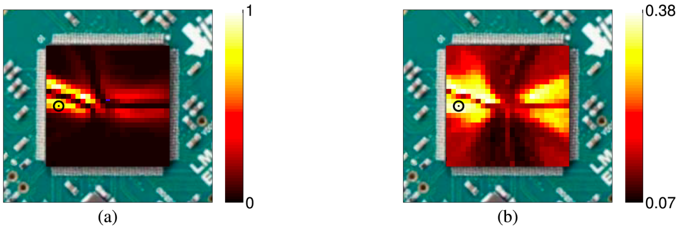

(a) Maximum PSD between 49 and 50 MHz.

(b) Mean magnitude of the maximum correlation coefficient for each of the 16 key bytes using the Hamming Weight with 1000 traces per location.

Figure 3.3: XY-scans using an oscilloscope with a spatial resolution of 25x25. The ⊙ indicates the chosen spot.

From the results from the XY-scans a position is chosen for the subsequent collections, where the leakage was sufficiently high. This spot is indicated with ⊙. On this spot first a baseline measurement is made to determine the exploitable frequencies and sample rate containing the AES secret key information for the SDR. These baseline measurements are performed using the oscilloscope set-up and because earlier experiments showed more leakages below 100 MHz a Chebyshev Type I filter with a cut-off of 100 MHz is added. The traces are then down sampled to a sampling frequency of 250 MSa/s in order to reduce the time required to filter the traces. To simulate the SDR behaviour a bandwidth of $W_{BW} = 2$ MHz was chosen for the filter. The center frequencies that were analysed are $f_c \in \{1, 2, ..., 99\}$, resulting in 99 different filters.

To be able to compare the effectiveness, a confidence is calculated based on the two highest correlating coefficients of a single key byte. A hypothesis test is performed using Fisher's transformation and a Z-test[17]. The results of the baseline measurement, containing of 1000 traces, can be seen in Figure 3.4.

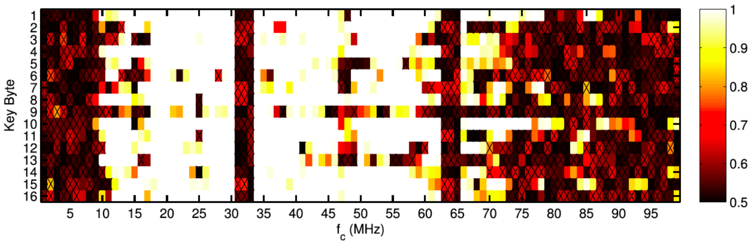

Figure 3.4: Baseline results using 1000 traces filtered with overlapping intervals. The boxes are centred at the correct f_c, but due to the overlap the width isn't corresponding to the correct bandwidth. Attacks with an incorrect key-byte are marked with a ×.

In the results can be seen that the confidence changes for both the frequency interval and the key byte. The confidence of some of the key bytes is high and correct, while some key bytes are much harder to attack using the same set of filtered traces. In Table 3.3, the attack can be seen against the unfiltered traces. Every key byte is correctly guessed, when the unfiltered traces are used.

The filtered traces between the frequency intervals of 12 MHz and 60 MHz, contain most information to extract all 16 bytes of the secret key with high confidence. But at 31.9 MHz and 63.7 MHz the confidence is lower and all key bytes are incorrect. At those frequencies signals were detected which are unrelated to the encryption operations.

Compared to the recordings with an oscilloscope the SDR can record traces in real-time. The SDR will

Key Byte	Confidence	Key Byte	Confidence	Key Byte	Confidence	Key Byte	Confidence
1	1.00000	5	0.99998	9	0.98726	13	0.99996
2	1.00000	6	0.99990	10	1.00000	14	1.00000
3	1.00000	7	0.99970	11	1.00000	15	1.00000
4	1.00000	8	1.00000	12	0.99979	16	1.00000

Table 3.3: Confidence using the 1000 unfiltered traces, decimated at $f_s = 250$ MSa/s.

record multiple traces at ones, called a collection. But in order to use the SDR without trigger, each individual encryption operation must be identified via signal processing. In order to simplify the task the PC will perform groups of $n_g = 250$ encryptions using a random plain text and fixed key. Each collection will contain n_g traces, and if the algorithm can't find that exact amount of traces the whole collection is discarded.

The ARM Cortex-M4F is only performing encryption operations when executed by the Control PC, which will result a full overview trace which can be seen in Figure 3.5a. A simple peak counting algorithm, with a minimum distance shown in 3.5a, is implemented in order to detect n_g traces within the collection. The amount of peaks found by the algorithm is defined as n_d. The algorithm starts with h_{min} initiated as the overall maximum value of the recorded collection. Then h_{min} is gradually lowered while $n_d < n_g$. If $n_d > n_g$, the whole collection is discarded. Else if $n_d = n_g$, the collection will be split according to those peaks into n_g separate traces.

The first trace will be used as a reference which is then cross-correlated to all other $n_g - 1$ traces to improve the alignment. These traces are then circularly shifted to correct the alignment between the different traces. The results of this alignment can be seen in Figure 3.5b, where 250 traces are superimposed. Multiple collections of n_g traces are collected until the total amount of traces exceeds the desired amount of traces, n_{total}.

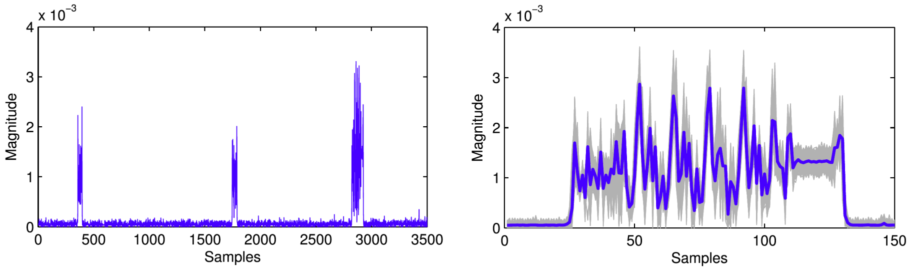

(a) The magnitude recorded includes the interrupt handling, command identification, receipt of the plain text, encryption of the plain text and returning the cipher text for a single AES encryption.

(b) Overlay of 1250 traces with the dark line showing the mean.

Figure 3.5: Traces and a magnitude recording by the USRP2 with $f_c = 22$ MHz and $f_s^D = 2$ MSa/s.

Since the RTL-SDR is a CMOS RF tuner with an automatic Direct Current (DC) offset compensator, some additional processing is needed. The RTL-SDR will apply different DC offsets for each new collection recorded. In order to compensate for this effect the mean across the whole collection is adjusted to zero. Because the RTL-SDR also sometimes overflows and doesn't report it, cross-correlations of the width of the encryption part are calculated. Traces with a standard deviation grater than 5 from the mean of the width are discarded from the trace sets.

In Figure 3.6 the results of the USRP2 with $f_s^D = 2$ MSa/s are shown and respectively in Figure 3.7 the results of the USRP2 with $f_s^D = 4$ MSa/s. Compared to the baseline measurements both results show similar effects of key bytes leaking in different frequencies. For the measurements two attacks were able to retrieve the full key with high confidence. In Figure 3.6c this can be seen at $f_c = 27$ MHz and in

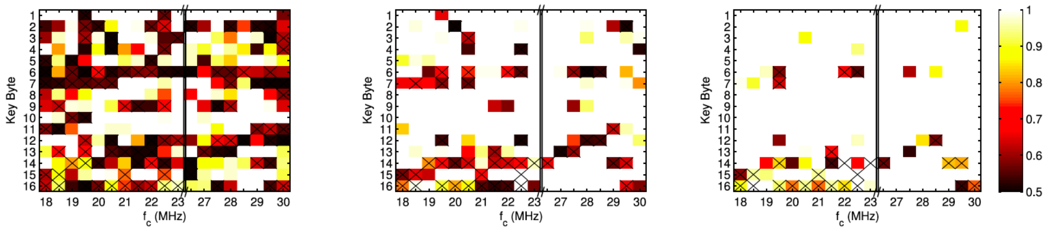

(a) $n_t = 5000$

(b) $n_t = 25000$

(c) $n_t = 100000$

Figure 3.6: CEMA confidence results of the USRP2 using $f_s^D = 2$ MSa/s. CEMA attacks with incorrect key-bytes are marked with a ×.

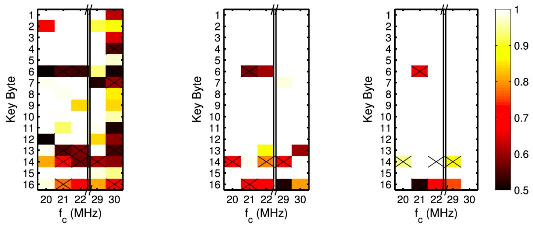

(a) $n_t = 5000$

(b) $n_t = 25000$

(c) $n_t = 100000$

Figure 3.7: CEMA confidence results of the USRP2 using $f_s^D = 4$ MSa/s. CEMA attacks with incorrect key-bytes are marked with a ×.

3.7c this can be seen at $F_c = 30$ MHz. This shows that it is possible to retrieve a full AES key using SDR with 10000 traces. With less traces it will become harder to retrieve the full key, and several different measurements are needed to retrieve the full key.

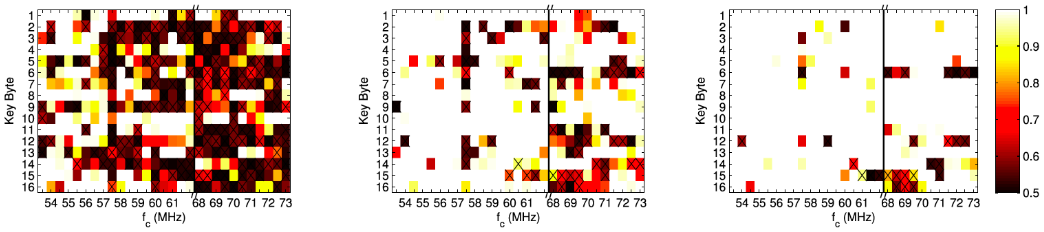

(a) $n_t = 5000$

(b) $n_t = 25000$

(c) $n_t = 100000$

Figure 3.8: CEMA confidence results of the RTL-SDR using $f_s^D = 2$ MSa/s. CEMA attacks with incorrect key-bytes are marked with a ×.

Similar attacks were performed using the RTL-SDR, and those results can be seen in Figure 3.8. Also these results are consistent with the baseline test. An unknown signal varies near 56.9 MHz which causes a reduction in attack with $f_c = 55$ MHz. With 25000 and 100000 traces there exist multiple frequencies at which the full key can be retrieved.

The baseline can't be directly compared to the SDR, because the baseline is not down-converted, low-pass filtered or decimated. The baseline has a total of 8450 samples for the whole encryption at $f_s = 250$ MSa/s. While the lower sampling rates of the SDR, 2 MSa/s and 4MSa/s only record 63 and 129 samples respectivly for the whole encryption. This means that multiple clock cycles are included in a single SDR sample. Comparing the baseline and SDR results, several key-bytes have a higher confidence using the SDR while others have a higher confidence in the baseline result.

The advantage of the SDR is that it can be used to collect trigger-free. While oscilloscopes without a trigger need a large amount of memory, to be able to capture multiple traces in a single collection. A high-end scope is needed to be able to achieve this, which are very expensive compared to the different SDRs. There are also scopes like the PicoScope which can capture in real-time, but it is likely that data aggregation will degrade the utility of the collected collections.

Additionally Montminy et al. [35] shows that compiling the code with different optimisation levels lead to different key-byte leakages. But in order to determine why key-bytes leak at different frequencies further research is needed. This could be used to optimize the code to reduce side-channel leakage. It would be extremely useful to compare the results to other targets, to verify that the key-byte leakage per frequency isn't restricted to the ARM Cortex-M4F.

The attacks in Montminy et al. [35] are extremely realistic since the need for an additional trigger is removed. Also removing the expensive oscilloscope costs makes the attack more accessible. But since they made use of a near-field probe this attack can be classified as a semi-invasive attack. Meaning that the device enclosure needs to be removed in order to measure on the surface of the chip. This can be detected by anti-tampering devices resulting in the removal of the secret key.

3.1.3. Far-field DEMA against AES using SDR

At Fox-IT already a lot of research has been done in the field of Far-Field DEMA attacks against FPGA's[45]. The main focus of the research at Fox-IT was to perform far-field DEMA attacks against devices running AES encryption using SDR. These DEMA attacks are performed on devices running a T-Tables implementation of the AES-256 algorithm from The OpenSSL Project [54]. During this section AES-256 is going to be refered as AES. The two targeted devices are a SF2 Starter Kit[12] from Microsemi and a PYNQ-Z1[9] from Xilinx. The attacks are performed in cooperation with Riscure[48].

Both targets are running the AES encryption in the ARM-core in software. A trigger is added to both devices with the help of a GPIO pin. This pin is switched at the center frequency of the SDR for a small amount of time just before the AES encryption. This will cause a spike in the SDR recording, which can easily be filtered to generate separate traces.

The encryption on the devices is either done on-line through connecting an Ethernet cable, or fully off-line feeding the output of the encryption as input for the next encryption. The on-line version makes use of a simple UDP protocol which can set the key during initialization. It encrypts the bytes send to the device and returns the encrypted bytes to the computer. The off-line set-up has a pre-programmed key and initialization input for the encryption. It performs 50000 encryptions and then resets the input to the pre-programmed value. During this reset an extra delay is added, which can later be used to find the start of the algorithm. All measurements are performed using the on-line method, unless stated otherwise.

Three different measurement set-ups were used consisting of different SDRs: SR-7100, USRP B200[15] and RTL-SDR[37]. The SR-7100 is a military grade SDR with a bandwidth of 500 MHz. The USRP in comparison only has a bandwidth of 56 MHz and the RTL-SDR a bandwidth of 2.4 MHz. An overview of how the work flow of the different attacks can be seen in 3.9.

The analogue measurement part consists of an antenna, filtering and amplification. The antennas used during the research in the near-field are a low sensitive Riscure probe[48] and a self-constructed loop antenna out of a shielded wire(see Figure 3.10). For the far-field measurements a directional log-periodic antenna is used with a frequency range of 400 till 1000 MHz and a yagi antenna focused at 144 till 146 MHz. Depending on the target different low-pass, high-pass and band-pass filters are added in order to filter out unwanted signals. For both the self-constructed loop and the far-field antennas amplifiers are added.

The Radio Recording consists of the different SDRs which will record the IQ samples. The recording is started before the device is instructed to encrypt and stopped after a predefined amount of encryptions are executed. This means that instead of separate traces a single recording is made of multiple encryptions, called collections.

These collections are then preprocessed in Matlab[32] to generate separate traces. In order to separate these traces at first the collection is filtered with a low-pass filter. After that a simple threshold is defined by the user, to separate the AES encryptions. This works due to the fact that the idle EM radiation is

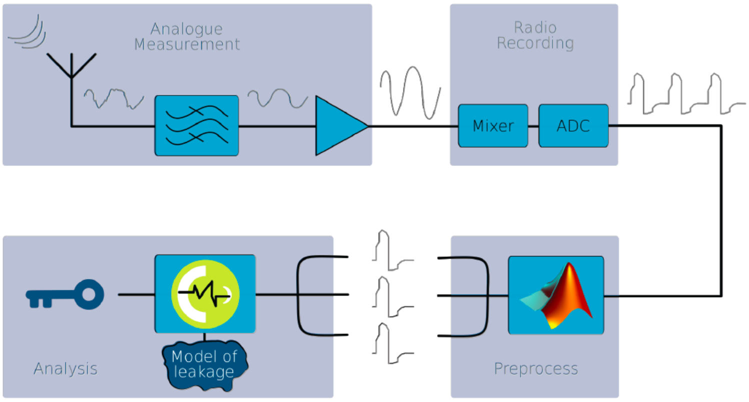

Figure 3.9: Overview of the attack work flow used at Fox-IT.

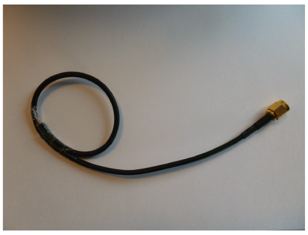

Figure 3.10: Self constructed loop antenna out of wire

lower than the EM radiation of the trigger. The second crossing with the threshold is used to define the start of the trace. An example of such a trace can be seen in 3.11.

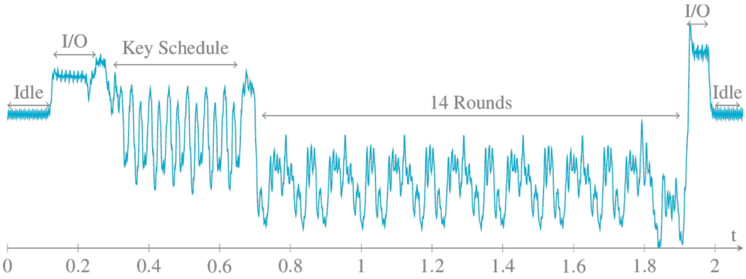

Figure 3.11: An overview trace recorded using the SR-7100 showing the different phases of the AES encryption.

The last step is the analysis process which is performed in Riscure's Inspector[48] Side-channel Analysis (SCA) tooling. This software can perform the DEMA and CEMA attacks performed in this research. Different leakage models can be defined and tested using the same input data. With a successful attack Inspector should give the correct key-bytes with the highest correlation. The leakage model used during all measurements is the Hamming Distance against the S-Box in the first round.
To evaluate the leakage of the target devices at first recordings were made using the SR-7100. Because of the high bandwidth the leakage frequencies could be easily verified. These recordings were made for the the SmartFusion2 and the Pynq using a near-field probe. The recordings from the SmartFusion2 showed leakage at 142 MHz, which is the exact clock frequency of the ARM core. The Pynq didn't have any leakage at the ARM core frequency but instead it leaked at the memory clock frequency of 400 MHz. This can be explained by the fact that the Pynq makes use of external DDR memory, while the SmartFusion2 has internal memory. Because IO operations use a lot of power and the traces of the memory are going through the Printed Circuit Board (PCB) these signals emit a lot of EM radiation.

For both platforms also measurements were made using the small loop antenna. Those showed high correlations and with only several seconds of recordings made using the RTL-SDR. For both devices the correct key could be extracted from 5 centimetres distance. But the recordings of the SmartFusion2 showed a very low SNR, because the leakage frequency was very close to several Frequency Modulated (FM) radio stations. Therefore further tests are only performed on the Pynq.

In order to analyse the far-field capabilities of the antennas the loop decay are measured for both the self constructed loop antenna and the log-periodic antenna. The results can be seen in Figure 3.12. The results show that the loop antenna decay follows a $\frac{1}{x}$, while the log-periodic antenna follows a more linear trend. Also the signal strength received by the log-periodic antenna is much higher and the antenna is also more directional. Therefore the log-periodic antenna is used during far-field measurements.

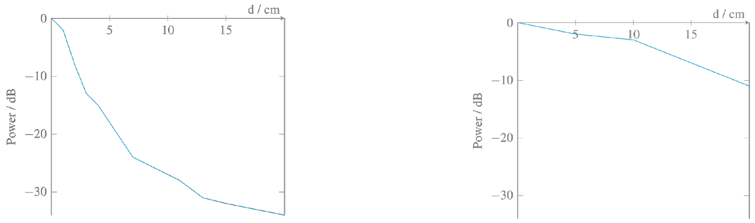

(a) The decay of the self constructed loop antenna.

(b) The decay of the log-periodic antenna.

Figure 3.12: Evaluation of the near-field and far-field antenna range for the different antennas.

Because Fox-IT didn't have any filtering available around 400 MHz a simple Faraday construction was made to shield from EM radiation. This shielded box was made using emergency blankets as can be seen in Figure 3.13a. Inside this box both the target device and the antenna were placed at a distance of 30 cm from each other. The paper in Figure 3.13b shows the distance between the target and the antenna. With a 50 second recording the key was successfully extracted from the device with the RTL-SDR. The size of the shielded box was limiting measurements from further distances.

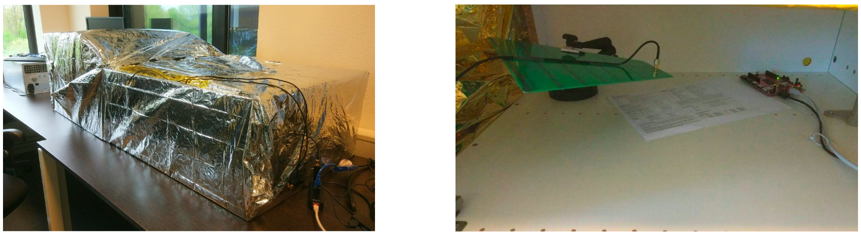

(a) Outside of the shielded box.

(b) The attack set-up inside the shielded box.

Figure 3.13: A self-made Faraday cage to shield from EM radiations out of emergency blankets.

At OSPL[42] more measurements were made at a distance of 1 meter inside an anechoic chamber. This anechoic provides an environment with almost no external electromagnetic radiation, resulting in high SNR. During these measurements a discone antenna was used, which was provided by OSPL. The discone antenna measures omnidirectional and has an extremely wide frequency range. In order to ensure a full isolation between the target and the measurement set-up a USB battery is used for the target and a separate battery is used for the amplifiers. Also instead of the on-line implementation, the off-line version is used. The antenna is converted to an optical signal and recordings outside the room are made with the RTL-SDR. The set-up in the anechoic chamber can be seen in Figure 3.14.

The results from the anechoic chamber shows that it is possible to retrieve the key successfully at a distance of 1 meter. In this ideal environment a recording of 5 minutes was needed to retrieve this key. Although also measurements were made with the SmartFusion2, attacking this target at larger distances wasn't possible.

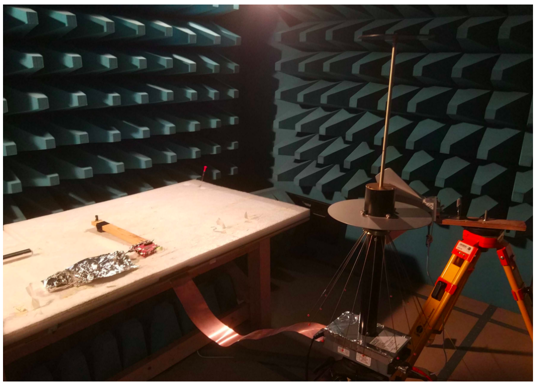

Figure 3.14: The set-up in the anechoic chamber using a discone antenna and batteries.

The research at Fox-IT shows attacks using the Hamming Distance of the S-Box in the first round. In a T-Tables implementation this leakage is most often not giving the best results. Due to the combination of the *SubBytes* and *MixColumns* in a single lookup, intermediate values after the *SubBytes* operation are expected to correlate less. Instead the last round should be attacked with the encrypted cipher text as input, as the last round in AES omits the *MixColumns* step. This would result in a normal S-Box leakage and results could have been improved.

Also in this work a trigger had been implemented to identify the separate AES encryptions and generate the traces. In real-world applications this is often not available, making these attacks less realistic. Also most measurements were made in a shielded environment, which means that access to the device and it's environment are needed to perform the attacks.

Instead of AES-128, which is often used for side-channel analysis research, the research at Fox-IT was using AES-256. But in side-channel analysis the amount of rounds and the length of the key aren't decreasing the effectiveness of the attack. The only difference between the attack is that instead of a single round consisting of 16 key bytes, two rounds need to be attacked. This results in a longer computing time for the statistical analysis, but doesn't affect the measurement time.

3.2. Smart triggering

In order to achieve a full non-invasive DEMA or CEMA attack, the trigger needs to be replaced with something else. The solution for this problem can be achieved by so called smart triggering[55]. A smart trigger makes use of the recorded samples and a pattern to genarate a trigger signal. The most commonly used in literature is the Sum of the Absolute Difference (SAD)[35, 39, 40]. In Beckers et al. [3], a comparison of three different algorithms are made and the interval method is evaluated.

For each of the three different pattern matching smart triggers, a reference signal g of N samples is defined. The continuously incoming signal is defined as h. The algorithms will calculate a value the represents the correspondence between g and (a part of) h. This value is defined as a score $T(k)$, where k is the time shift from the starting point of h. Since oscilloscopes and SDRs make use of ADC conversions these calculations are all done in the discrete time domain.

The research of Beckers et al. [3] is focussed on finding algorithms for low-latency matching which can be efficiently implemented in hardware. This is needed for the oscilloscope, because it has limited a memory.

3.2.1. Cross-correlation

The cross-correlation makes use of the product of the samples to measure the resemblance. It is a commonly used algorithm in statistical analysis and in Equation 3.2 the calculation for this smart trigger can be seen.

$$T(k) = \sum_{m=1}^{N} g(m)h(m+k) \tag{3.2}$$

In Beckers et al. [3], cross-correlation is discarded as a solution due to its use of multiplications. The amount of multipliers available in a platform would limit the length of the reference signal g. Also the score calculation can grow significantly with N, which would require large adders.
When looking at SDR most these effects can be neglected, as triggers will be calculated after the full recording was made. This makes it possible to use a high-performance computer to perform these calculations. Since cross-correlation makes use of difficult operations and is still scaled by N, it would result in a slower preprocessing phase.

3.2.2. Sign comparison

This method makes use of the mean μ of g, which can be calculated after the reference signal is set. The reference signal is transformed into a binary sequence g'. Where $g'(m) = 1$ if $g(m) > \mu$ and else $g'(m) = 0$. A similar transformation is performed on h resulting into h'. In Equation 3.3 the calculation of the score can be seen. In this equation $g'(m)h'(m+k)$ relates to the fact that $g'(m) == h'(m+k)$.

$$T(k) = \sum_{m=1}^{N} g'(m)h'(m+k) \tag{3.3}$$

The results from Beckers et al. [3], show that the algorithm has an unreliable trigger behaviour when used in noisy measurements recorded by an oscilloscope. Therefore this method is also discarded.
For SDR the noise behaviour is similar to an oscilloscope, as both methods are measuring the same signal. Although SDR makes use of a down-converted signal, this method would probably also generate a unreliable trigger for SDR.

3.2.3. SAD

The Sum of the Absolute Difference performs a sample-wise subtraction between g and h. The score is than calculated as the sum of absolute value of this difference. In Equation 3.4 this calculation is shown.

$$T(k) = \sum_{m=1}^{N} |g(m) - h(m+k)| \tag{3.4}$$

SAD is used on icWaves[47] device from Riscure, which is used by multiple researchers[35, 39, 40]. This is due to the effect that the SAD can be efficiently implemented on an FPGA, resulting in a minimum delay between the detection of the pattern, and the generated smart trigger.
In Beckers et al. [3] is noted that a single sample with a large difference between g and h could have a significant impact on the result of $T(k)$. Next to that the flexibility compared to Interval matching is less, discarding also this option.

3.2.4. Interval matching

This algorithm makes use of an interval defined by an *offset*, o, above and below the reference trace. The score is also based on a binary sequence, which checks if the signal lies within the interval. For each sample $IM(m,k)$ is calculated as in Equation 3.5

$$IM(m,k) = \begin{cases} 1 & \text{if } g(m) + o \geq h(m+k) \geq g(m) - o \\ 0 & \text{otherwise} \end{cases} \tag{3.5}$$

Then using this equation the score is calculated according to Equation 3.6.

$$T(k) = \sum_{m=1}^{N} IM(m,k) \tag{3.6}$$

In Beckers et al. [3] the decision was made to implement the Interval matching algorithm in their waveform-matching trigger system. They state that the algorithm can be efficiently implemented, which could potentially result in a low latency trigger. Compared to SAD it has an extra degree of freedom by defining the *offset*, giving it more flexibility. It also is better resistant against sample outliers, which have a significant impact on the score of SAD.

3.3. Conclusion

It has been shown by Kim et al. [21] and Ramsay [45] that far-field DEMA attacks are possible against symmetric ciphers. The attack from Kim et al. [21] could be categorised as a passive semi-invasive side-channel attack. In Kim et al. [21] a trigger is added to the device by adding a GPIO to the oscilloscope. In a real-world situation this would often mean that the device enclosure of the target needs to be removed. Although they state that it should be possible to make the attack fully non-invasive by adding a smart trigger, this hasn't been researched in the far-field.

In Ramsay [45] a fully passive non-invasive side-channel attack against AES is shown. But in order to achieve this fully non-invasive attack a trigger was added by emitting a signal in the recorded frequencies. These are then used as a threshold and used to align the different traces. Since these triggers are often not present in real-world situations, this attack could also be considered as semi-invasive.

In Montminy et al. [35] a semi-invasive near-field DEMA attack is shown. Due to the high SNR when using the near-field a simple threshold can be used to identify single traces. Because the removal of the device enclosure is often needed to be able to record the near-field, these attacks are categorised as semi-invasive. The target is also programmed to perform only the AES-encryption and go to idle in the meantime. In several real-world platforms this is not realistic, as most platforms perform multiple tasks.

Also in Montminy et al. [35] observations are made that key bytes are leaking at different frequencies. This effect can't be explained yet. In order to verify if this effect is platform specific, measurements should be made on different platforms. This effect is even occurring when the AES algorithm is compiled using different optimisation settings. This could also easily be verified on other platforms.

Smart triggering used for side-channel analysis is often aimed at low latency[3]. This is done in order to be able to trigger an oscilloscope. Because recordings made with SDR can record in real-time, this triggering can be calculated after the measurements are made. In real-world situations a fast computer can be used to calculate those triggers. This opens up new possibilities for smart triggering algorithms, which might achieve better results. When approaching the far-field, the SNR becomes lower and detecting a trigger becomes harder. Simple triggering methods like SAD and interval matching could be compared to cross-correlation, in order to improve the results.

The research related to DEMA attacks shows that it is extremely hard to perform passive non-invasive attacks. A full non-invasive DEMA attack hasn't been achieved without several compromises like requiring a trigger. Combining the research of SDR DEMA attacks and smart-triggers isn't researched and could potentially lead to a full non-invasive side channel attack. In order to make this attack fully non-invasive this attack should be performed in the far field, to prevent the removal of the device enclosure.

4
Measurement set-ups

During this thesis two different measurement set-ups are used, the baseline and SDR set-up. The baseline set-up is used to identify the leakage and as a comparison to the SDR set-up. The main difference between the two set-ups is the recording device and both set-ups are described in detail in this section. For each of these measurements set-ups different analog front-ends are used in order to assure a good leakage assessment of the different devices under attack. The differences and decisions about the analog front-ends are described in this section. As last we will describe and discuss the analysis software and the different target platforms.

4.1. Baseline

The baseline set-up makes use of a LeCroy HDO6104 oscilloscope. This oscilloscope has a bandwidth of 1.25 GHz can record up to 10 GSa/s with a 12-bit resolution. Because the devices' leakage frequencies are below 500 MHz and to speed up the recording process for most recordings the oscilloscope was configured at 2.5 GSa/s at 8-bit resolution. This means that frequencies lower than 1.25 GHz can easily be reconstructed according to Nyquist [38]. The lower 8-bit resolution was chosen in order to minimize the recorded data and to accelerate the required preprocessing and analysis phase.

All baseline recordings make use of a trigger signal implemented in the device under attack. A GPIO trigger pin is set to high before the first round of the AES encryption. After the last round of the encryption is finished the trigger pin is set to low. This way we are able to verify that we have recorded the full AES encryption and make sure that all recordings are aligned at the first AES round.

Some baseline recordings make use of an XY-table, which can move an electromagnetic probe at a precision of less than 50 µm to a specific spot above a chip. With this tool automated recordings can be made at multiple positions above the target device. These recordings give a good representation on which particular position a device leaks, and how the positioning of the probe affects the leakage results.

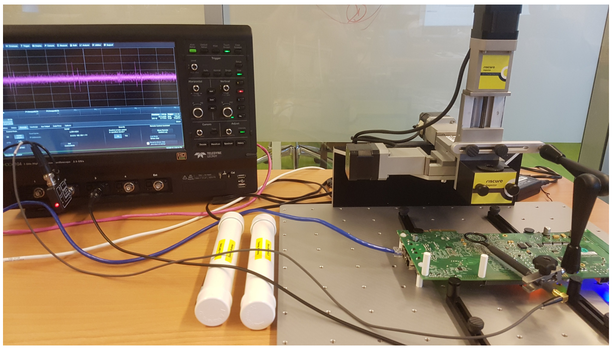

Figure 4.1: An example of the baseline set-up with the oscilloscope, XY-table, Microsemi SF2 Advanced and a Langer probe.

4.2. SDR

During the SDR set-up two different SDR's are used with different characteristics. The specifications and differences of the devices will be shortly discussed in this section. Because the SDR set-ups don't make use of a trigger input, pre-processing is required in order to split a single recording, collection, into separate traces. This pre-processing will be explained at the end of this section.

4.2.1. RTL-SDR

The NooElec [37] is one of the cheapest SDR's available, and was originally designed as DVB-T receiver. An overview of the internals of the RTL-SDR can be seen in Figure 4.2. It makes use of an RTL2832U demodulator which has a DSP inside.

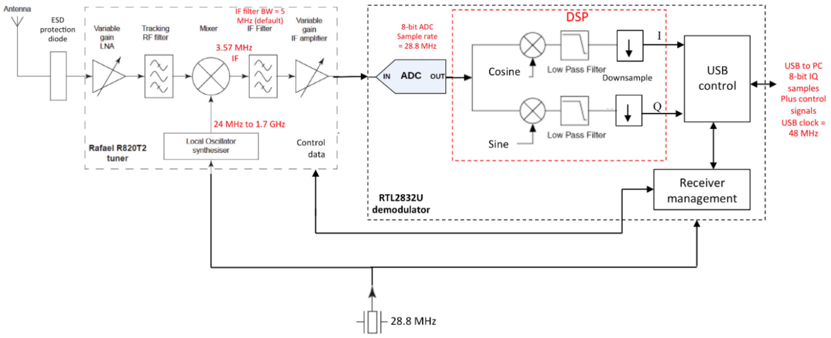

Figure 4.2: Simplified block-diagram of the RTL-SDR from Scher [50]

The RTL-SDR RF hardware can generate baseband frequencies from 24 MHz until 1.7 GHz. Two amplifiers, a Low Noise Amplifier (LNA) and a Intermediate Frequency (IF) Amplifier, with a variable gain will amplify the analog signal before the ADC. The IF is filtered with a bandwidth of 5 MHz.

In the RTL-SDR the I and Q samples are generated after the ADC conversion inside the DSP. The samples are then down-sampled to limit the data rate for the USB 2.0 connection. even though the ADC can sample at 28.8 MHz resulting in 14.4 MSa/s the maximum transfer rate the USB controller can achieve successfully is 2.8 MSa/s without any sample loss.

This device was chosen because it is relatively cheap, around 17 Euro, and has a wide range of baseband frequencies. The RTL-SDR is also widely available and in Montminy et al. [35] it is shown that it can perform side-channel attacks against AES.

4.2.2. USRP B200

The USRP B200 is made by Ettus Research [15]. The consists of the Spartan FPGA and AD 9364 RFIC analog front-end. An overview diagram of the USRP B200 can be seen in Figure 4.3.

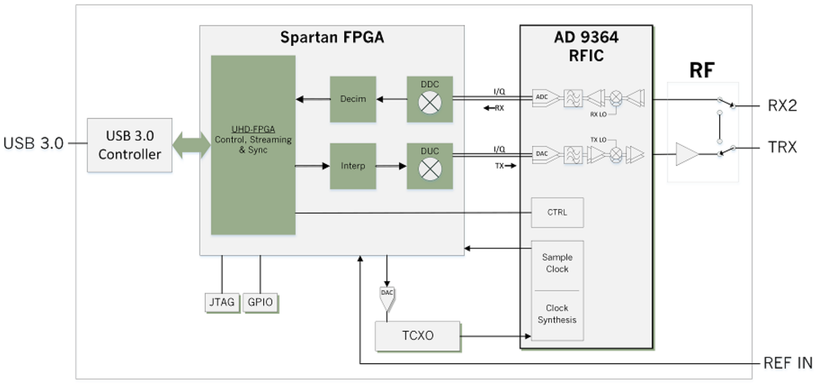

Figure 4.3: Simplified block-diagram of the USRP B200 (mini) from Research [46]

The main difference in the analog front-end, compared to the RTL-SDR, is that the USRP performs the phase shifting inside the AD 9364 RFIC. Also different baseband frequencies can be set, namely from 70 MHz until 6 GHz. Both devices can also record outside the defined baseband frequencies, but this could lead to a lower SNR.
Because of the Spartan FPGA and USB 3.0 the USRP can record at 61.44 MSa/s. Depending on the throughput of the USB-controller on the computer this speed can vary. Because of the higher sample rate and price of around 750 Euro the USRP was chosen. This makes it possible to record more samples per AES encryption which could improve the SNR.

4.2.3. Pre-processing

For the SDR set-ups pre-processing is required to separate a single large recording into multiple traces. This pre-processing is performed in MATLAB [32] using the raw I and Q samples from the SDR. The input and output data of the AES-encryptions during the recordings are imported in MATLAB in order to link the trace with the in- and output data. Since most recordings are large, more than 1 GB, chunks of 200 MB are processed in series for each of the pre-processing steps.
The first step of the pre-processing is to convert the I and Q samples to power using $P = \sqrt{I^2 + Q^2}$. In this step the the complex samples, with an 90 degrees phase shift, are converted and only the power of the signal remains. Since most leakage is inside the power, the signal does not need to be up-converted which saves a lot of computing power.
The recorded signal contains a lot of noise and a lot of oscillations at high frequencies. This makes it extremely hard to perform some sort of smart triggering without generating a lot of false positives. In order to improve the signal a moving average filter with a variable window size is implemented. The window size is chosen based on the sampling frequency and the different analog front-ends.
In order to split the large recording into multiple traces different techniques are implemented. Three different approaches are tested and compared, SAD, correlation and a simple trigger based on power. Each of these functions take the filtered signal as input and will output the positions of the found triggers. Each of the functions has different type of variables which will be described.
The SAD needs an input trace, which is selected from the large recording by hand. This means that the user must be able to visually analyse the recording and see where the AES-encryption is taking place. For the SAD also a maximum distance is required, which has to be given by the user as input. At first a high maximum distance needs to be chosen to assure all traces are found. Then the maximum distance needs to be tuned in order to detect exactly the amount of traces available in the recording.
For the correlation based smart-trigger a similar procedure as the SAD is required. At first an input trace is selected by hand and the maximum distance is also tuned similarly to the SAD.
For the simple trigger, a peak finder, multiple inputs are needed in order to find all traces. At first a threshold level must be selected and a width of the peak. This makes sure that only peaks with the correct amount of width and height are found. The last input is the gap length, which is the minimum amount of distance between two peaks. This is all selected by hand based in the visually analysed recording and an AES-encryption. The parameters are then tuned by hand in order to split the recording in the exact amount of traces.

4.3. Analog Front-end

Since the measured electromagnetic signals are very weak, amplification is required to increase the signal strength. This is done to make sure that the measurements equipment can record the signal. Sometimes unwanted signals can reduce the SNR and in order to improve the results filtering is added in certain situations. The last part of the analog front-end is the measurement probe or antenna. Based on the measurement situation an appropriate solution needs to be chosen. In this section we will shortly discus the different solutions used during measurements.

4.3.1. Amplifier

Since the signals coming from a probe or antenna are very weak and the voltage range of an oscilloscope or SDR is much wider and higher than the measured signal, amplification is required for the recordings. An analog amplifier is chosen depending on the probe or antenna and the target platform. For the

measurements LNA amplifiers will be chosen based on the measured signal strength and SNR.
The near-field probes from Riscure have an included amplifier. The high-sensitive probe from Riscure has a 3 stage amplification with 100 mV / 1 µT and only 30 pT / $\sqrt{Hz}$ of noise. The low sensitive probe also has a 3 stage amplification but only with 20 mV / 1 µT and only 15 pT / $\sqrt{Hz}$ of noise. The amplifier has a bandwidth of 1 GHz for both probes. This is more than enough for the measurements on top of the surface of the IC.
The Langer probes do not have any internal amplification. For the measurements with the Langer probes an external amplifier is used, the Langer PA-203[25]. The Langer PA-203 has an amplification of 20 dB and a frequency range of 100 kHz to 3 GHz. It has a noise figure of 4.5 dB, making it a suitable LNA. The range and amplification is well suited for all of our target platforms.
For the far-field, low-noise ZFL-1000+ amplifiers from Mini-Circuits are used. Together with the appropriate filters the SNR is increased for better analysis. The amplifiers used during the experiments will be described in the next sections.

4.3.2. Filters

For the far-field measurements filtering is added to the analog front-end. For most measurements filtering is added to filter out unwanted frequencies and improve the SNR. Depending on the target and the leakage frequency the filters are chosen. These filters include Band-pass filters, High-pass filter and Low-pass filters. All of the filters are from Mini-Circuits and have different frequency ranges. The filtering frequencies and filters will be described for each of the measurements in the next sections.

4.3.3. Probes and Antennas

Three different probes and two different antennas are used during the experiments. We describe a probe as an antenna with a range up to 10 cm. The probes are used for the baselines measurements in the EM near-field and the antennas for the EM far-field. In the near-field antennas in the shape of coils are used, to capture the EM-radiation up to the distance of the diameter of the coil. these circular loop antennas are called probes during this thesis. The far-field antennas used during this research are log-periodic dipole antennas. These are chosen because of the directionality and relatively wide bandwidths.
The near-field probes from Riscure are the Riscure EMPHS and Riscure EMPLS. Respectively with the high-sensitive and low-sensitive amplification inside the probe. The coils have an inner diameter of 1.13 mm and an outer area of 1.60 mm and are orientated horizontally with the surface of the IC. This probe is only used for measurements performed on the surface on the IC, since it can only measure EM-radiation up to 0.57 mm.

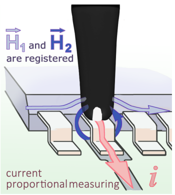
$\vec{H}_1$ and $\vec{H}_2$ are registered
current proportional measuring
i

Figure 4.4: Langer RF-U 2.5-2 probe from LANGER EMV-Technik [27]

Another near-field probe is the Langer RF-U 2.5-2[27]. The orientation of the coil inside the Langer RF-U is orientated vertically and has an angle of 90 degrees with the surface of the IC. An overview of the probes orientation can be seen in Figure 4.4. This probe is used to measure the EM-radiation of capacitors or wires, which will generate a vertical EM field. Instead of measuring on the surface of the IC, the decoupling capacitors on the PCB will be measured. This is used for shielded devices such as the Microsemi SF2 Advanced.
The last probe is the Langer RF-R 400-1[26]. The probe has a diameter of 25 mm, which makes it suitable for measurements up to 10 cm of distance. This probe can thus measure both the near- and far-field EM-radiation. This probe is used for both the baseline measurements and the far-field measurements. An overview of the probe and the measurement directions can be seen in Figure 4.5.

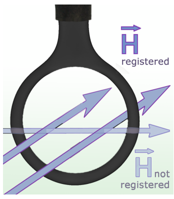

Figure 4.5: Langer RF-R 400-1 probe from LANGER EMV-Technik [26]

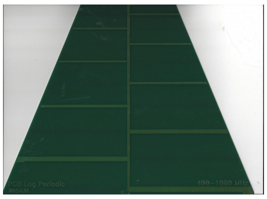

Figure 4.6: WA5VJB 400-1000 MHz log periodic antenna from Electronics [11]

For the far-field two log-periodic PCB antennas from WA5VJB[11] are used. The difference between the two antennas are the frequency ranges, which are 900-2600 MHz and 400-1000 MHz. This gives us a wide frequency range which can be captured by the antennas. Because the antennas grow in size when the frequencies are lower, we have only chosen to include antennas of frequencies higher than 400 MHz. The 400-1000 MHz antenna can be seen in Figure 4.6. Lower frequencies require antennas which are up to a couple of meters in size, which makes it extremely difficult to perform a stealth real-world attack.

4.4. Analysis Software

The last part of the side-channel attacks and TVLA tests is the analysis, which is performed in the Riscure Inspector tooling. The traces are split in the pre-processing step in MATLAB and converted to Inspector trace sets for the far-field results. While the baseline set-up makes use of Inspector to record the EM-radiation using an oscilloscope.

The baseline set-up has a self-made module which communicates with the target devices to be able to send data to encrypt. This module first communicates over User Datagram Protocol (UDP) to the target device to set-up the key. Then for each encryption the scope is armed, then the input data is send to the target and waits for the response data. The target device will trigger the oscilloscope with a GPIO pin and will send the data back to Inspector. Both the EM-radiation and the in- and output data of the AES are saved to disk.

For the far-field set-up a different module is made which loads the traces from a binary file exported from Matlab. In Matlab the recorded EM-radiation is split and matched with the correct in- and output data. This is then exported into a self-designed binary format, to reduce the disk size. The module in Inspector reads the exported file and generates a trace set, which can then be used for further analysis.

For both set-ups different modules from Inspector are used to analyse the data and generate the results from the next section. These modules allows us to easily perform attacks such as DEMA or CEMA, but can also perform the TVLA tests.

4.5. Target Devices

The chosen target device influences the electromagnetic leakage substantially[44]. In order to generate a more general analysis about the far-field electromagnetic leakage, multiple targets are analysed. In total three different target devices are analysed and compared. All three selected FPGA devices include an internal ARM-core, thus both software and FPGA implementations can still be compared. Because the need of higher bandwidth cryptographic devices and the increasing use of FPGA's[51], we chose to only include FPGA devices in the analysis.

For all target devices an AES-encryption implementation is made. These implementations communicate to the computer over ethernet using UDP. Both the key and input data will be send from the com-

puter. The device will then run the AES-encryption and return the encrypted result to the computer over UDP as well.
Since manufactures use different architectures for their IC's two different vendors are chosen, to be able to compare and generalize the far-field electromagnetic leakage. Hence we selected both Xylinx and Microsemi FPGA's, which will be discussed in this section in more detail. We we start looking at the Xylinx Pynq-Z1[9] and then look into the Microsemi SmartFusion2 and Microsemi SmartFusion2 Advanced.

4.5.1. Pynq

The Xylinx Pynq-Z1[9] is contains a ZYNQ XC7Z020-1CLG400C System on a Chip (SoC). It includes a 650MHz dual-core Cortex-A9 processor and an Artix-7 FPGA with 13300 logic slices. The 512MB memory is externally connected with an 16-bit DDR3 bus. An overview of the Pynq-Z1 board can be seen in Figure 4.7.

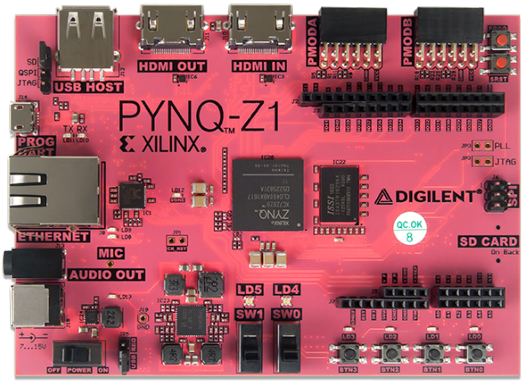

Figure 4.7: PYNQ-Z1 Zynq 7000, top view from Digilent [10]

For the Pynq-Z1 only software implementations are made, and the FPGA logic is only used to route the Cortex-A9 signals to the outside. The software implementation used in this target is the The OpenSSL Project [54] implementation, because the Cortex-A9 can make use of the 32-bit optimization.

4.5.2. Microsemi SmartFusion2

The Microsemi SmartFusion2 starter kit contains a Microsemi M2S050 SoC. The SoC includes a 166 MHz Cortex-M3 with internal memory and 56340 logic elements. The board consists of the main System on Module (SOM) containing the SoC and an extension board containing the IO. An overview of the board can be seen in Figure 4.8.

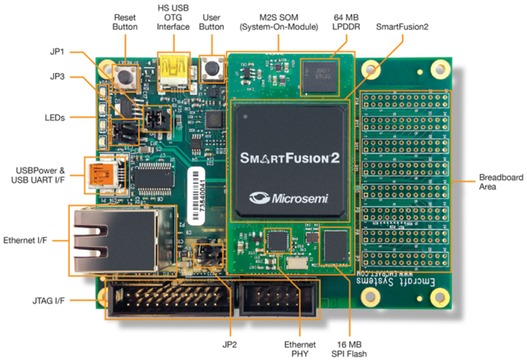

Figure 4.8: SmartFusion2 Starter Kit from Emcraft Systems [14]

For the Microsemi SmartFusion2 starter kit both a software and the combinatorial FPGA AES implementations are made. The communication through ethernet is implemented in the FPGA fabric for both implementations. The Cortex-M3 is also a 32-bit processor, and the same The OpenSSL Project [54] implementation is used for the software version.

4.5.3. Microsemi SmartFusion2 Advanced

The Microsemi SmartFusion2 Advanced has a M2S150 SoC, with an internal 166 MHz Cortex-M3 processor similar to the M2S050. The FPGA has 150000 logic elements and the SoC is part of the main board. The SoC has a metal shielding attached for distributing the heat. The board overview can be seen in Figure 4.9.

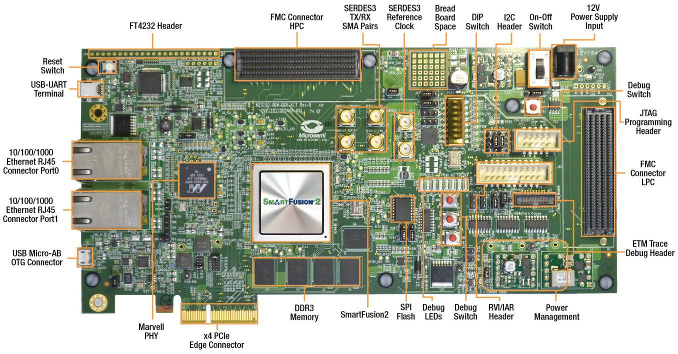

Figure 4.9: SmartFusion2 Advanced Development Kit from Emcraft Systems [13]

The Microsemi SmartFusion2 is also capable of running both the combinatorial and the software AES implementations. The ethernet communication for both implementations is implemented in the Cortex-M3 processor, because the FPGA logic from the starter kit was not compatible. Since it makes use of the same Cotex-M3 processor the same The OpenSSL Project [54] implementation is used for the software AES.

5 Baseline Results

Each implementation and target platform leaks differently. In order to attack platforms in the far-field using SDR an analysis on how the platform leaks must be made first. This analysis is performed using the baseline set-up described in Section 4.1. Using an XY-table an TVLA-test is performed at multiple locations above or under the main IC. These tests are performed at different heights from the IC using different analog front-ends. The t-test analysis is performed in both the time and frequency domain. These results show both the alignment problems in the time domain and which frequencies leak per position above the IC.

A TVLA-test is chosen in order to be able to do a quick quantitative leakage assessment of the IC. This makes it easy to compare different platforms and perform a lot of measurements with different conditions. In order to analyse what is leaking inside the chip and at which position a lot of measurements are required. Performing DEMA or CEMA attacks at all positions takes a lot of time, and doesn't give a good quantitative measurement.

The test is performed at different heights to analyse how the leakage changes over distance. This gives an indication of what is required to perform attacks in the EM far-field. It also shows how precise the antenna or probe needs to be placed at a certain distance.

The measurements are converted to the frequency domain using a FFT. This is performed just before the TVLA inside Inspector. This gives a good indication on which frequencies the target is leaking. Alignment problems aren't occurring in the frequency domain. Comparing the results to the time domain would give a good indication about alignment issues.

Two different leakage points, *SubBytes* and *MixColumns*, in the AES encryption will be analysed for the different platforms. These refer to the intermediate values which are set during the T-test. The *SubBytes* leakage point is defined as the intermediate value after the *SubBytes* operation of round 5. The *MixColumns* leakage point is respectively defined as the intermediate value after the *MixColumns* operation of round 5. All implementations make use of a loop in the AES implementation, meaning that each round reuses the same code or combinatorial circuit.

For all TVLA measurements 4000 traces will be recorded per position. It is important to keep this similar for each recording in order to be able to compare the results. The results from Ramsay [45] showed significant leakage which would require less than 4000 measurements for a TVLA test.

The baseline results for the three different targets are discussed in this section.

5.1. Pynq

The Pynq ARM IC is configured to run at a frequency of 100 MHz and the external Double Data Rate (DDR) memory is configured to run at 400 MHz. The implementation that is measured in this section is running fully inside the ARM in software. All measurements are made using 2.5 GSa/s to ensure frequencies up to 1.25 GSa/s can be reconstructed[38]. This is kept the same for all measurements to ensure consistency between measurements.

In Figure 5.1 can be seen how the XY-table is orientated. The probes are scanning the full surface of the

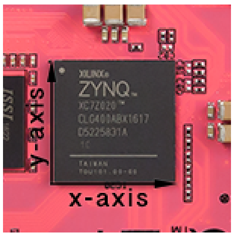

Figure 5.1: PYNQ-Z1 reference for the XY-table measurements.

IC. The centre of the probe will reach all edges and corners, to ensure the full IC is scanned. The area that is being scanned is 17 by 17 mm, the exact size of the IC. The IC will be scanned from the top side of the PCB, since this would give the shortest distance to the silicon inside the IC.

5.1.1. On surface

For the Pynq the low-sensitive probe from Riscure is chosen, since the EM leakage was expected to be high. Since the probe from Riscure includes an amplifier no extra amplifier is added on the analog side. The probe is directly on top of the IC's surface and scans the full surface of the IC. A resolution for the XY-table of 25 by 25 is chosen, because the size of the chip is only 17 by 17 mm. This leads to a resolution of 680 by 680 µm per measurement. A TVLA test using only the *MixColumns* as target is chosen. Previous research by Ramsay [45] has shown that the *MixColumns* leak more than the *SubBytes* intermediate value.

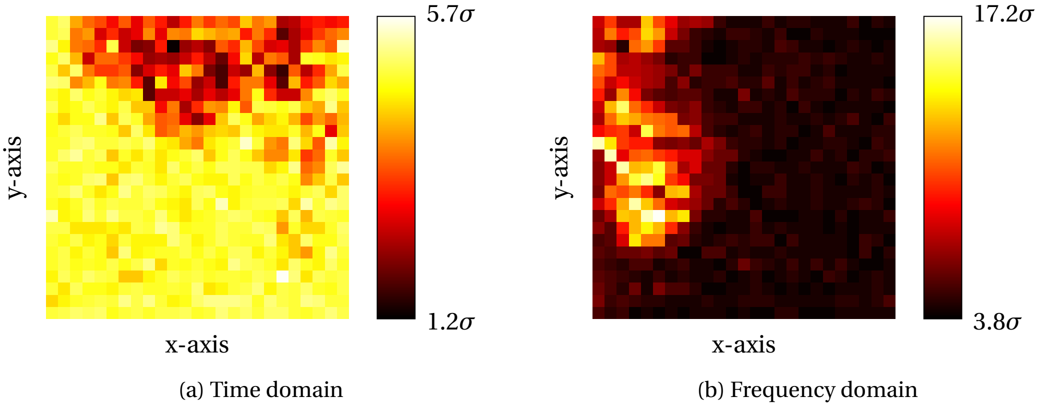

(a) Time domain

(b) Frequency domain

Figure 5.2: Frequency and time-domain TVLA tests in σ on the Pynq using MixColumns at the IC surface.

In Figure 5.2 the results from the TVLA tests in both the time- and frequency domain can be seen. It shows a clear distinction between the time- and frequency domain. The time domain has a much lower overall standard deviation, which is abnormal. This can be explained by the fact that the traces aren't aligned during round 5, which can be seen in Figure 5.3. This figure shows two traces from the highest leakage point in the frequency domain (4,9). Even though the leakage in time domain is less it still exceeds the 4.5σ at multiple locations.

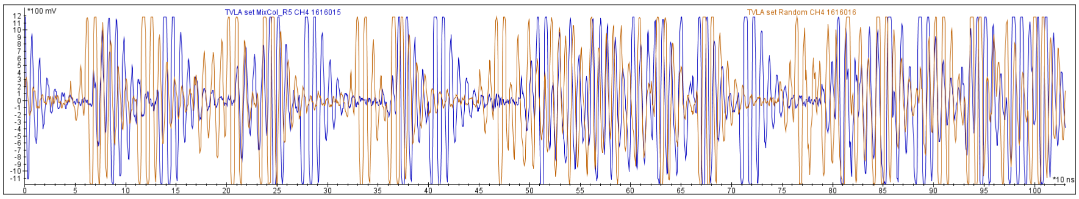

Figure 5.3: Pynq misalignment in the time domain at position (4,9).

The misalignment in time-domain can be caused by multiple factors but is extremely hard to determine. The most likely cause is the unreliability of the external memory latency and speeds. Due to the fact that we are looking at the values of *MixColumns* round 5 and the trigger is activated before round 0, alignment is more difficult. This can be explained by the fact that the distance between the trigger and leakage is larger, which causes jitter inside a single round to add up over five rounds.
The results in frequency domain shows that there is leakage, which is clearly exceeding the 4.5σ. At position (4,9) a maximum leakage of 17.2σ is measured. An overview of that position can be seen in Figure 5.4. It shows that multiple frequencies are exceeding the 4.5σ line. Around 400 MHz a clear peak can be seen, which is exactly at the same frequency of the memory bus. This means that the connection between the processor and the external memory is leaking information.

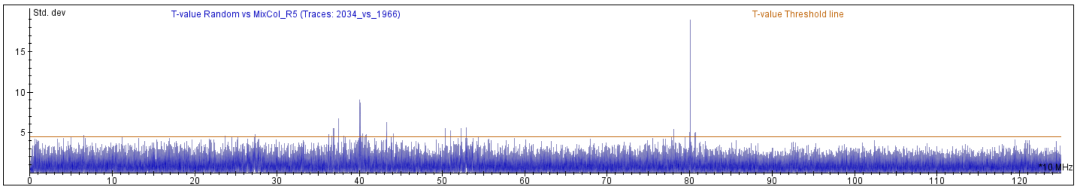

Figure 5.4: Pynq TVLA with MixColumns in frequency domain results at position (4,9).

The highest peak is at 800 MHz, which is exactly twice the frequency of the external DDR memory. Since the external memory is DDR, this means that data is transferred twice a clock cycle. The leakage of 800 MHz is occurring in our measurements because the memory clock is configured at 400 MHz and data over that bus is transferring at 800 MHz.

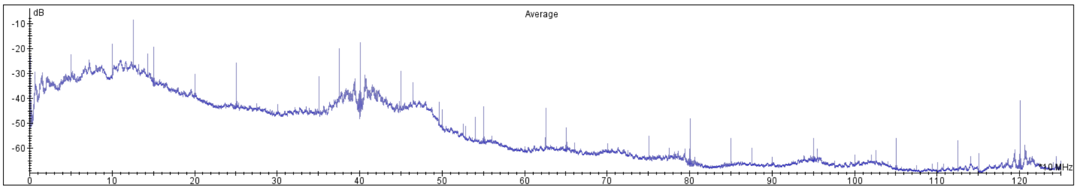

Figure 5.5: Pynq FFT at position (4,9) averaged 4000 traces.

Leakage is expected to be present in the clock frequencies of the IC[28, 59], due to modulation on the core clock. The results from the TVLA test show that no significant leakage, more than 4.5σ, is measured around the IC frequency. In Figure 5.5 an FFT is shown from position (4,9) which is averaged over 4000 traces to give a more clear result. It shows that the IC frequency is measured and significantly present.

5.1.2. Langer probe at 1 cm

For these tests the test set-up is adapted to measure from a 1 cm distance from the surface of the IC. The Riscure probe has a diameter of 1.6 mm which makes it unable to record the EM radiation at 1 cm. For these TVLA tests the Langer RF-R probe is selected, together with an external amplifier. The resolution of the XY-table is changed to 15 by 15, because the Langer probe covers a larger EM area. This will generate measurements with a resolution of 1.133 by 1.133 mm.
In Figure 5.6 the TVLA results are shown of the *MixColumns* intermediate. Similar to the results from the surface measurement it shows a clear distinction between time- and frequency domain. The trigger was not changed during this test, which explains why there are still alignment problems in the time domain. Both in the time- and frequency domain there are a lot of positions on which 4.5σ is exceeded. In the time domain there seem to be more leakage in the bottom left corner.
The results in the frequency domain show that most leakage is present in the upper right and lower left corners. Compared to the results from the previous surface scan the leakage is present in different locations. This can be explained by the fact that other components and traces from the PCB are also emitting EM radiation. The larger antenna range also covers parts of the PCB.
When we look at the result at position (11,12) in frequency domain in Figure 5.7, we can clearly see at

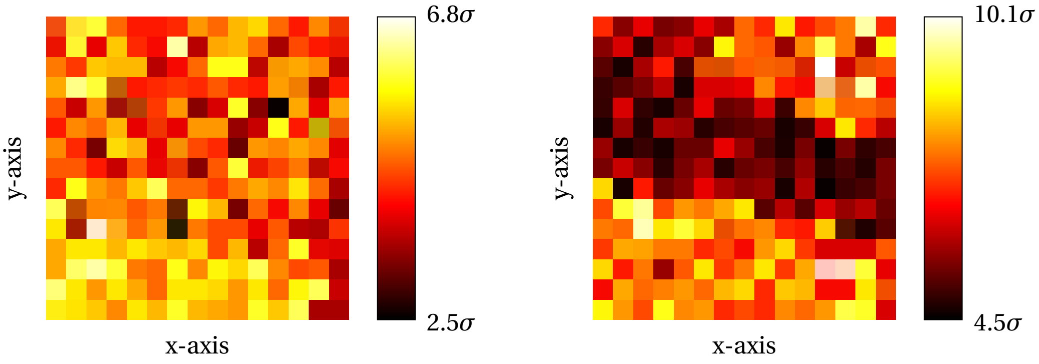

(a) Time domain

(b) Frequency domain

Figure 5.6: Frequency and time-domain TVLA tests in σ on the Pynq using MixColumns at 1 cm height.

400 MHz peaks exceeding the 4.5σ. the results on the surface showed more leakage around the 800 MHz caused by the DDR memory, which is not visible at 1 cm. This can be explained by the fact that only inside the IC the DDR of the memory is present. Since the Langer probe covers more of the PCB we are now recording the traces from the IC to the external DDR memory.

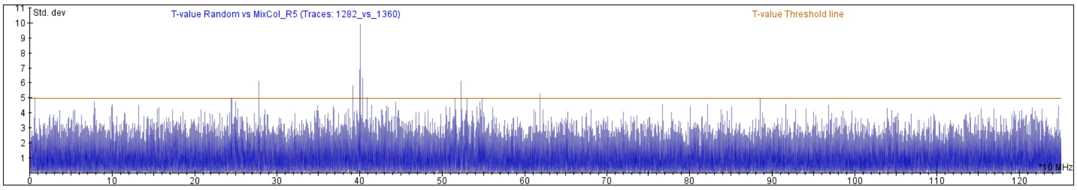

Figure 5.7: Pynq TVLA with MixColumns in frequency domain results at position (11,12) at a height of 1 cm.

The leakage around the 400 MHz has a similar standard deviation compared to the surface measurements. This can be explained by the frequency of the leakage and the length of the PCB traces between the memory and IC. With the surface measurement only the connections inside the IC are emitting EM radiation. These connections are much smaller than the traces of the PCB, which are thus emitting much more EM radiation. The other frequencies that exceed the 4.5σ are similar to the measurements from the surface.

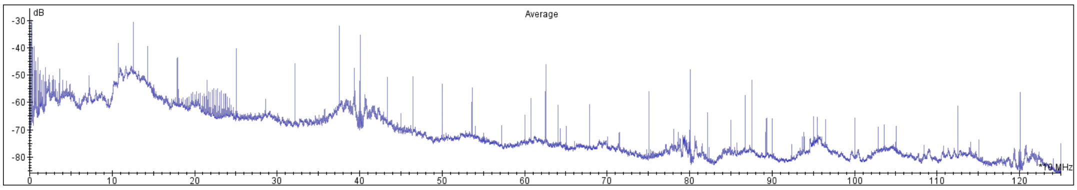

Figure 5.8: Pynq FFT at position (11, 12) averaged 4000 traces at a height of 1 cm.

In Figure 5.8 the averaged FFT of 4000 traces from position (11, 12) is shown. Compared to the surface results it shows much more peaks at different frequencies. Because the range of the antenna is much further it will record a lot more signals from other sources inside the IC and from the PCB. This also means that the SNR will decrease if those new frequencies are not leaking the *MixColumns* intermediate, which should be visible in the time based TVLA test. Even though the 800 MHz frequency is still present in the measured traces no significant leakage was measured, meaning that the *MixColumns* intermediate is not modulated on the 800 MHz frequency measured from 1 cm distance.
The TVLA results from the *SubBytes* intermediate can be seen in Figure 5.9. As expected the TVLA measurements in time- and frequency from the *SubBytes* intermediate are significantly lower than the *MixColumns* intermediate. The results from Ramsay [45] in the near-field are also valid for measurements

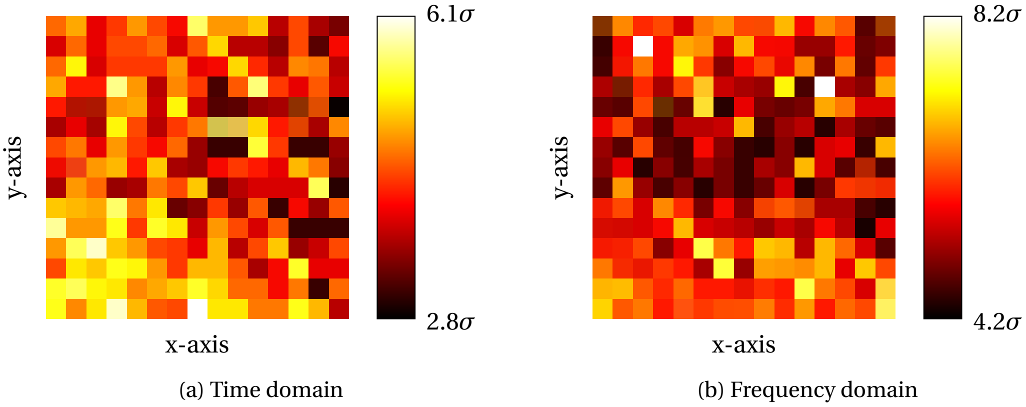

(a) Time domain

(b) Frequency domain

Figure 5.9: Frequency and time-domain TVLA tests in σ on the Pynq using SubBytes at 1 cm height.

from 1 cm distance. The results still show that both in the time- and frequency domain there are still positions which exceed the 4.5σ. This means that even with bad assumptions it could still be possible to extract the secret key.

5.1.3. Langer probe at 2 cm

For these tests the test set-up is adapted to measure from a 2 cm distance from the surface of the IC. The same resolution and Langer probe with amplification is used as the previous tests at 1 cm. Only the height of the probe is changed to 2 cm above the surface of the IC.

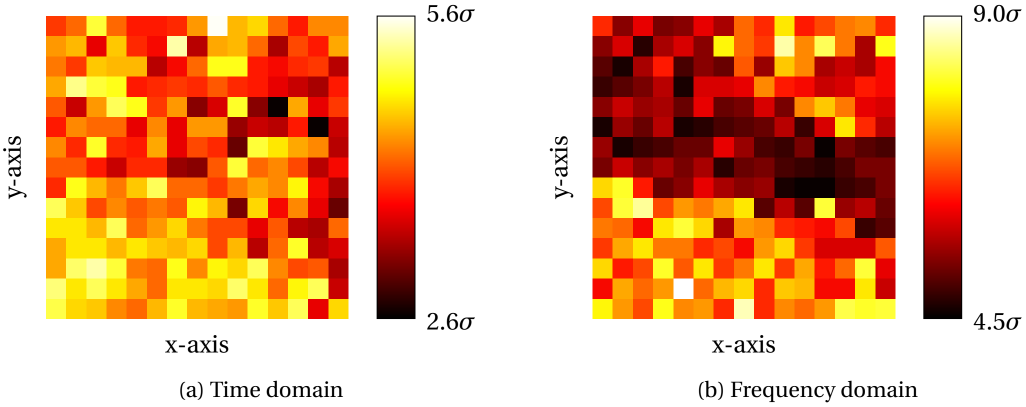

(a) Time domain

(b) Frequency domain

Figure 5.10: Frequency and time-domain TVLA tests in σ on the Pynq using MixColumns at 2 cm height.

The *MixColumns* frequency domain TVLA test at 2 cm in Figure 5.10 shows that there is up to 9σ. This means that there are still a lot of positions on where the TVLA tests exceeds the 4.5σ. The overall results of the measurements at 1 cm look very similar and the main difference is the standard deviation, which is significantly decreased over the whole measured surface. the locations on where the leakage is present are still similar to the results from 1 cm. This means that when measuring with a larger range antenna it becomes easier to record at a EM leaking position.

The measured frequencies and the leaking frequencies from a distance of 2 cm are similar to the measurements from 1 cm. The overall SNR is decreased, which is expected when the distance to the IC increases. To increase the SNR for CEMA or DEMA attacks, the amount of measured traces can be increased. This would mean that attacks from 1 cm can easily also be achieved at 2 cm, without changing any parameters except the amount of required traces.

The results from the *SubBytes* intermediate at 2 cm can be seen in Figure 5.11. These results still show a lower standard deviation, similar to the measurements from 1 cm and on the surface. The T-tables AES

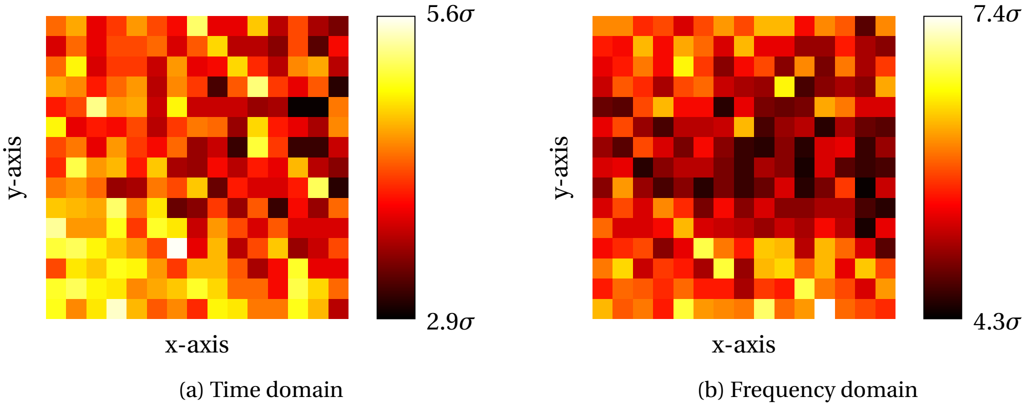

(a) Time domain

(b) Frequency domain

Figure 5.11: Frequency and time-domain TVLA tests in σ on the Pynq using SubBytes at 2 cm height.

implementation does not directly include the *SubBytes* intermediate, which explains the lower standard deviation compared to the *MixColumns* intermediate. These tests show that when the distance increases the primitives of the EM leakage still hold.

5.2. Microsemi SmartFusion2

The Microsemi SmartFusion2 starter kit ARM is configured to run at a frequency of 142 MHz. The T-tables implementation is running in the software ARM core. All measurements are made using 2.5 GSa/s to ensure frequencies up to 1.25 GSa/s can be reconstructed[38]. This is kept the same for all measurements to ensure consistency between measurements.

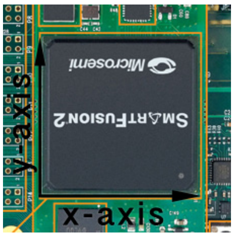

Figure 5.12: SmartFusion2 Starter kit reference for the XY-table measurements.

In Figure 5.12 the orientation of the XY-table is shown. The XY-table is configured to scan the full surface of the IC, where the center of the probe is reaching the edges and corners of the IC. The Microsemi SF2 M2S050 IC has a size of 23 by 23 mm. The IC has no shielding on top and will be scanned from the top side of the PCB.

5.2.1. On surface

Because the Microsemi SF2 does not include a heat-sink or other type of shielding, the low-sensitive probe from Riscure was chosen for the measurements. This probe includes an amplifier and is directly connected to the oscilloscope. The shortest path from the silicon to the probe is achieved by measuring from the surface of the IC. A resolution for the XY-table of 25 by 25 is chosen, because the size of the chip is only 23 by 23 mm. This measures squares of 920 by 920 µm, which are well within the probes range. Since the leakage of the Microsemi SF2 was analysed in detail before, only the *MixColumns* intermediates are used during the TVLA tests which showed more leakage.

In Figure 5.13 can be seen that there is a clear distinction between the time- and frequency domain TVLA tests. This means that also the Microsemi SF2 suffers from trace alignment problems. Even though the misalignment problem, the time domain leakage still exceeds the 4.5σ at certain positions. In the top-middle, there is leakage in the frequency domain which is also measured in the time domain.

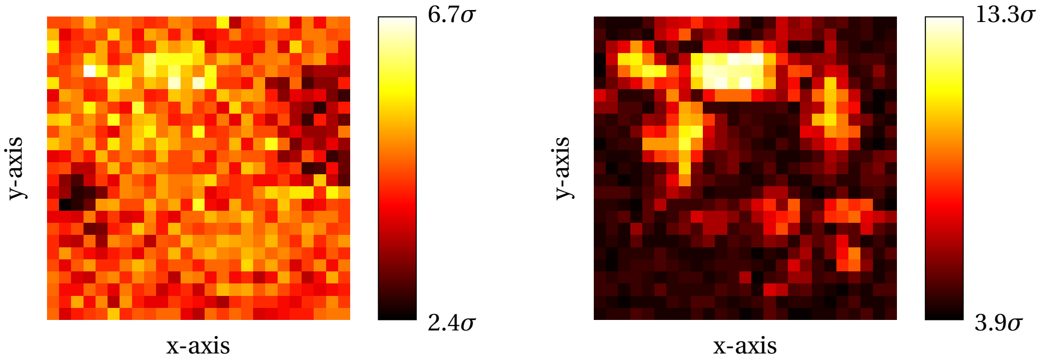

(a) Time domain

(b) Frequency domain

Figure 5.13: Frequency and time-domain TVLA tests in σ on the Microsemi SF2 using MixColumns at the IC surface.

This is where the ARM core is expected to be inside the SoC.

For the Microsemi SF2 it seems that the misalignment is caused by the generation of the trigger. This means that in between AES rounds no extra jitter is added, and only initial jitter before the first round is causing the misalignment. After analysis of the separate traces at position (13, 5) only misalignment problems can be seen before the first round. The trigger misalignment can be explained by the fact that the generation of the trigger is performed in the FPGA clock domain, while the AES encryption is performed in the ARM clock domain. The switching between the various clock domains could cause significant timing delays between the generation of the trigger and the actual AES encryption.

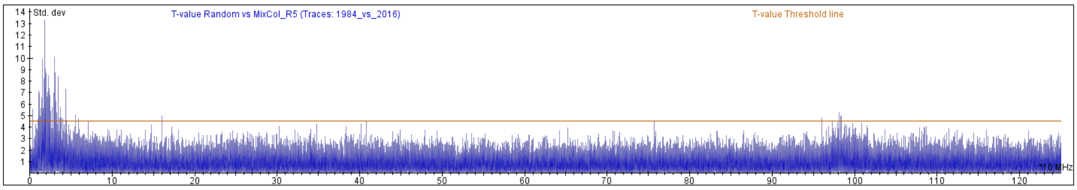

Figure 5.14: Microsemi SF2 TVLA in frequency domain results at position (13,5).

In the frequency domain TVLA test the highest deviation is measured at position (13, 5), which is 13.3σ. In Figure 5.14 the results of this TVLA test can be seen. It shows that there is a lot of leakage in the frequencies below 50 MHz. Around 987 MHz there is also leakage exceeding the 4.5σ. This is caused by the Phase-locked loop (PLL), which has a divider of 7 and 987/7 = 141 MHz. The actual measured clock frequency is 141 MHz instead of the set clock frequency of 142 MHz.

The leakage in the lower frequencies can be explained as direct emanations, because all clock frequencies are higher than 50 MHz. The leakage around 987 MHz is clearly modulates leakage of the main ARM clock frequency PLL. Even though the results below 50 MHz show a higher standard deviation, the 987 MHz is much more interesting for the far-field. Because of the higher frequency it can be recorded much easier using SDR and also requires a much smaller antenna.

5.2.2. Langer probe at 1 cm

TVLA tests are performed at a 1 cm height from the surface of the IC. To be able to compare the results to the other targets, the same set-up is used as the Microsemi SF2 with the Langer probe at 1 cm. The XY-table is set-up to generate measurements of 15 by 15 positions, which will generate a resolution of 1.533 by 1.533 mm. Since the far-field was not analysed in detail before we chose to include both the *MixColumns* and *SubBytes* intermediates for the TVLA test.

The results in time domain for *MixColumns* in Figure 5.15 shows us that the average leakage has dropped. But it clearly shows a vertical line in the centre, where the leakage is exceeding the 4.5σ. Compared to the near-field results it shows a significant change in leakage positions. In the near-field

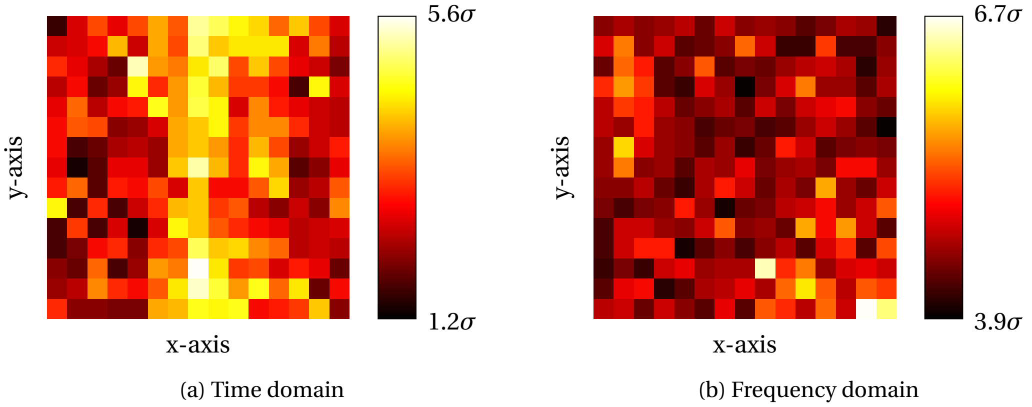

(a) Time domain (b) Frequency domain

Figure 5.15: Frequency and time-domain TVLA tests in σ on the Microsemi SF2 using MixColumns at 1 cm height.

the ARM processor was clearly recognizable, but the far-field shows an unrecognisable leakage pattern. But even though the Microsemi SF2 has trigger alignment problems, the time domain exceeds the 4.5σ at a lot of positions. This means that the device is clearly leaking secret information.

The frequency domain has an even lower average, as can be seen in Figure 5.15. But it has higher peaks in specific spots and the lower right corner. These peaks exceed the 4.5σ, but seem a bit random and could be considered false positives. The results do show that even though the maximum leakage in the frequency domain is higher, the leakage in the time-domain seem to be more consistent.

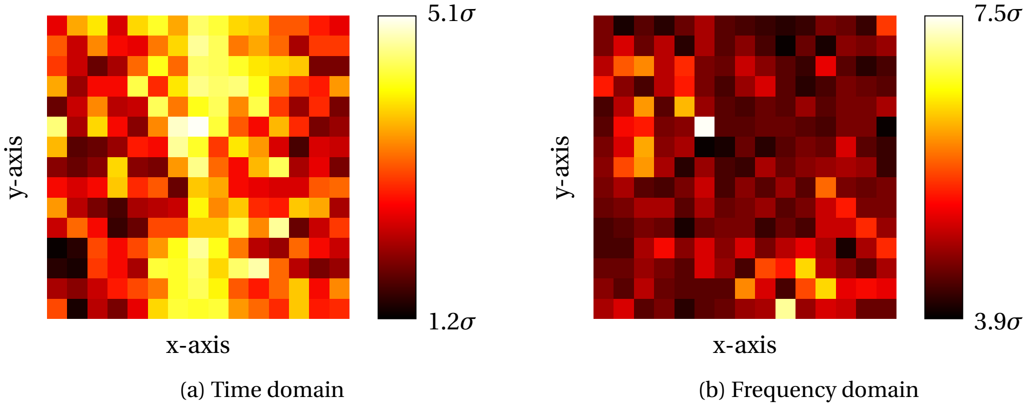

(a) Time domain (b) Frequency domain

Figure 5.16: Frequency and time-domain TVLA tests in σ on the Microsemi SF2 using SubBytes at 1 cm height.

As expected from the near-field the results in time-domain show a lower standard deviation for the *SubBytes* compared to the *MixColumns* intermediates. The results from the *SubBytes* TVLA test can be seen in Figure 5.16. It follows the same pattern as the *MixColumns* with the vertical centre exceeding the 4.5σ.

Even though the results in the frequency domain show a higher maximum leakage, it seem to still be caused by the small peaks at specific locations. Those peaks also look a lot like false positives, which could explain the higher maximum leakage.

5.2.3. Langer probe at 2 cm

The TVLA tests are preformed at a height of 2 cm from the surface of the IC. A similar set-up is used as the Langer probe at 1 cm, and only the hiehgt is changed for this measurement.

The results in the time- and frequency domain from the *MixColumns* intermediate are shown in Figure 5.17. The time domain results clearly show a similar pattern, compared to the results from 1 cm height. The overall standard deviation is lower, but not very significantly. The clear vertical centre line is still

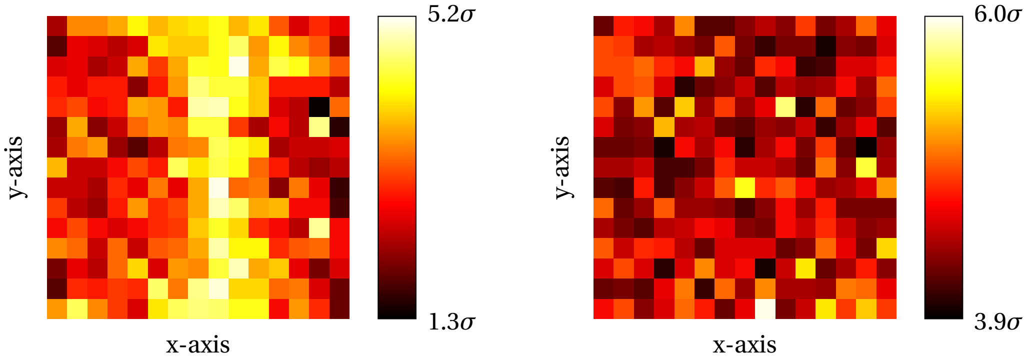

(a) Time domain

(b) Frequency domain

Figure 5.17: Frequency and time-domain TVLA tests in σ on the Microsemi SF2 using MixColumns at 2 cm height.

exceeding the 4.5σ and the device is still clearly leaking secret information.

The frequency domain results are similar in pattern to the 1 cm height results. The overall standard deviation is also reduced, which is expected when the distance to the silicon increases. But since the pattern is still similar and it is still exceeding the 4.5σ this means that with more traces it could still be possible to extract the secret key.

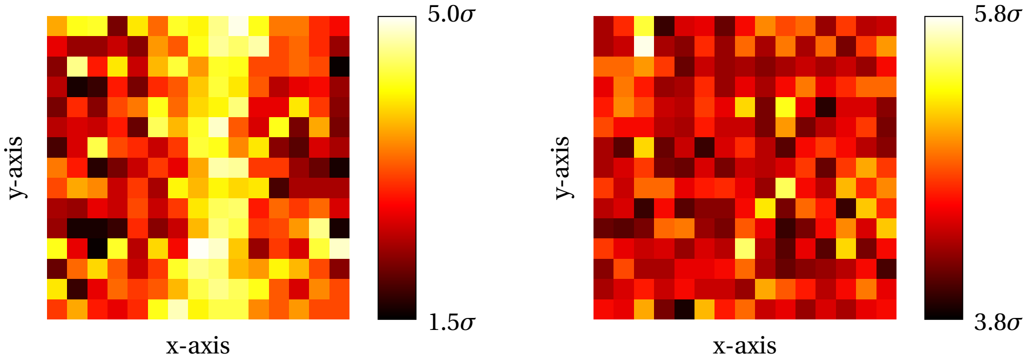

(a) Time domain

(b) Frequency domain

Figure 5.18: Frequency and time-domain TVLA tests in σ on the Microsemi SF2 using SubBytes at 2 cm height.

The *SubBytes* intermediate leakage at 2 cm is shown in Figure 5.18. It shows an overall lower standard deviation, which is similar to the result from 1 cm. This means that when the distance increases the leakage from the *SubBytes* intermediate is always lower than the *MixColumns*. This is expected as the intermediate results from *MixColumns* are directly present in the T-tables implementation, while the *SubBytes* are not.

5.3. Microsemi SmartFusion2 Advanced

The Microsemi SmartFusion2 Advanced ARM is configured to run at a frequency of 142 MHz. Similar to the other set-ups the scope is configured to record at 2.5 GSa/s to be able to compare the results. The SmartFusion2 Advanced board is a single PCB containing the SoC.

The reference and orientation of the XY-table is shown in Figure 5.12. The Microsemi SF2 M2S150 SoC has a size of 35 by 35 mm and a metal heat-sink on the top surface. Since the heat-sink on the surface of the SoC is made out of metal and grounded it acts as EM shielding. Therefore the measurements are performed in the back side of the PCB, where the decoupling capacitors of the SoC are placed. The XY-table is configured to scan the full surface of the SoC, where the center of the probe is reaching the edges and corners of the SoC.

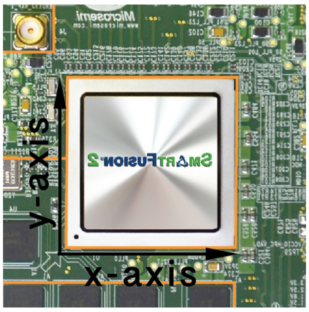

Figure 5.19: SmarFusion2 Advanced kit reference for the XY-table measurements. Measurement on the back, which makes the images flipped horizontal from the front.

5.3.1. On surface

The TVLA test measurements are preformed using the Langer RF-U probe with the external Langer amplifier. The height of the XY-table is set to the surface of the highest capacitor. The XY-table is configured to take 30 steps per axis, which will generate measurement of 1.17 by 1.17 mm resolution. The Microsemi SF2 Advanced was never analysed before, therefore both *MixColumns* and *SubBytes* intermediates are analysed.

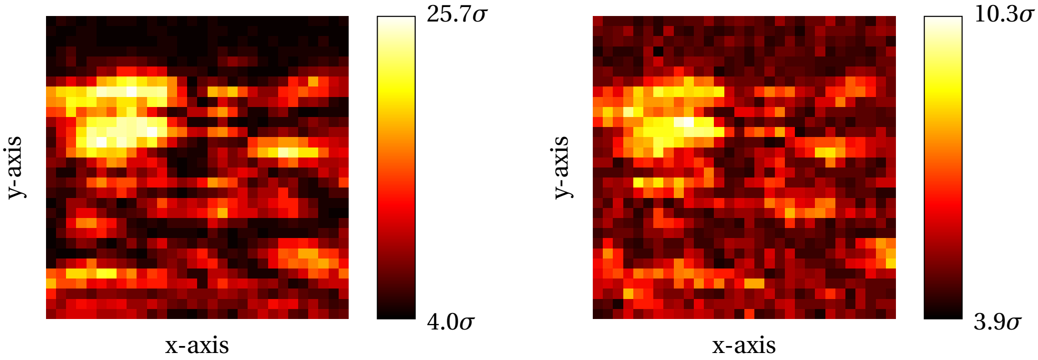

(a) Time domain

(b) Frequency domain

Figure 5.20: Frequency and time-domain TVLA tests in σ on the SF2 Advanced using MixColumns at the IC surface.

The comparison between the results in time- and frequency from *Mixcolumns* intermediate in Figure 5.20 show a very similar pattern. The overall leakage in time domain is higher, which is normally expected behaviour. The results from the frequency domain add up several P_{data} and P_{op} from different points in time, causing the noise floor to rise. This also means that the alignment of traces in the time domain is correct, which can be explained by the fact that both the trigger and AES encryption are performed inside the ARM core.

The time domain results show a lot of leakage in the top left corner, where the ARM core capacitors are expected to be. It also shows a lot of other locations where the leakage is exceeding the 4.5σ, which are more spread around the SoC. This can be explained by modulated leakage on other signals inside the SoC or a direct connection between the capacitor and the ARM core. The maximum leakage is extremely high 25.7σ, which means that the ARM core is leaking the secret information clearly.

We will take a closer look at position (9, 20) where the maximum leakage in the time domain is. The frequency TVLA test results at (9, 20) can be seen in Figure 5.21. It shows that most leakage is occurring around the 50 MHz, which is mostly caused by direct emanations. Around the 100 MHz also leakage is exceeding the 4.5σ, which is much easier to attack with SDR and is caused by modulated emanations. There are peaks at higher frequencies, but they are evaluated to be false-positives.

The results from the *SubBytes* intermediate can be seen in Figure 5.22. It follows the exact same pattern as the *MixColumns* leakage, except that the overall leakage is lower. This is expected behaviour, due to

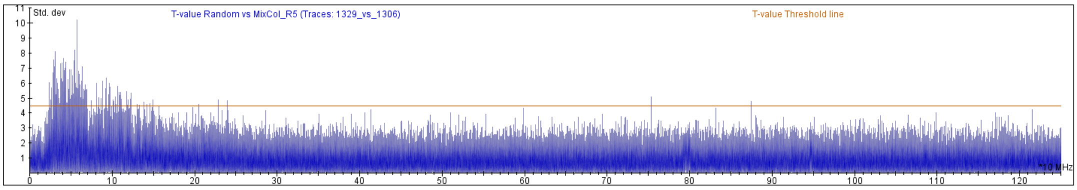

Figure 5.21: Microsemi SF2 Advanced TVLA in frequency domain results at position (9, 20).

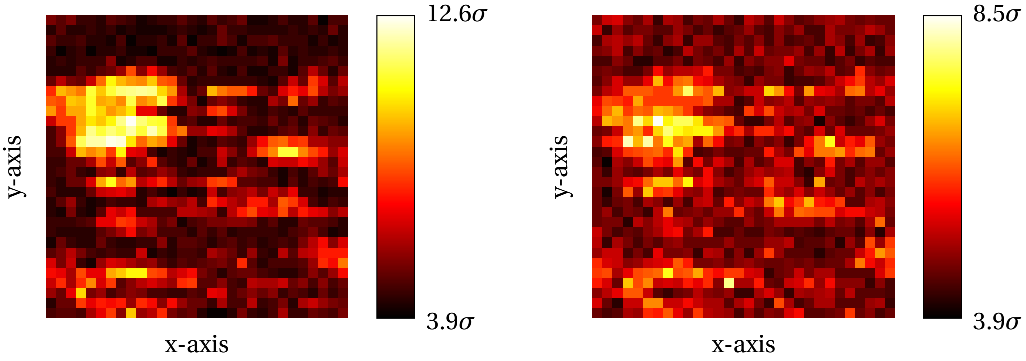

(a) Time domain

(b) Frequency domain

Figure 5.22: Frequency and time-domain TVLA tests in σ on the SF2 Advanced using SubBytes at the IC surface.

the T-tables implementation no directly using the *SubBytes* intermediate. Even though the *MixColumns* leakage is higher, the *SubBytes* leakage still exceeds the 4.5σ for a lot of positions. This means that an attack using the *SubBytes* intermediate can also be successful.

5.3.2. Langer probe at 1 cm

A similar set-up as the Pynq and Microsemi SF2 will be used. The Langer probe is set to a 1 cm height from the back of the PCB. The Langer RF-R probe is selected together with an external amplifier. The XY-table is configured to perform 15 by 15 measurements, which will generate measurements of 2.34 by 2.34 mm resolution. The *MixColumns* and *SubBytes* intermediates will be used for the TVLA measurements.

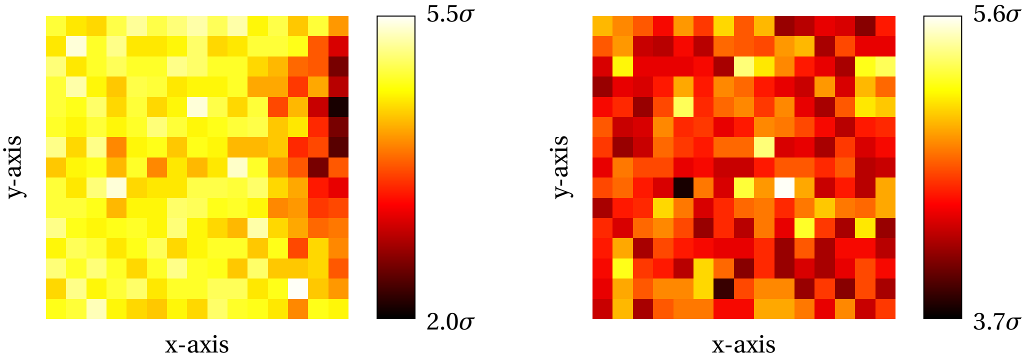

(a) Time domain

(b) Frequency domain

Figure 5.23: Frequency and time-domain TVLA tests in σ on the SF2 Advanced using MixColumns at 1 cm height.

In Figure 5.23 the TVLA results from the *MixColumns* intermediate can be seen. Compared to the surface measurements it shows a significant difference in shape and standard deviation. In both the time- and frequency domain the ARM processors can not be recognized and the EM leakage is much lower.

The results in time domain show that there are still a lot of locations on which the 4.5σ is exceeded, while the frequency domain only has very specific locations. The lower leakage is expected due to the increase of distance, but can also be explained by the directionality of the Langer probe. The Langer RF-U probe from the surface measurements, mainly measures EM radiation in the horizontal direction. While the Langer RF-R measures EM radiation main in the vertical direction, which is expected to radiate less due to the EM field of the capacitors is mainly in the horizontal direction.

In the time domain the leakage shows a clear pattern on the right where the leakage is significantly lower. Most of the IC is exceeding the 4.5σ, meaning that an attack could still be possible. This can be explained by the bigger antenna range, which is covering much more capacitors at a single position. This means that the placement of the antenna can be less precise compared to the surface results.

The frequency domain only has very specific positions where the leakage is exceeding the 4.5σ. The patterns of the leakage locations are not as clear as the surface results and show a lot more noise. This is also caused by the larger size of the antenna, which is decreasing the SNR and recording a lot of non leaking frequencies. Compared to the time domain the standard deviation is on average lower, as was expected from the surface measurements.

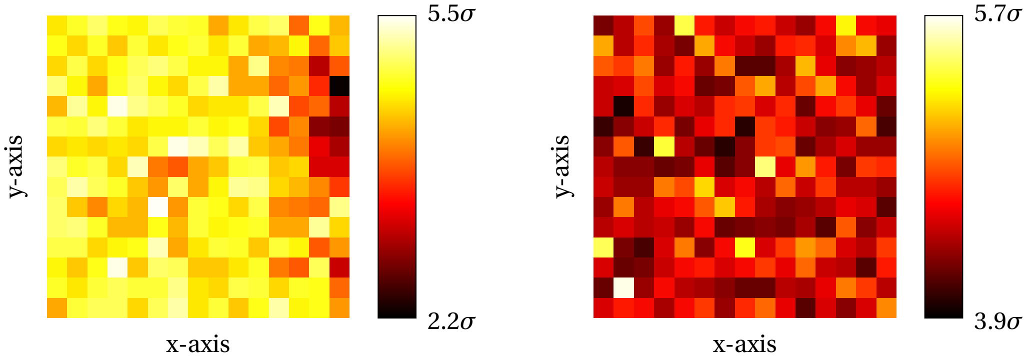

(a) Time domain

(b) Frequency domain

Figure 5.24: Frequency and time-domain TVLA tests in σ on the SF2 Advanced using SubBytes at 1 cm height.

The results from the *SubBytes* intermediate from Figure 5.24 show a lower standard deviation. This was expected compared to the results from the surface measurements. Similar patterns as the *MixColumns* intermediate are visible in the time domain. The frequency domain shows a more random leakage and looks like random noise. Even though the leakage is less than the *MixColumns* intermediate, in the time domain a lot of positions are exceeding the 4.5σ. The frequency domain only has certain spikes exceeding the 4.5σ, which can not be clearly linked to frequencies of the IC.

5.3.3. Langer probe at 2 cm

For the measurements at 2 cm the same set-up from 1 cm was used. Only the height of the XY-table is changed to 2 cm from the back of the PCB. The same Langer RF-R probe and amplifier is used and the XY-table is configured with 15 by 15 steps. Both the *MixColumns* and *SubBytes* intermediate values are evaluated during the TVLA tests.

The results from the *MixColumns* intermediate can be seen in Figure 5.25. Both the leakage in time- and frequency domain has decreased compared to the results from 1 cm. The time domain exceeds the 4.5σ on a lot of positions, similar to the results from 1 cm. It shows a clear spot in the top right corner with less leakage, which on 1 cm was measured more to the middle right. This could be due to the fact that the probe was slightly changed in orientation in between the tests while changing the height. Since most of the positions are exceeding the 4.5σ it seems that the positioning of the probe does not need to be as precise as the surface measurements.

The frequency domain shows on average a lower leakage than the 1 cm measurements, but the maximum leakage is exceeding the results from 1 cm at a single position (14, 14). After further investigation of the results at (14, 14) it is clearly a false positive, as the 6.0σ is only measured at a single point in

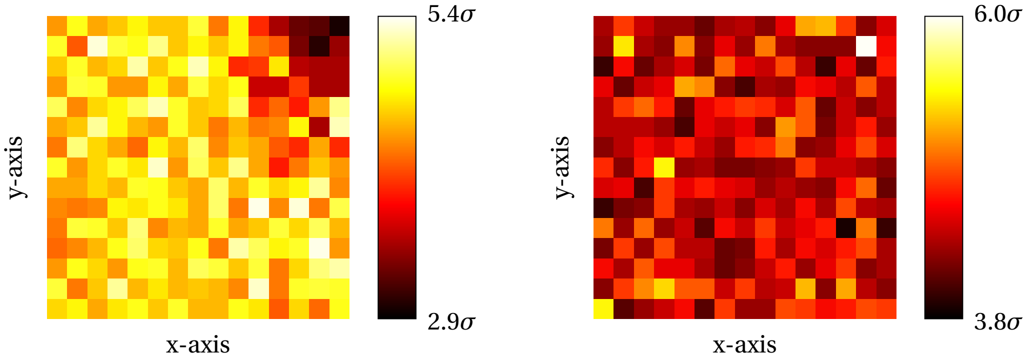

(a) Time domain

(b) Frequency domain

Figure 5.25: Frequency and time-domain TVLA tests in σ on the SF2 Advanced using MixColumns at 2 cm height.

the FFT and around the frequency the leakage is not exceeding the 5.2σ. The other positions where the 5.0σ is exceeded also show similar results, with a very high noise floor. Even though the 4.5σ is exceeded at several positions, it seems that most positions are false positives and the results do not seem to be significant.

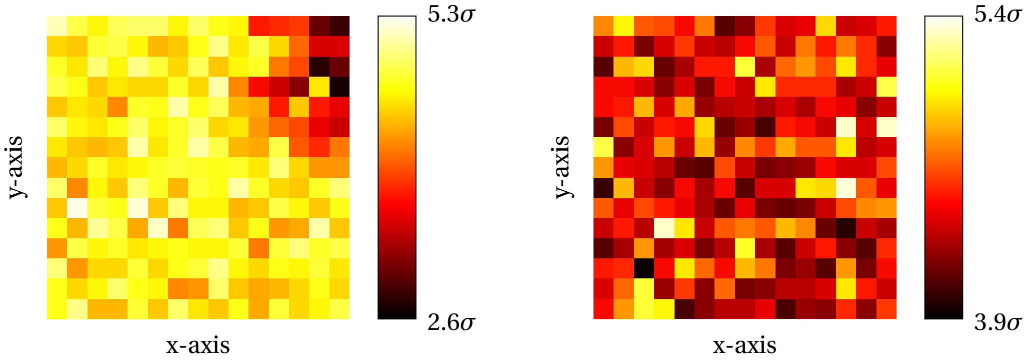

(a) Time domain

(b) Frequency domain

Figure 5.26: Frequency and time-domain TVLA tests in σ on the SF2 Advanced using SubBytes at 2 cm height.

The results from the *SubBytes* intermediate can be seen in Figure 5.26. The results in the time domain show a similar pattern as the results from 1 cm. Most positions exceed the 4.5σ except the top right corner. The overall standard deviation is lower than the *MixColumns* intermediate, as expected. The frequency domain results are similar to the results at 1 cm and are lower on average. Compared to the *MixColumns* intermediate, the frequency domain shows similar false positives and a very high noise floor.

5.4. Conclusion

Based on the baseline TVLA results we can conclude that all platforms are leaking secret information from the surface up to 2 cm of height. This can be explained by the fact that no countermeasures are implemented and the SNR of the set-up at small distances is high enough. But even though TVLA results show results exceeding the 4.5σ for all targets the σ decreases when the distance increases.

For all platforms we can conclude that the *MixColumns* intermediate leaks more than the *SubBytes* intermediate. This is caused by the T-tables implementation of the AES algorithm, which does not directly use the *SubBytes* intermediate. But since the *MixColumns* operation is skipped in the last AES round, the *SubBytes* intermediate is used in round 10. In the case of an attacker which cannot provide chosen plaintext inputs, the *MixColumns* intermediate requires2^{32} key guesses for DEMA or CEMA attacks.

While the *SubBytes* intermediate only requires 2^8 key guesses. Therefore CEMA attacks based on the *SubBytes* intermediate of round 10 are assumed to receive the best results.

When looking at the results in time domain in general both the Pynq and the Microsemi SF2 Starter kit have problems with misalignment. Which could make far-field attacks much more difficult, since these attacks are performed in the time domain. Smart-triggering could solve the problem of misalignment which are caused by the inaccuracy of the GPIO trigger, because the jitter could be excluded from the pattern trace. The Pynq target also has jitter in between rounds, caused by memory accesses, which could cause problems with smart-triggering due to the fact that the pattern changes between traces.

The Pynq target clearly shows leakage from the surface up to 2 cm of distance in both the time- and frequency domain. The frequencies at which the Pynq are leaking the most are around 400 and 800 MHz, which are modulated on the DDR memory frequency. Both frequencies are interesting for far-field analysis, but from the measurements at 1 cm distance the 800 MHz frequency is not exceeding the 4.5σ. The positioning of the probe also becomes less important when the distance to the target increases. This confirms the measurements from Ramsay [45], which showed that the Pynq target leaks secret information in the far-field around 400 MHz.

The Microsemi SF2 Basic shows also leakage from the surface up to a distance of 2 cm in both time- and frequency domain, although the 1 cm and 2 cm results are only just exceeding the 4.5σ. The overall results from the Microsemi SF2 compared to the Pynq are much lower and most of the secret information is leaked at very low frequencies due to direct emanations. Around 987 MHz also the 4.5σ is just exceeded, which due to its higher frequency is much easier to attack in the far-field. Since the TVLA results show a much lower leakage than the Pynq especially around the higher frequencies, it would seem that this target is much harder to attack in the far-field.

The last target, the Microsemi SF2 Advanced, also shows leakage from the surface up to a distance of 2 cm in both time- and frequency domain. But the overall results are worse than the Microsemi SF2 Basic, except for the time domain results because of the alignment problems. The Microsemi SF2 Advanced only shows direct emanations and frequencies around 100 MHz exceeding the 4.5σ, which makes far-field attacks extremely hard. Based on the baseline results this target seems to be the hardest to attack.

6 Far-Field Results

As defined by De Mulder [8] the EM far-field as $r > \lambda/(2\pi)$, where r is the distance from the target and λ is the wavelength. The distance of the target r decreases when the frequency increases. This means that the far-field definition is depending on the frequency of the EM radiation. Using the speed of light 299792458 m/s in vacuum, the wavelength is calculated based on the measured frequency. The experiments are not performed in vacuum, but the effects on r are not significant for the performed measurements at these short distances.

To be able to perform far-field TVLA tests and CEMA attacks we make use of the SDR set-up. Measurements will be made in the EM near field at 2 cm distance, to verify the SDR set-up for each target and to be able to compare the results to the baseline measurements. Next to that far-field TVLA tests are performed and CEMA attacks are attempted and described in this section.

For each of the targets different attack strategies will be used based on the baseline results from the previous chapter. The far-field results make use of SDR, which makes it important to select a baseband frequency at which secret information is leaking, since the SDR can only record a small bandwidth compared to an oscilloscope. TVLA tests in the frequency domain were performed during the baseline measurements, which are used to select baseband frequencies for the SDR set-up. Analog filtering is selected based on the FFT analysis of the baseline results, to improve the SNR and eliminate unwanted frequencies in the recordings. Frequencies above 400 MHz are chosen to reduce the antenna size and make the attack portable, meaning that the attack set-up is capable of fitting in a backpack. Antennas ranging from 400 MHz up to 2600 MHz are available, which further limits the possible attack frequencies in the far-field.

The RTL-SDR is configured at a rate of 2.4 MSa/s with IQ samples for all far-field measurements to ensure that no samples are lost during the recording. Problems with recovering traces from the recording could occur when higher sample rates are selected and samples are dropped. For the attacks all traces are needed and alignment in between traces must be guaranteed, which could be impossible when samples are lost. This sample rate will result in a bandwidth of 2.4 MHz, but the effective bandwidth which is recorded is around 80% of 2.4 MHz caused by internal filtering and aliasing. This means that the the leaking frequency must be within 1.92 MHz of the frequency. For measurements with the USRP the sample rate is configured at 4 MSa/s with IQ samples and a respective bandwidth of 4 MHz, which is the highest sample rate the USRP does not drop any samples with our current set-up. Similar to the RTL-SDR, the USRP has an effective bandwidth of around 80% of 4 MHz caused by internal filtering and aliasing, which means that the effective bandwidth is 3.2 MHz. The RTL-SDR and USRP are compared, to verify if the more expensive USRP is needed for far-field attacks.

All measurements are performed in normal office conditions without any shielding, unless stated otherwise. These recording will produce results as close to real-world situations as possible, where external signals are interfering with the measurements. Measurements are also performed inside an anechoic room, which provides results in a more ideal scenario where almost no external interference is present. The set-up inside the anechoic room can be seen in Figure 6.1. These measurements will provide an

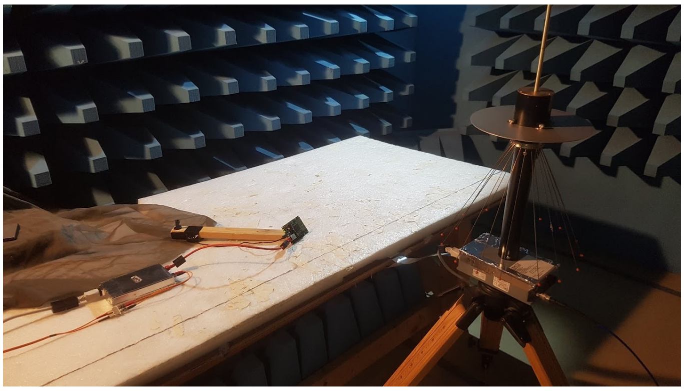

Figure 6.1: Set-up inside the anechoic room with the Microsemi SF2 and the discone antenna at 50 cm. The Ethernet connection is converted to fibre to avoid EM leakage to outside the room.

indication of the average and best conditions for performing a CEMA or DEMA attack in the far-field using SDR.

For all measurements using SDR the input data is send to the device over UDP using a python script. This input data is an array of 16 randomly generated bytes, or TVLA based input on the selected intermediate, and recorded together with the output data send back from the device. The input- and output data is saved to a Matlab file and later combined with the correct traces. During this communication the SDR is recording the I and Q samples and saves them in binary format to a file. These files are then together loaded into Matlab and used for further analysis and processing, which is described in the previous sections.

Since the SDR set-up does not make use of an artificial GPIO trigger, a smart trigger must be selected to generate traces from the EM recording. Based on the visual analysis of the EM recording an appropriate smart trigger is selected and evaluated. The advantages and disadvantages of the different triggers will be discussed in this section. The GPIO trigger is not disabled in the FPGA or software, and will generate a pulse on the GPIO port. But no connection between the target device and the computer will be made except for the Ethernet cable, which is used for communication. This means that it is still possible to record EM radiation from the communication or the GPIO pin, which is also plausible in real-world situations. For example due to a 'busy' or 'enable' line of an cryptographic engine.

TVLA tests and CEMA attacks will be performed after all traces are successfully extracted from the EM recording and combined with the input and output data. For these CEMA attacks the results from the baseline are used to optimize the EM leakage. The results from the baseline measurements show at which position in time the leakage is occurring and which leakage model is expected to have the highest correlation. Verification of the results from the near-field are performed using the new SDR set-up. Further analysis is done to verify the portability of assumptions from the near-field EM analysis to the far-field EM attacks.

For each of the three targets TVLA tests and CEMA attacks will be attempted at various distances using SDR. The measurement set-ups used and the results of these attacks will be discussed in this section for each of the three targets.

6.1. Pynq

The baseline results from the Pynq showed the most leakage in the frequency based TVLA tests around 800 MHz on the surface, but from further distances the highest leakage was around 400 MHz. The antennas available for the far-field cover both of these frequencies and both frequencies will be evaluated in this section. The far-field of 400 MHz starts at $r > 11.93$ cm in a vacuum and for 800 MHz the far-field starts at $r > 5.96$ cm in a vacuum. Other frequencies like the harmonics of the IC frequency were also analysed, but did not give any significant result. First we will look at the measurements using a Langer RF-R probe at 2 cm, which for both frequencies is defined as the near-field. Then we will look at the far-field results and verify if far-field EM side-channel attacks against the Pynq are possible in real-world conditions.

6.1.1. Langer probe at 2 cm

For these measurements a Langer RF-R probe is used similar to the baseline measurements at 2 cm. No external analog amplification is added, because of the internal SDR amplifiers which can amplify the signal in the measurement range of the SDR. The Langer probe is positioned 2 cm above the target IC surface and 200000 traces are recorded based with the same TVLA input as the baseline results. Based on the baseline measurements the amount of traces needed for an attack should not exceed the 200000 traces. The positioning of the probe is placed as closely as possible to the highest leaking spot from the baseline measurements, but due to the difference in measurement set-up this is only an approximation by hand. Because both the 400 MHz and the 800 MHz showed leakage in the baseline results, both frequencies are evaluated as frequency for the SDR. For smart-triggering different approaches are evaluated based on the measured EM radiation and the availability of a repeating pattern or peak.

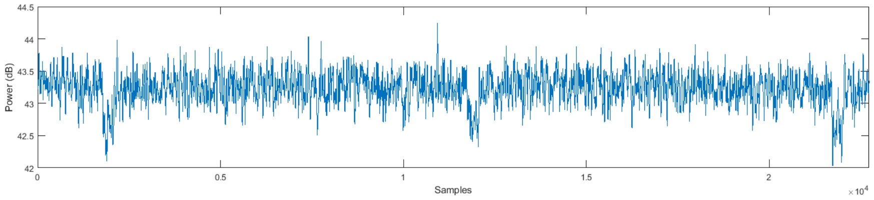

Figure 6.2: Pynq EM recording at 400 MHz frequency from 2 cm distance showing a clear AES trace pattern.

As can be seen in Figure 6.2 a clear repeating pattern is visible in the measured EM radiation of the 400 MHz frequency measurement. Since these patterns do not show a clear peak compared to the average measured EM radiation, the simple peak finding algorithm used by Ramsay [45] could not find all the needed traces. The digital moving average filtering of the signal improved the peak detection, but due to the low SNR it was still not possible to retrieve all traces. This could be explained by the fact that the pattern is identifiable and the peak finding algorithm triggers also on false positives. Instead the cross-correlation smart-trigger was used, which could successfully extract all traces when the correct threshold and pattern was selected. Since cross-correlation takes into account multiple samples and also peaks inside the pattern it did not trigger on any false-positives. Compared to the results from the SAD smart-trigger no significant difference in the detection of the traces could be found, which could be explained by the the fact that the patterns of the AES encryption had a small deviation.

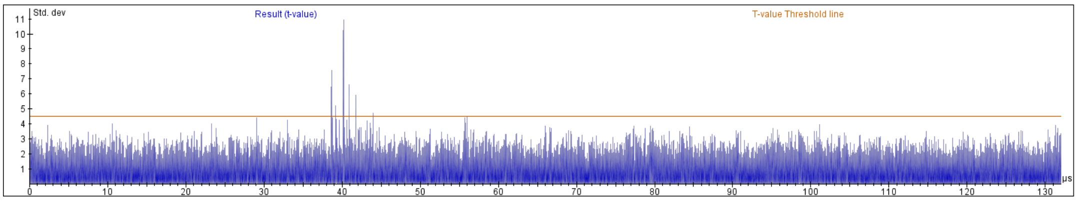

Figure 6.3: Pynq TVLA test at 400 MHz frequency from 2 cm distance, where a clear leakage is measured using the *MixColumns* round 5 intermediate.

After extracting all traces and performing a second alignment using Inspector a TVLA analysis is performed which can be seen in Figure 6.3. This clearly shows that the target is exceeding the 4.5σ at the expected point in time, around round 5, for the 400 MHz frequency. This means that our set-up is functioning correct and that even with a bandwidth of 2.4 MHz significant leakage is measured. Further analysis by performing a CEMA attack on the 10th round *SubBytes* intermediate of AES, shows that the full key can be recovered with only 80000 traces. This verifies the assumption that the target is leaking secret information.

When looking at the measurements from the 800 MHz frequency in Figure 6.4, no repeating pattern is visible in the measured EM traces. Auto-correlation is a technique which could reveal repeating patterns. While performing auto-correlation on a part of the 800 MHz frequency recording there were no

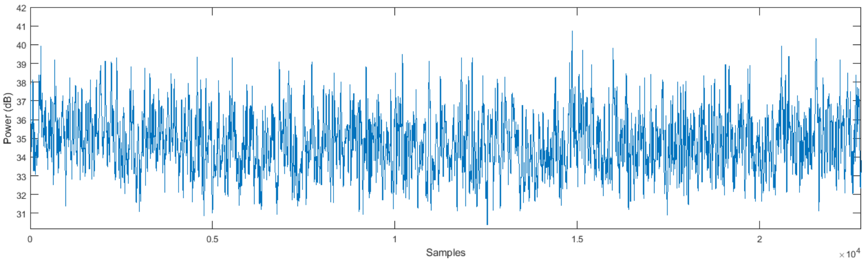

Figure 6.4: Pynq EM recording at 800 MHz frequency from 2 cm distance missing a clearly visible AES pattern.

new patterns revealed. This means that smart-triggering can not be used to extract the traces from the recording and also that the Pynq can not be attacked using the SDR set-up at 800 MHz baseband frequency without an artificial trigger. These results could be explained by the SNR of the leakage signal and the fact that the baseline TVLA measurements from 2 cm also did not show any leakage.
Verification with the USRP at 4 MSa/s with IQ samples and a bandwidth of 4 MHz for both the 400 MHz and 800 MHz frequency was performed. For all of the above measurements similar results are achieved with the USRP. Only at 400 MHz the amount of traces needed to recover the full key using a CEMA attack is slightly decreased, with the USRP only 77000 traces were needed to recover the full key. The amount of traces needed is only decreased by 3000, which could be explained by a higher SNR or a change in conditions in-between changing the set-ups. This means that the bandwidth and sample speed could only slightly increase the CEMA attacks performance, but did not give any improvement for smart-triggering.

6.1.2. Far-field radiation

For the measurements in the far-field both the 400 MHz and the 800 MHz baseband frequencies are analysed, starting with a distance of 15 cm. For both frequencies the WA5VJB 400-1000 MHz antenna is used, since it includes both frequencies. No appropriate filtering was available and the FFT of the Pynq did not show any interfering frequencies, which is why no extra filtering was added to the analog system. All measurements are made with 1 million traces using TVLA input on the round 5 *MixColumns* intermediate. More traces could be recorded, but if based on the TVLA test significant leakage is measured at exactly 1 million traces the CEMA attack typically would require more traces, depending on the leakage model used. The recording of 1 million traces for the Pynq target takes 5 minutes, which would mean a CEMA attack would at least take 10 minutes to record. Recording more traces increases the processing time of the analysis step and environmental changes in between traces could affect the performance of the attack. Therefore we have chosen to limit the recording time to 10 minutes.
The measurements from the 400 MHz baseband frequency are showing a clear increase in measured EM power when the device is performing AES encryptions. But similar to the results at 800 MHz baseband frequency from 2 cm, no visible repeating pattern is identified from the measured EM radiation. The results from an autocorrelation also did not show any repeating pattern. When the distance is decreased to less than 10 cm the repeating pattern, similar to the measurements from 2 cm, becomes visible. This indicates that smart-triggering is only possible up to 10 cm of distance, which for 400 MHz is considered the near-field. Since these results are extremely close to the far-field we will continue to analyse the recorded data from a 10 cm distance. Similar to the results from 2 cm correlation smart-triggering was used, which could retrieve all traces. The peak finding algorithm was also not successful at a 10 cm distance, as expected from the 2 cm results and the decreasing SNR due to distance. The SAD smart-trigger achieved similar results to the correlation smart-trigger, which can be explained due to the fact that the pattern is not changing when the distance to the target is increased.
When performing a TVLA test at 400 MHz from 10 cm after extracting all the traces, it is clear that that the target is exceeding the 4.5σ at the expected moment in time. The results of the TVLA test can be seen in Figure 6.5. After performing a CEMA attack on the same traces the full key can be recovered with only 210000 traces. This verifies the TVLA test and confirms that the target is leaking the secret

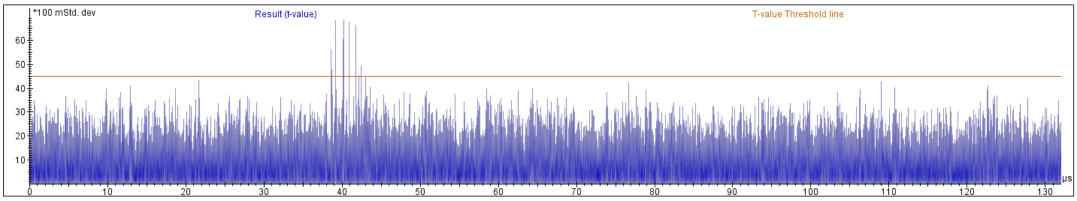

Figure 6.5: Pynq TVLA test at 400 MHz baseband frequency from 10 cm distance, where a clear leakage is measured using the *MixColumns* round 5 intermediate.

key. Even though the measurement was not performed in the far-field, it is only 2 cm from the far-field definition. The decrease of the SNR which can also be seen in the TVLA results from 2 cm and 10 cm, explains why the attacks can not be performed in the far-field. Increasing the distance with another 2 cm would even further decrease the SNR, which makes it impossible to distinguish the noise from the trace and make use of smart-triggering.

In order to improve the measured results we have performed measurements in an anechoic room at multiple distances for the 400 MHz baseband frequency, starting at a 15 cm distance. Because of the higher SNR and lower interference of outside signals, similar results are achieved as the office environment measurements from 10 cm. The correlation smart-trigger was used and could retrieve all traces. Similar to the results from 10 cm in the office environment, also the SAD smart-trigger could retrieve all traces and the peak finding algorithm could not be used as a smart-trigger. The TVLA results showed a maximum of 6.7σ and the full key could be retrieved using a CEMA attack against the 10th round *SubBytes* intermediate with 261000 traces. Compared to the results from 10 cm in the office environment only 51000 extra traces were needed. This proves that far-field attacks are possible in a controlled environment without an artificial trigger. Even though the anechoic room does not emulate a real-world environment an increase in SNR with less interference could potentiality be possible. Further increasing the distance from the target in the anechoic room made it impossible to find a repeating pattern to retrieve all recorded traces. This means that inside an anechoic room the Pynq could only be attacked up to 15 cm.

For the measurements at the 800 MHz baseband frequency in office environment, no repeating pattern could be detected using autocorrelation or human identification. Even at distances up to 6 cm no pattern was found and attacks could not be achieved. Which could be explained by the further decrease of SNR when the distance increases, compared to the results of 2 cm. This made it impossible to retrieve the traces and perform TVLA tests or CEMA attacks using the SDR set-up without an artificial trigger. Measurements with 800 MHz baseband frequency were are performed inside the anechoic room, but achieved similar results at 6 cm. This means that we can conclude that far-field attacks at 800 MHz are not possible for the Pynq target.

In the far-field the USRP was also evaluated for the office environment measurements at 400 Mhz and 800 MHz baseband frequencies. Similar to the near-field results the USRP only slightly improved the CEMA attack at 10 cm distance where only 257000 traces were needed to recover the full key. The difference of 5000 traces could be explained by the higher SNR of the USRP compared to the RTL-SDR caused by the device. No improvements for smart-triggering were measured and thus a higher bandwidth or sample rate could also not achieve EM far-field side-channel attacks in the office environment.

Comparing the results to the results from Ramsay [45], we see that that without the trigger we can only retrieve the traces up to a distance of 15 cm in an anechoic room. While the target is leaking secret information up to a 1 m distance in an anechoic room, according to the results from Ramsay [45]. This means that the smart-trigger could be a decisive factor in the maximum leakage distance for the Pynq target.

6.2. Microsemi SmartFusion2

The baseline results from the Microsemi SF2 stater kit showed mostly direct emanations at very low frequencies, which can not be captured by our SDR set-up due to the antennas and SDR frequency ranges. Also around the 987 MHz frequencies were found which exceeded the 4.5σ and can be recorded by our SDR set-up. Based on previous experience from Ramsay [45] also other harmonics of the IC frequency are analysed in the far-field, which are within the range of our antennas. The most interesting frequencies that were measured are 423 MHz, the third harmonic of the IC frequency, and at 987 MHz, the 7th harmonic of the IC frequency. The far-field of 423 MHz starts at $r > 11.28$ cm and for 987 MHz at $r > 4.83$ cm in a vacuum. First we verify the baseline results at 2 cm with the new SDR set-up and then continue with far-field TVLA tests and CEMA attacks.

6.2.1. Langer probe at 2 cm

For these measurements a Langer RF-R probe is used similar to the baseline measurements at 2 cm, without any external amplification. The TVLA input is provided to the target and 20000 traces are generated and recorded using the RTL-SDR. Based on the results from the near-field the amount of traces needed will not exceed the 200000 traces. No extra filtering is added, to provide a similar set-up as the baseline. The positioning of the probe is placed as closely as possible to the highest leaking spot from the baseline measurements.

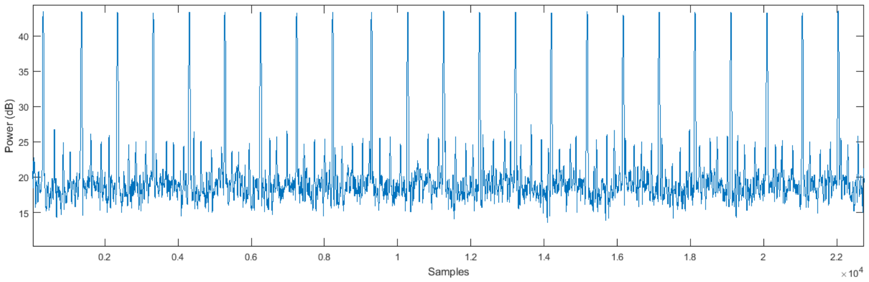

Figure 6.6: Microsemi SF2 EM recording at 987 MHz baseband frequency from 2 cm distance showing a repeating pattern.

Starting at the most promising frequency from the near-field, the 987 MHz baseband frequency, we can clearly see the repeating pattern of a single spike in Figure 6.6. A simple peak finding algorithm is used in MATLAB [32] which is capable of finding all AES traces. This means that more advanced smart-triggering such as correlation or SAD is not needed for the Microsemi SF2 at this frequency. In order to evaluate and compare the different smart-trigger both the correlation and SAD smart-triggers are evaluate. Both the cross-correlation and SAD are also capable of retrieving all traces from the recording. But compared to the peak finding algorithm, both the correlation and SAD smart-trigger give feedback about the quality of the match with the pattern. This feedback is used to verify the SNR between real traces and false-positives, confirming that the traces are successfully extracted. With the peak finding algorithm only the derivative could give feedback about a false-positive trigger. Due to the inconsistency of the connection between the PC and the target the variation of the AES encryption timing could exceed multiple milliseconds similar to a missed trigger found by the derivative. This is why we chose to use the cross-correlation smart-trigger instead.

When looking at the TVLA results in Figure 6.7 from 2 cm at 987 MHz baseband frequency, the 4.5σ threshold is not exceeded. This means that the target is not leaking any secret based on our leakage model at 987 MHz from a distance of 2 cm using the RTL-SDR set-up. This is verified by performing a CEMA attack using 1 million traces, where not a single key byte could be recovered. These results can be explained by the leakage signal, the lower sample rate or the lower bandwidth of the RTL-SDR set-up compared to the baseline set-up or the misalignment in between traces in the time domain. A correctly aligned trace set has the measured leaking sample at the same time for each trace. For a proper alignment in the time domain, the trigger must be consistent and a lower sample rate could cause aliasing problems or ADC quantisation noise preventing correct alignment.

To analyse if a higher bandwidth of 4 MHz or a sample rate of 4 MSa/s is capable of retrieving secret information the RTL-SDR is exchanged with the USRP. Also with the USRP 1 million traces are recorded at 987 MHz with a distance of 2 cm. Similar to the RTL-SDR the cross-correlation smart-triggering is used, but the peak-finding algorithm could also successfully extract all the traces. These TVLA results also did not exceed the 4.5σ and a CEMA attack was also attempted without any successful results.

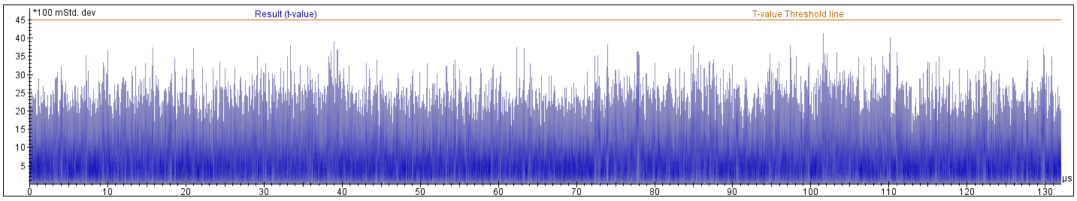

Figure 6.7: Microsemi SF2 TVLA test at 987 MHz baseband frequency from 2 cm distance showing no significant leakage exceeding the 4.5σ.

After human analysis of the spectrum using the SDR the third harmonic at 423 MHz was selected as a potential leakage frequency. This was based on the change in spectral intensity between AES encryptions and idle time of the target device, which showed a significant difference in the FFT analysis. Similar to the results at 987 MHz the peak finding algorithm was capable of retrieving all traces, but the cross-correlation smart-trigger was used to verify the traces.

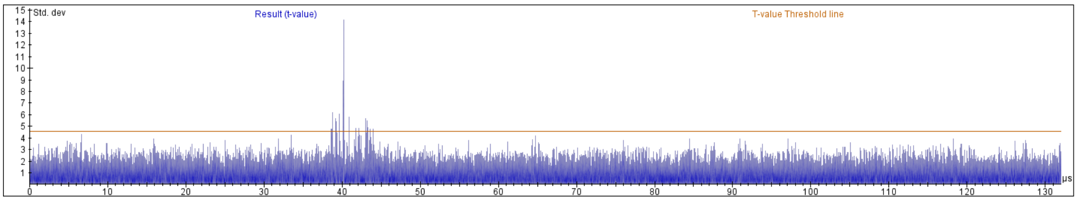

Figure 6.8: Microsemi SF2 TVLA test at 423 MHz frequency from 2 cm distance showing significant leakage exceeding the 4.5σ caused by the *SubBytes* round 5 intermediate.

The result from the TVLA test at 423 MHz baseband frequency are shown in Figure 6.8. This shows a very significant peak with a standard deviation of 14.2σ, which is exceeding the 4.5σ. Based on the baseline results this peak was not expected, since the frequency TVLA test did not show a leakage exceeding the 4.5σ at this frequency. This could be explained by the misalignment of the antenna in between tests, or the introduction of noise by the FFT during the baseline measurements. Based on the same input traces a CEMA attack is performed against the round 10 *SubBytes* intermediate, to verify if the leakage results could be used for an attack. This attack could successfully extract the full secret key with only 18000 traces, which correlates clearly with the results from the TVLA test.

Instead of the RTL-SDR, also for the tests at the 423 MHz baseband frequency the USRP is used as comparison. The TVLA results from the USRP showed a maximum deviation of 14.1σ and a CEMA attack required 20000 traces to successfully recover the full key. This is a slight decrease in performance, which is most likely caused by the margin of error in between the two different measurements and interference of external signals. Since the measurements were performed in an office environment, small changes to the environment could cause the external EM radiation to change, and potentially lead to better or worse results. Since both the TVLA result and the amount of traces are not significantly higher than the results from the RTL-SDR, we can conclude that the increase in bandwidth or sample rate did not improve our overall results.

6.2.2. Far-field radiation

For the measurements in the far-field both the 423 MHz and the 987 MHz baseband frequencies are analysed at various distances. For both frequencies the WA5VJB 400-1000 MHz antenna is used, which covers both baseband frequencies. Since specific filtering for these frequencies is not available and the bandwidth of the SDR's are relatively low, no analog filtering is added. The antenna is aimed by looking at the real-time FFT of the SDR and for all measurements recordings are made with 1 million traces using TVLA input. The measurements from the 987 MHz baseband frequency and 423 MHz baseband frequency are discussed in this section.

Measurements with the 987 MHz bandwidth are made at a distance of 5 cm in the office environment, which is considered to be the far-field. Both the peak-finding and SAD/correlation smart-triggers are capable of recovering all traces and for verification the correlation smart-trigger will be used. The TVLA test show no leakage exceeding the 4.5σ, similar to the results from 2 cm distance. This is expected, because of the lower SNR caused by the increase in distance. When a CEMA attack is performed, not a single key-byte is successfully extracted, as expected from the TVLA results. But to analyse the smart-triggering possibilities at 987 MHz the distance is increased up to 1 m, which was the maximum distance of our set-up. From this distance all traces could easily be recovered by both smart-triggers and the peak-finding algorithm. This shows that for the Microsemi SF2 no artificial trigger is needed for distances up to 1 m.

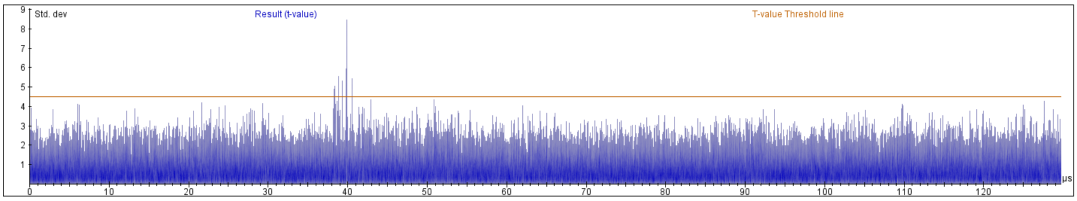

Figure 6.9: Microsemi SF2 TVLA test at 423 MHz baseband frequency from 15 cm distance showing significant leakage exceeding the 4.5σ caused by the *SubBytes* round 5 intermediate.

Further measurements are performed at the 423 MHz baseband frequency, starting at a distance of 15 cm. Also at this distance and frequency all traces could be extracted using both smart-triggering and peak-finding methods and again the cross-correlation smart-trigger is used for verification. The TVLA results can be seen in Figure 6.9, which shows a clear maximum leakage of 8.8σ exceeding the threshold line. Compared to the results from 2 cm distance, the maximum leakage value has dropped significantly, which is expected when the distance increases and the SNR decreases. The CEMA attack could successfully recover the full key with 781000 traces, which means that far-field side-channel attacks without an artificial trigger are possible in an office environment for the Microsemi SF2 target.
When looking at the average of 1000 traces we can see that the SNR of the used trigger is extremely high, which can be seen in 6.10. After further analysis using an oscilloscope, we can conclude that the trigger is caused by the Ethernet communication with the computer. From further distances than 15 cm this peak could also be used as a smart-trigger for distances up to 1 m. The SNR will drop when the distance increases, but even in the normal office environment it can be used as a trigger, instead of an artificial trigger.

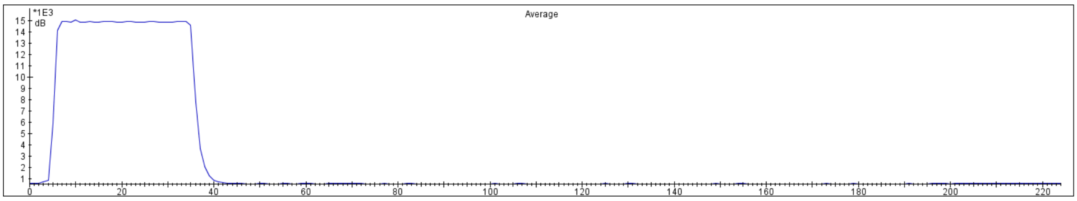

Figure 6.10: Microsemi SF2 average of 10000 traces at 423 MHz baseband frequency from 15 cm distance showing how significant the trigger is compared to the AES encryption. The AES encryption starts around sample 40.

Measurements further than 15 cm in office environment are performed, but the TVLA results did not show any leakage of secret information. In Figure 6.11 a zoomed in average trace is shown after the measured peak, which was used for smart-triggering. A pattern of the AES encryption is visible in this trace, starting with the key scheduling and after that 5 repeating patterns representing the 10 AES rounds. This pattern is similar to the results from Ramsay [45] with on surface measurements in Figure 3.11. When the distance is increased above 15 cm this pattern is not visible any more and the TVLA tests maximum deviation is not exceeding the 4.5σ. This could be explained by the decrease of SNR caused by the increase of distance and the capabilities of the SDR capturing both the trigger and AES with the same gain. CEMA attacks which are performed at further distances than 15 cm did not retrieve a single key byte. The results show that far-field side-channel attacks against the Microsemi SF2 are possible up to a distance of 15 in the office environment with the current set-up.

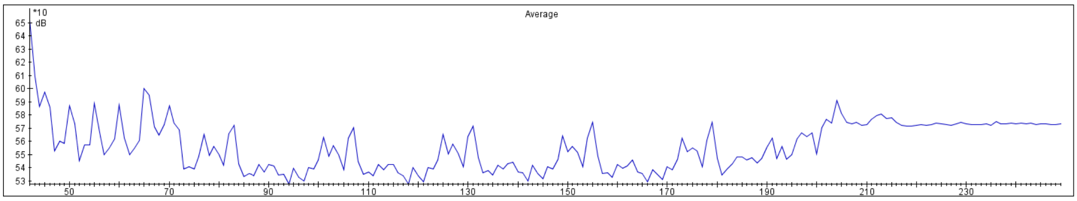

Figure 6.11: Microsemi SF2 average of 10000 traces at 423 MHz baseband frequency from 15 cm distance showing the AES encryption in more detail.

In order to compare results with higher sample rates and bandwidth, measurements were made using the USRP with a bandwidth of 4 MHz and sample rate of 4 MSa/s from a distance of 15 cm. All traces could be recovered with the smart-trigger and the TVLA test showed a maximum deviation of 8.9σ. The CEMA attack could also retrieve all key bytes with 767000 traces, which is slightly less than with the RTL-SDR. This shows that higher sample rates and bandwidth can improve the amount of traces needed to attack the Microsemi SF2. The higher sample rate and bandwidth improves the SNR and trace alignment, which explains the higher TVLA deviation and lower amount of traces for a CEMA attack. Increasing the distance and exceeding the 15 cm showed similar results as the RTL-SDR, on which we can conclude that the higher sample rate and bandwidth can not increase the attack distance for CEMA attacks. Even though the USRP has a higher SNR, this is not significant for increasing the distance.

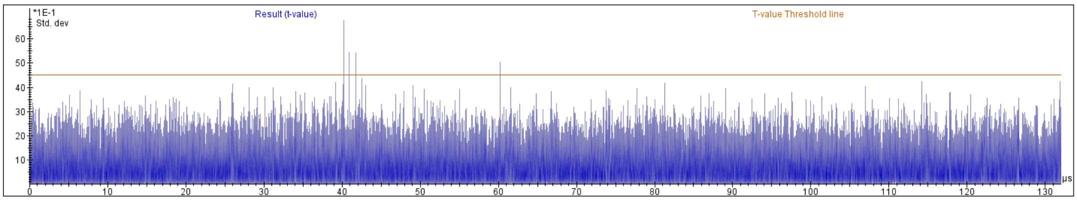

Figure 6.12: Microsemi SF2 TVLA test at 423 MHz baseband frequency from 50 cm distance in an anechoic room showing significant leakage exceeding the 4.5σ caused by the *SubBytes* round 5 intermediate.

For the Microsemi SF2 measurements were made using the RTL-SDR inside an anechoic chamber, which removes the environmental noise. Measurements up to 1 m were performed, but only up to 50 cm the TVLA test exceeded the 4.5σ. Compared to the office environment an increase in SNR is measured, caused by the lower noise floor inside the anechoic chamber. The TVLA test results from 50 cm distance inside the anechoic room can be seen in Figure 6.12. It shows only some slight peaks exceeding the 4.5σ line and a maximum standard deviation of 6.7σ is measured. The CEMA attack could retrieve the full key with 913000 traces, which is close to our total amount of measured traces and could explain why measurements from further distances were not successful. These results show that in more ideal scenario, an anechoic room, even far-field EM side-channel attacks up to 50 cm can be achieved against the Microsemi SF2.

The baseline frequency based TVLA test results did not show all frequencies on which the target was leaking in the far-field. Especially the 423 MHz was not visible during the baseline measurements, while the far-field results show the best results are 423 MHz. The positioning of the probe could change the measured field of the PCB, causing certain leaking frequencies to appear. The baseline measurements make use of an FFT, which add up all the samples inside a trace with similar frequencies. This means that the baseline TVLA measurements, which are performed in the frequency domain contain much more noise compared to time based TVLA tests. This shows that baseline measurements from the near-field can not identify all leakage frequencies which can be exploited in the far-field.

6.3. Microsemi SmartFusion2 Advanced

The baseline results from the Microsemi SF2 Advanced shows mostly direct emanations and modulated emanation around 100 MHz. Since both of these EM emanations are below our minimum antenna frequency of 400 MHz, we chose to record at the harmonics of the IC frequency above 400 MHz. The lowest frequency is the third harmonic at 426 MHz, similar to the Microsemi SF2. These frequencies are selected because of previous results with the other targets and measuring the whole frequency range would take months, which is not possible for this master thesis. First we verify the baseline results with the new SDR set-up and then continue with far-field TVLA tests and CEMA attacks.

6.3.1. Langer probe at 2 cm

For these measurements a Langer RF-R probe is used similar to the baseline measurements at 2 cm. The amplification from the RTL-SDR and USRP are used and no external amplifier is connected. No extra filtering is added, to ensure the results are similar to the baseline. The input is based on the *MixColumns* and *SubBytes* round 5 intermediate and 1 million traces are recorded. All the IC harmonics up to 994 MHz are analysed and the positioning of the probe is placed as closely as possible to the highest leaking spot from the baseline measurements. This is all done to ensure that the measurements are as similar to the baseline measurements as possible.

For all the different baseband frequencies no repeating pattern was found which could be used as a smart-trigger. Further analysis with autocorrelation did not show any repeating pattern, which was similar to the analysis by hand. Because no pattern could be found it is impossible to trigger with a smart-trigger for these baseband frequencies. No comparison with the baseline results can be made, because a trigger is needed for both TVLA tests and CEMA attacks. These results could be explained by the fact that the measured leakage in the baseline TVLA tests only exceeded the 4.5σ below the 150 MHz and frequencies above 400 MHz are not correlated to the AES encryption.

The baseline results also showed that the leakage at 2 cm is significantly less than the other targets, which could further explain the inability to detect the AES encryptions triggers. Even though the RTL-SDR could go to lower frequencies than 400 MHz and maybe succeed at around 100 MHz, the far-field antennas selected for the measurements are not capable of retrieving signals from these low frequencies. Selecting other antennas with lower frequencies would significantly increase the size of our far-field radiation antenna, which makes the attack less portable and unable to fit inside a backpack. This is why we chose not to include these results.

Measurements using the USRP are performed at 4 MSa/s and 4 MHz bandwidth, from all IC harmonics from 426 MHz up to 994 MHz baseband frequencies. These results showed that there was no pattern in the measured EM radiation which could be used to extract the AES traces. Similar to the results from the RTL-SDR, the TVLA test and CEMA attacks could not be measured. This means that the increased bandwidth and sample rate in this case did not provide any improvement for achieving our smart-triggers based in the EM radiation.

6.3.2. Far-field radiation

Far-field measurements are attempted at the IC harmonic frequencies above 400 MHz, but similar to the near-field results no smart-trigger pattern could be extracted. Measurements starting from a distance of 15 cm inside an anechoic room were performed to lower the noise floor and increase the SNR. These measurements showed that in ideal conditions no repeating pattern could be found and thus the leakage of the Microsemi SF2 Advanced could not be confirmed. Based on these measurements we can conclude that the Microsemi SF2 Advanced target does not provide a significant EM footprint for smart-triggering in these frequencies. This means that a non-invasive far-field side-channel attack is not possible with the used techniques against the Microsemi SF2 Advanced development board.

6.4. Conclusion

Based on the results from this chapter we can conclude that far-field SDR EM side-channel attacks can be performed on different targets without the use of an artificial trigger. In the office environment, the maximum distance from a target on which a successful CEMA attack could be performed was 15 cm. While inside an anechoic room up to 50 cm successful CEMA attacks were achieved. For both the Pynq and Microsemi SF2 Basic far-field CEMA attacks are possible with the current set-up.

The best results were achieved with the Microsemi SF2 Basic development kit. This can be explained by the fact that an EM trigger was available, caused by the UDP communication. The communication caused an increase EM radiation at the exact frequency were also leakage was found. This made it possible to use smart-triggers to extract the traces from the recording and perform successful CEMA attacks. The limiting factor of increasing the distance for the Microsemi SF2 Basic was the actual EM leakage at that frequency. The other targets were limited in the distance because of the current set-up was not able to extract the traces using smart-triggers.

Different smart-triggers were tested for each of the targets. With the current set-up it did not matter which smart-trigger was used for extracting the traces. This could be explained by the fact that for all smart-triggers a template trace is needed, which is limiting a comparison with high noise recordings. For the Microsemi SF2 Basic even a simple peak-finding algorithm is capable of recovering all traces. But the advantage of using smart-triggers such as correlation is that misidentified triggers can easily be identified, and the retrieved traces can thus be verified.

The results from both the Pynq and Microsemi SF2 Basic targets clearly show that the attacks inside an anechoic room can be achieved at further distances compared to the office environment. This can be explained by the lower noise floor inside the anechoic room, which makes it possible to extract the traces at further distances. These results show that in more ideal scenarios, where less interference of external signals are within the bandwidth of the SDR, better results can be achieved. Changes to the set-up could also improve the SNR and make attacks possible at further distances.

Compared to the results from Ramsay [45], which also used the Pynq target, results were only achieved up to 15 cm inside an anechoic room instead of 1 m. This clearly shows that smart-triggering can be a limiting factor for achieving far-field non-invasive EM side-channel attacks. From the results from Ramsay [45] we know leakage is still present up to 1 m, but attacks are not possible with the current set-up at these distances. Further the analysis of the near-field improved the attack results in the far-field by attacking the 10th round. This shows that a good leakage model and analysis can increase the attack opportunities.

The Microsemi SF2 Advanced was not possible to attack using the current set-up. This could be explained by the fact that no template trace could be found inside the recording, and thus no traces could be extracted. Autocorrelation was attempted to achieve a higher possibility of detecting repeating patterns, but this was also not successful. From this we can conclude that with the current set-up a frequency must be found on which the target leaks and also provides a clear pattern for smart-triggering. For this platform we confirmed with the near-field analysis that the platform is leaking at multiple frequencies, but could not verify this with the SDR set-up.

7 Discussion and Future work

Many devices are connected to the internet and make use of cryptographic functions to protect the data from being read when the data is intercepted. A lot of these devices make use of the The OpenSSL Project [54] AES encryption implementation and it is known that certain implementations of AES are vulnerable to EM side-channel attacks. These side-channel attacks are performed by recording the power or EM radiation directly on the surface of a chip. This requires physical access and often modifications to the device, which limits their use in real world scenarios due to enclosures or other obstructions. The target devices can be modified by adding an artificial trigger to identify the AES encryption when recording the EM radiation, or the device might make use of a trigger such as an LED to identify the AES encryption. Both methods require physical access to the device, which can not be guaranteed in real world scenarios. No research has been performed in more realistic scenarios with further distances from the target, the EM far-field, without an artificial trigger to attack symmetric cryptographic algorithms through differential side-channels.
In this section we will revisit the research question, together with the sub-questions from the first chapter:

Can far-field EM side-channel attacks be achieved against an IC running AES in real world scenarios?

We will discuss the achieved results based on the the research question, and will provide future research directions by identifying further improvements.

7.1. Discussion

In this thesis far-field EM side-channel attacks against AES are evaluated in real world scenarios. First we will discuss how different platforms are leaking secret information in the near-field. Then we will discuss if smart-triggers can be used for far-field EM side-channel attacks when using an SDR. And as last we will discuss if information from the near-field can be used to determine the far-field leakage.
The near-field EM leakage of the Pynq, Microsemi SF2 Basic and Microsemi SF2 Advanced is analysed using TVLA tests. This showed us in which frequencies the three different targets leak secret information. For the Pynq the 400 MHz and 800 MHz are exceeding the 4.5σ and are viable frequencies for far-field attacks. The Microsemi SF2 Basic mostly has direct emanations, but the 987MHz also exceeds the 4.5σ and can easily be analysed in the far-field. The Microsemi SF2 only has direct emanations and EM radiations below 150 MHz, which are difficult to analyse in the far-field. These frequency TVLA tests shows us that each target IC leaks at different frequencies. Analysis is also performed on which intermediate value inside the AES cipher is leaking the most EM radiation. All targets showed that the *MixColumns* intermediate had the highest correlation, which is caused by the The OpenSSL Project [54] T-Tables implementation of AES. This shows that the used IC's did not have any implication on the intermediate leakage compared to other intermediates leakage, which verifies the assumptions of the leakage for the T-Tables implementation.

Smart-triggers can be used for far-field EM side-channel attacks when using SDR for specific targets. But depending on the target and the frequency of the leakage this could vary. For the Pynq only far-field side-channel attacks inside an anechoic chamber can be achieved, because the smart-triggers could not be extracted in the far-field measurements from the office environment. The results showed that in the near-field smart-triggers can be used to attack the Pynq IC. For the Microsemi SF2 Basic far-field side-channel attacks could also be performed with smart-triggers in an office environment. It also clearly showed that certain frequencies could provide a clear trigger, but did not leak any secret information. The last IC is the Microsemi SF2 Advanced which could not be attacked using smart-triggers in both the near- and far-field. These results showed that the possibility of using smart-triggers varies mostly on the target IC, but is also influenced by the baseband frequency of the SDR.

The advantage of smart-triggers is shown by the fact that no artificial trigger was needed to attack both the Pynq and Microsemi SF2 basic. Different smart-triggers were tested, but no significant difference was measured with the current test set-up. Even though the peak-finding algorithm could also extract all traces successful, it can not be used to verify if the traces were extracted correctly and no false-positives were detected. Therefore the smart-triggers were chosen, which can give a good indication if the triggers are successfully extracted.

The near-field analysis using TVLA tests in both time- and frequency domain from the baseline results chapter, was used to perform the far-field side-channel attacks in the far-field results chapter. For the Pynq target a clear correlation between the near- and far-field can be seen in the leakage frequencies. This shows us that the information from the near field TVLA tests can be used for far-field EM analysis. The Microsemi SF2 Basic showed leakage in the far-field at specific frequencies, which were not identified in the near-field TVLA tests. This shows that even though the near-field measurements give a good indication of the far-field radiations, the measurements are not concluding that other frequencies are not emitting any secret information. The last Microsemi SF2 Advanced target had no successful results in the far-field because no traces could be identified using smart-triggering. This shows that no comparison between the near- and the far-field can be made for the Microsemi SF2 Advanced. The results of near-field analysis portability vary depending on the target, but give a good indication of the expected leaking frequencies.

Based on the results from the three different targets we can conclude that far-field EM side-channel attacks can be achieved in real-world scenarios, without modifying the target device. But these results vary extremely based on the target IC, leaking frequency and AES implementation used.

7.2. Future work

These far-field EM side-channel attacks against AES are, to the best of our knowledge, the first attacks performed in a real world scenario without making use of an artificial trigger. These attacks form a realistic threat when developing secure encryption for a data processing IC. There are however improvements possible which could extend the possible scenarios for the current side-channel attacks. For example attacks with intermediate missing traces or multiple processes running on the the target IC. The analysis of far-field EM leakage could also be improved to simplify the attack and achieve faster results, requiring less user input during the attack.

7.2.1. Frequency analysis

In this thesis we used the baseline frequency TVLA results to select the baseband frequencies for far-field analysis. But from the far-field results of the Microsemi SF2 Basic we have concluded that certain leakage frequencies are only measured in the far-field. The baseline frequency TVLA tests also requires extensive analysis of the target before an attack in the far-field can be achieved. In order to improve this analysis and verify the results, a more advanced frequency analysis could be developed to determine the leakage frequencies. Instead of verifying by hand using an real-time FFT when the EM field is changing and could leak secret information, tooling can be developed to automatically evaluate all measurable EM radiation frequencies.

This could lead to new discovery of frequencies radiating secret information in the far-field which are not present in the near-field. It could also lead to improvements where the near-field analysis is not required any more for far-field attacks, which could improve the speed of far-field attacks. It could also significantly improve the work flow of the attacker, requiring less input from the user and thus further automating the attack.

7.2.2. Multiple frequencies

From the far-field results of the Pynq target we have seen that smart-triggers from a distance using an identified leaking frequency is not possible for large distances. Ramsay [45] showed leakages up to 1 m distance with artificial triggers, confirming that the target is leaking at further distances. The Microsemi SF2 Basic on the other hand did have frequencies where the smart-triggering was capable of retrieving all traces. But at some of those frequencies no significant leakage was measured.

Further research can be performed in the recording of two synchronized down-converted frequencies at the same time, using SDR. Where the first frequency is used for smart-triggering, and the second frequency is used for CEMA of DEMA attacks. This could potentially lead into far-field attacks possible from further distances and other devices such as the Microsemi SF2 Advanced.

7.2.3. Trace recovery

When far-field attacks are performed using the SDR and make use of a smart-trigger, the smart-triggering can not always retrieve all recorded traces. Missing single or multiple traces, makes it impossible to align the input and output data with the measured EM radiation. Without this alignment no further analysis is possible and CEMA or DEMA are most likely not to succeed. Algorithms can be designed to identify these missing traces and correct the alignment of the input and output data with the traces. Measurements such as input- and output correlation give a good indication if the correct input- and output data is supplied with the EM trace. Combining this knowledge and performing and providing small sets of traces, the point of misalignment could be determined.

This trace recovery could improve the usability of far-field side-channel attacks for various targets and more realistic scenarios. The results could also improve the distance of far-field CEMA or DEMA attacks, when smart-triggering is not sufficient of recovering all traces. Even though the results from this thesis showed that either all traces or not a single trace could be recovered, it could be very useful when the target is not reliable computing the AES encryption and traces are missed.

7.2.4. Multiple processes

The target devices in this thesis are currently only running the AES encryption algorithm, but in most situations the IC is also performing a lot of other tasks. Further analysis should be performed to verify if smart-triggering is capable of recovering all traces when the IC is performing multiple varying tasks. Also verification should be performed how the multi-task system performs compared to a single task system when performing side-channel attacks again AES. This further analysis and verification could broaden the scope of possible targets for CEMA and DEMA attacks and would represent even more realistic real-world applications.

7.3. Concluding Remarks

The research objective was to perform far-field EM side-channel attacks in more realistic scenarios. Previous research has shown that far-field side-channel attacks against symmetric cryptographic algorithms are possible[21, 45]. But these attacks make use of artificial triggers and did not provide a full analysis of possible leaking frequencies. Also no comparisons has been made in between different platforms for far-field EM side-channel leakage. As last no analysis has been made on how near-field leakage analysis can be used for far-field attacks.

This thesis provides a full analysis of three target platforms and describes in detail the near-field leakage detected using TVLA. With this analysis the research objective of far-field EM side-channel attacks in more realistic scenarios is achieved, by porting the knowledge to the far-field. The artificial trigger used by Ramsay [45] and Kim et al. [21] is also removed and replaced with a smart-triggering system. Compared to Ramsay [45] also far-field attacks in an office environment without any shielded enclosure are achieved, which makes the attack possible for real-world applications.

The results from this thesis shows that even in an office environment the Microsemi SF2 Basic development board can be attacked non-invasive up to a distance of 15 cm, where only the EM radiation needs to be recorded. The SDR attack set-up with small PCB antenna is capable of fitting inside a backpack, which makes this attack more realistic and stealth. Inside an anechoic room even attack distances up to 50 cm are achieved, which means that in almost ideal scenarios without any outside noise this result can also be achieved in real-world applications.

Depending on the capabilities of the attacker, the environment and the target platform, it is clear that the AES cypher can be attacked from the far-field. This means that cryptographic devices must protect themselves against these attacks in order to prevent an attacker from stealing the secret key. This thesis showed that these attacks are possible without having access to the target device or the removal of the target enclosure. Often these attacks were not considered, but the results from the thesis clearly show that far-field side-channel attacks in real-world applications are possible.

Bibliography

[1] Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao, and Pankaj Rohatgi. The em side—channel(s). In Burton S. Kaliski, çetin K. Koç, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2002*, pages 29–45, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-36400-9.

[2] G Becker, J Cooper, E DeMulder, G Goodwill, J Jaffe, G Kenworthy, T Kouzminov, A Leiserson, M Marson, P Rohatgi, et al. Test vector leakage assessment (tvla) methodology in practice. In *International Cryptographic Module Conference*, volume 1001, page 13, 2013.

[3] Arthur Beckers, Josep Balasch, Benedikt Gierlichs, and Ingrid Verbauwhede. Design and implementation of a waveform-matching based triggering system. In *International Workshop on Constructive Side-Channel Analysis and Secure Design*, pages 184–198. Springer, 2016.

[4] Eli Biham and Adi Shamir. *Differential cryptanalysis of the data encryption standard.* Springer Science & Business Media, 2012.

[5] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 16–29. Springer, 2004.

[6] Joan Daemen and Vincent Rijmen. The block cipher rijndael. In *CARDIS*, volume 1820, pages 277–284. Springer, 1998.

[7] Joan Daemen and Vincent Rijmen. Aes proposal: Rijndael. Technical report, FIPS, 1999.

[8] Elke De Mulder. *Electromagnetic techniques and probes for side-channel analysis on cryptographic devices.* PhD thesis, PhD thesis, KU Leuven, 2010.

[9] Digilent. PYNQ-Z1 Python Productivity for Zynq, 2017. URL `http://store.digilentinc.com/pynq-z1-python-productivity-for-zynq/`. [Online; Accessed December 21, 2017].

[10] Digilent. PYNQ-Z1 Zynq 7000, top view, 2018. URL `https://cdn6.bigcommerce.com/s-7gavg/products/529/images/4414/PYNQ_-_Top_-_New_Jumpers_-_Academic_-_600__89514.1488572177.1280.1280.png?c=2`. [Online; Accessed March 12, 2018].

[11] Kent Electronics. Wa5vjb, 2018. URL `http://www.wa5vjb.com/`. [Online; Accessed April 11, 2018].

[12] Emcraft Systems. SF2-STARTER-KIT, 2017. URL `https://www.emcraft.com/products/153`. [Online; Accessed December 21, 2017].

[13] Emcraft Systems. SmartFusion2 Advanced Development Kit, 2018. URL `https://www.microsemi.com/images/soc/products/smartfusion2/M2S150-ADV-DEV-KIT_Callout_Rev2.jpg`. [Online; Accessed March 12, 2018].

[14] Emcraft Systems. SmartFusion2 Starter Kit, 2018. URL `https://www.microsemi.com/images/soc/products/hardware/SmartFusion2_StarterKitBoard.jpg`. [Online; Accessed March 12, 2018].

[15] Ettus Research. USRP B200, 2017. URL `https://www.ettus.com/product/details/UB200-KIT`. [Online; Accessed December 28, 2017].

[16] FIPS. Specification for the Advanced Encryption Standard (AES). Federal Information Processing Standards Publication 197, 2001. URL `http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf`.

[17] Ronald A Fisher. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika*, 10(4):507–521, 1915.

[18] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic analysis: Concrete results. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 251–261. Springer, 2001.

[19] Benjamin Jun Gilbert Goodwill, Josh Jaffe, Pankaj Rohatgi, et al. A testing methodology for side-channel resistance validation. In *NIST non-invasive attack testing workshop*, 2011.

[20] Benjamin Jun and Gary Kenworthy. Is your mobile device radiating keys. In *RSA Conference*, volume 2, page 2012, 2012.

[21] ChangKyun Kim, Martin Schläffer, and SangJae Moon. Differential side channel analysis attacks on fpga implementations of aria. *ETRI journal*, 30(2):315–325, 2008.

[22] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *Advances in cryptology—CRYPTO'99*, pages 789–789. Springer, 1999.

[23] Markus Guenther Kuhn. *Compromising emanations: eavesdropping risks of computer displays.* PhD thesis, University of Cambridge, 2002.

[24] Susan Landau. Standing the test of time: The data encryption standard. *Notices of AMS*, 47(3): 341–349, 2000.

[25] LANGER EMV-Technik. PA 203 SMA set, 2018. URL `https://www.langer-emv.de/en/product/preamplifier/37/pa-203-sma-set-preamplifier-100-khz-up-to-3-ghz/518`. [Online; Accessed April 20, 2018].

[26] LANGER EMV-Technik. RF-R 400-1, 2018. URL `https://www.langer-emv.de/en/product/rf-passive-30-mhz-3-ghz/35/rf-r-400-1-h-field-probe-30-mhz-up-to-3-ghz/13`. [Online; Accessed April 20, 2018].

[27] LANGER EMV-Technik. RF-U 2.5-2, 2018. URL `https://www.langer-emv.de/en/product/rf-passive-30-mhz-3-ghz/35/rf-u-2-5-2-h-field-probe-30-mhz-up-to-3-ghz/11`. [Online; Accessed April 20, 2018].

[28] Thanh-Ha Le, Cécile Canovas, and Jessy Clédière. An overview of side channel analysis attacks. In *Proceedings of the 2008 ACM Symposium on Information, Computer and Communications Security*, ASIACCS '08, pages 33–43, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-979-1. doi: 10.1145/1368310.1368319.

[29] Stefan Mangard. Hardware countermeasures against dpa–a statistical analysis of their effectiveness. In *Cryptographers' Track at the RSA Conference*, pages 222–235. Springer, 2004.

[30] Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards (Advances in Information Security)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007. ISBN 0387308571.

[31] Martin Marinov. Remote video evesdropping using a software-defined radio platform. Master's thesis, St Edmund's College, November 2014.

[32] MATLAB. *version 9.2.0 (R2017a)*. Natick, Massachusetts, 2017.

[33] J. C. Maxwell. *A Treatise on Electricity and Magnetism*, volume 2. Clarendon Press, 1873.

[34] Lukáš Mazur and Martin Novotnỳ. Differential power analysis on fpga board: Boundaries of success. In *Embedded Computing (MECO), 2017 6th Mediterranean Conference on*, pages 1–4. IEEE, 2017.

[35] David P Montminy, Rusty O Baldwin, Michael A Temple, and Mark E Oxley. Differential electromagnetic attacks on a 32-bit microprocessor using software defined radios. *IEEE Transactions on Information Forensics and Security*, 8(12):2101–2114, 2013.

[36] Edwin NC Mui, R Custom, and D Engineer. Practical implementation of rijndael s-box using combinational logic. *Custom R&D Engineer Texco Enterprise Pvt. Ltd*, 2007.

[37] NooElec. RTL-SDR (RTL2832U), 2017. URL `https://www.rtl-sdr.com/`. [Online; Accessed December 21, 2017].

[38] Harry Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928.

[39] Colin O'Flynn and Zhizhang David Chen. Side channel power analysis of an aes-256 bootloader. In *Electrical and Computer Engineering (CCECE), 2015 IEEE 28th Canadian Conference on*, pages 750–755. IEEE, 2015.

[40] Syed Rafay Hasan Oluwadara Adegbite. A novel correlation power analysis attack on pic based aes-128 without access to crypto device. *Circuits and Systems (MWSCAS)*, 2017.

[41] Siddika Berna Örs, Elisabeth Oswald, and Bart Preneel. Power-analysis attacks on an fpga-first experimental results. In *CHES*, volume 2779, pages 35–50. Springer, 2003.

[42] OSPL. Tempest specialists, 2017. URL `http://www.ospl.nl/`. [Online; Accessed December 21, 2017].

[43] Eric Peeters, François-Xavier Standaert, and Jean-Jacques Quisquater. Power and electromagnetic analysis: Improved model, consequences and comparisons. *Integration, the VLSI journal*, 40(1): 52–60, 2007.

[44] Jean-Jacques Quisquater and David Samyde. *ElectroMagnetic Analysis (EMA): Measures and Counter-measures for Smart Cards*, pages 200–210. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001. ISBN 978-3-540-45418-2. doi: 10.1007/3-540-45418-7_17.

[45] Craig Ramsay. TEMPEST attacks against AES. Technical report, Fox-IT, June 2017.

[46] Ettus Research. USRP B200mini Series Architecture, 2018. URL `https://www.ettus.com/content/USRP_B200mini_BD_925x422.png`. [Online; Accessed March 8, 2018].

[47] Riscure. icWaves Datasheet, 2017. URL `https://www.riscure.com/security-tools/hardware/icwaves`. [Online; Accessed March 8, 2018].

[48] Riscure. Inspector SCA, 2017. URL `https://www.riscure.com/security-tools/inspector-sca/`. [Online; Accessed December 28, 2017].

[49] S. Markgraf and D. Stolnikov. Rtl-sdr osmosdr, 2017. URL `http://osmocom.org/projects/sdr/wiki/rtl-sdr`. [Online; Accessed 28 December, 2017].

[50] Aaron Scher. Simplified block diagram of nooelec rtl-sdr, 2018. URL `http://aaronscher.com/wireless_com_SDR/docs/rtl_sdr_block.jpg`. [Online; Accessed March 8, 2018].

[51] O. X. Standaert, E. Peeters, G. Rouvroy, and J. J. Quisquater. An overview of power analysis attacks against field programmable gate arrays. 94(2):383–394, 2006. ISSN 0018-9219. doi: 10.1109/JPROC.2005.862437.

[52] Data Encryption Standard et al. Federal information processing standards publication 46. *National Bureau of Standards, US Department of Commerce*, 4, 1977.

[53] NIST-FIPS Standard. Announcing the advanced encryption standard (aes). *Federal Information Processing Standards Publication*, 197:1–51, 2001.

[54] The OpenSSL Project. OpenSSL: The open source toolkit for SSL/TLS. `www.openssl.org`, April 2003.

[55] Pol Van Aubel, Kostas Papagiannopoulos, Łukasz Chmielewski, and Christian Doerr. Side-channel based intrusion detection for industrial control systems. *arXiv preprint arXiv:1712.05745*, 2017.

[56] Wim Van Eck. Electromagnetic radiation from video display units: An eavesdropping risk? *Computers & Security*, 4(4):269–286, 1985.

[57] Ingrid Verbauwhede. *Secure integrated circuits and systems.* Springer, 2010.

[58] Wikipedia, the free encyclopedia. Software defined radio concept, 2018. URL `https://nl.wikipedia.org/wiki/Software-defined_radio#/media/File:SDR_et_WF.svg`. [Online; Accessed March 8, 2018].

[59] Cheuk Wong. Analysis of DPA and DEMA Attacks. *Master's Projects*, 2012.