INVESTIGATING AND LEVERAGING EM AND BACKSCATTERING SIDE CHANNELS FOR HARDWARE SECURITY

A Dissertation
Presented to
The Academic Faculty

By

Frank Thompson Werner

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Electrical and Computer Engineering

Georgia Institute of Technology

August 2022

Copyright © Frank Thompson Werner 2022

INVESTIGATING AND LEVERAGING
EM AND BACKSCATTERING SIDE CHANNELS FOR HARDWARE SECURITY

Approved by:

Dr. Alenka Zajić, Advisor
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Andrew F. Peterson
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Milos Prvulovic, Co-Advisor
School of Computer Science
Georgia Institute of Technology

Dr. Gregory D. Durgin
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Morris B. Cohen
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Hyesoon Kim
School of Computer Science
Georgia Institute of Technology

Date approved: May 25, 2022

Every great advance in science has issued from a new audacity of imagination.

John Dewey

This dissertation is dedicated to my friends and family. Without their love and support, I would have never made it this far.

ACKNOWLEDGMENTS

I would first like to express my deep gratitude to my advisers, Dr. Alenka Zajic and Dr. Milos Prvulovic. Their guidance has helped me grow as an engineer and improve as a researcher. I have greatly enjoyed my time here and appreciate the opportunities they have given me. I would also like to express my gratitude to Dr. Morris B. Cohen, Dr. Andrew F. Peterson, Dr. Gregory D. Durgin, and Dr. Hyesoon Kim, for serving on my dissertation committee.

Furthermore, I would like to thank my labmates, including Jinbang Fu, Elvan Ugurlu, Moumita Dey, Mine Kerpicci, Erik Jorgensen, Andrew Kacmarick, Giannis Stamoulis, Jonas Theumer, Dr. Oluwaseun (Seun) Sangodoyin, Dr. Prateek Juyal, Dr. Baki Yilmaz, Dr. Nader Sehatbakhsh, Dr. Monjur Alam, Dr. Chia-Lin (Woody) Cheng, Dr. Alireza Nazari, Dr. Nguyen (Pavel) Luong, Dr. Sinan Adibelli, and Dr. Haider Khan. Not only have they been great collaborators, but good friends.

In addition, I would like to thank my adviser for my MS at Auburn University, Dr. Robert N. Dean, and my former labmates, Dr. Benjamin Keaton Rhea and Dr. R. Chase Harrison. I greatly enjoyed working with them.

Finally, I would like to thank the friends and family that have supported me all of this time. I would have never gotten this far without their support and I cannot express how much I appreciate them. While there are too many names to list here, I would like to specifically list my parents, Dr. Isabelle K. Thompson and Dr. Warren W. Werner. They instilled the importance of education and a love of learning. More importantly, however, their unconditional love and endless support helped me overcome multiple challenges and become the person I am today.

TABLE OF CONTENTS

Acknowledgments v

List of Tables xi

List of Figures xii

Summary xvi

Chapter 1: Introduction 1

1.1 Motivation 1

1.2 A Method for Efficient Localization of Magnetic Field Sources Excited by Execution of Instructions in a Processor 3

1.3 An Efficient Method for Localization of Magnetic Field Sources that Produce High-Frequency Side-Channel Emanations 4

1.4 Leveraging EM Side Channels for Recognizing Components on a Motherboard 5

1.5 Detection of Recycled ICs Using Backscattering Side-Channel Analysis . . 5

1.6 Research Contributions 6

1.7 Outline 7

Chapter 2: Background 9

2.1 Side-Channel Analysis 9

2.2 Electromagnetic Side Channels 9
2.2.1 Applications 10
2.2.2 Physical Sources of EM Side Channels 10
2.2.3 The Influence of Program Activity 12
2.2.4 Locating the Sources of the Side Channel: Near Field Scanning . . 13
2.3 Backscattering Side Channels 15
2.4 Hardware Authentication 16
2.4.1 Common Component Authentication Approaches 17
2.5 Recycled ICs 19
2.5.1 Mechanisms of IC Aging 20

Chapter 3: A Method for Efficient Localization of Magnetic Field Sources Excited by Execution of Instructions in a Processor 22

3.1 Overview 22
3.2 Creating System Activity at Controlled Frequencies 23
3.3 Localization of Magnetic Field Sources 25
3.3.1 Problem Statement 25
3.3.2 Forward Model 26
3.3.3 Optimization Procedure 28
3.4 Model Verification and Robustness 29
3.4.1 Verification of Localization Algorithm via Simulated Data 30
3.4.2 Localization with Unknown Number of Sources 34
3.4.3 Localization Using Simulated Data with Noise 34
3.5 Experimental Results 36

3.5.1 Measurement Setup 37

3.5.2 Devices Under Test 38

3.5.3 Localization of Sources Created by On-Chip and Off-Chip Instructions on the DE1 39

3.5.4 Localization of Sources Created by Instructions Modulated onto FPGA Processor Clock 42

3.5.5 Localization of Sources on the A13-MICRO 44

3.6 Conclusions 46

Chapter 4: An Efficient Method for Localization of Magnetic Field Sources that Produce High-Frequency Side-Channel Emanations 47

4.1 Overview 47

4.2 Localization System 48

4.2.1 Side-Channel Signal Generation 48

4.2.2 Near-Field Scanner Setup 49

4.2.3 Probe Design 52

4.2.4 Localization Algorithm and EM Source Model 54

4.3 Validation of Measurement Setup 56

4.4 Localization of Sources of High-Frequency Emanations on PCBs - Experimental Results 59

4.5 DE1 Experiments 60

4.5.1 Baseband DE1 Measurements 61

4.5.2 50 MHz DE1 Measurements 63

4.5.3 1 GHz DE1 Measurements 65

4.6 A13-MICRO Experiments 68

4.6.1 Baseband A13-MICRO Measurements 69
4.6.2 A13-MICRO Processor Measurements 70
4.6.3 A13-MICRO Memory Measurements 72
4.7 Conclusions 73

Chapter 5: Leveraging EM Side Channels for Recognizing Components on a Motherboard 75

5.1 Overview 75
5.2 Signals Carrying Information About Device Side-Channel Signatures 76
5.2.1 Spectral Features 77
5.2.2 Device Excitation 77
5.3 Signal Compression and Processing 80
5.4 Training and Testing Process 83
5.5 Experimental Validation 84
5.5.1 Measurement Setup 84
5.5.2 Test Devices 87
5.5.3 Measurement Projection 89
5.5.4 Recognition of Memory Components 93
5.5.5 Recognition of Processor Components 97
5.5.6 Recognition of Ethernet Transceiver Components 100
5.6 Conclusions 102

Chapter 6: Detection of Recycled ICs Using Backscattering Side-Channel Analysis 104

6.1 Overview 104

6.2 Understanding the Relationship Between Aging and BSCA 105

6.2.1 Advantages Over Other Detection Methods 107

6.2.2 Impact of Aging on the Backscattered Signal 108

6.2.3 Distinguishing Differences Between Unaged and Aged Measurements 111

6.3 Simulating Aging 111

6.4 Detecting Aging Using BSCA 116

6.5 Experimental Validation 119

6.5.1 Setup 119

6.5.2 Impact of Circuit Size 122

6.5.3 AES Circuit Results 123

6.6 Conclusions 125

Chapter 7: Summary and Future Work 127

7.1 Summary 127

7.2 Future Work 129

References 131

Vita 146

LIST OF TABLES

3.1 Parameters of simulated loop sources. 31

5.1 List of Tested Components 90

5.2 Average Distances Between Select Memory Components 93

5.3 Average Distances Between Memory Components 96

5.4 Confusion Matrix for the Processors (in %). 100

5.5 Average Distances Between Processor Components 100

5.6 Confusion matrix for Ethernet Transceivers (in %) 102

5.7 Average Distances Between Ethernet Transceivers 102

LIST OF FIGURES

3.1 Pseudo-code for generating the X/Y alternation activity. 24

3.2 Source space (ν) with several current-carrying loops and test points (the light blue squares). 25

3.3 (left) Local Cartesian coordinate system attached to a magnetic dipole with the corresponding spherical system; (right) magnetic field lines of a small loop. 27

3.4 Normalized intensity of the magnetic field for (a) one loop, (b) two loops, and (c) three loops. 32

3.5 Results of simplex optimization launched from randomly selected points for (a) one loop, (b) two loops, (c) three loops, and (d) three loops with an increased number of test points. 33

3.6 Results of simplex optimization when the number of assumed dipoles does not match the actual number of loops: (a) one loop and three dipoles, (b) three loops and one dipole, and (c) three loops and five dipoles. 35

3.7 Optimization for the case of three loops when the SNR is (a) $30\ \mathrm{dB}$, (b) $20\ \mathrm{dB}$, and (c) $10\ \mathrm{dB}$. 36

3.8 The measurement setup. 37

3.9 (left) Total measured magnetic field for DIV/ADD with locations of the sources; (right) locations of the sources for DIV/ADD on the DE1. 40

3.10 The measured total magnetic field and the x, y, and z components of the magnetic field generated by DIV/ADD on the DE1. 41

3.11 The simulated total magnetic field and the x, y, and z components of the magnetic field generated by DIV/ADD on the DE1 using the optimal location of the sources found by the algorithm. 42

3.12 (left) Total measured magnetic field for MUL/ADD with the location of the source; (right) the location of the source for MUL/ADD on the DE1. 43

3.13 (left) Total measured magnetic field for LDM/DIV with locations of the sources; (right) locations of the sources for LDM/DIV on the DE1. 43

3.14 (left) Total measured magnetic field for modulated DIV/ADD with locations of the sources; (right) locations of the sources for modulated DIV/ADD on the DE1. 44

3.15 (left) Total measured magnetic field for modulated LDM/LDL1 with locations of the sources; (right) locations of the sources for modulated LDM/LDL1 on the DE1. 45

3.16 (left) Total measured magnetic field for MUL/ADD with locations of the sources; (right) locations of the sources for MUL/ADD on the A13-MICRO. 45

3.17 (left) Total measured magnetic field for MUL/SUB with locations of the sources; (right) locations of the sources for MUL/SUB on the A13-MICRO. 46

4.1 The measurement setup. The DUT is the aluminum plate used in Section 4.3. "Probe 1" is the moving probe and "Probe 2" is the stationary probe. 50

4.2 Diagram of the measurement setup. 52

4.3 The shield-loop probe used for measurements. 53

4.4 Comparison of the components of the measured (top) and simulated (bottom) magnetic fields for the first test case. 58

4.5 Comparison of measured (top) to simulated (bottom) magnetic fields for the second test case. 58

4.6 Comparison of measured (top) to simulated (bottom) magnetic fields for the two-source test case. 58

4.7 The measured magnetic field in the first test case before correction. (Compare with Fig. 4.4.) 59

4.8 The (a) DE1 development board and (b) A13-MICRO with relevant components labeled. 60

4.9 The magnetic fields measured at the baseband (156 kHz) for LDM/LDL1, STM/DIV, DIV/ADD, and MUL/SUB. 62

4.10 Magnetic fields measured at the upper sideband (+20 kHz) of the 50 MHz FPGA clock. 64

4.11 Magnetic field measured at the 50 MHz FPGA clock. 65

4.12 Magnetic fields measured at the upper sideband (+20 kHz) of the 1 GHz harmonic of the FPGA clock. 66

4.13 Magnetic field measured at the 20th harmonic (1 GHz) of the FPGA clock. . 68

4.14 Example of the baseband (156 kHz) A13-MICRO measurements for LDM/ADD, LDL1/ADD, STM/ADD, and MUL/SUB. 69

4.15 Examples of the upper sideband (+20 kHz) of the 1.008 GHz processor clock for different instruction pairs. 71

4.16 Magnetic field measured at the 1.008 GHz processor clock. 72

4.17 Magnetic fields measured at the upper sideband (+20 kHz) of the 816 MHz memory clock for different instruction pairs. 73

4.18 Magnetic field measured at the 816 MHz memory clock. 74

5.1 Example of the spectrum produced by the excitation program. 79

5.2 Measurement setup used for the experiments. 85

5.3 Diagram of the measurement setup. 85

5.4 The seven IoT devices (not to scale). 88

5.5 The DE0-CV Cyclone V and DE1 Cyclone II development boards (not to scale). 88

5.6 Example of the memory measurements projected into the new feature space. 91

5.7 Comparison of the EM signatures from MEM1 (top in blue), MEM2 (second from the top in red), MEM3 (third from the top in green), and MEM4 (bottom in magenta). 95

5.8 Comparison of the MEM4 EM signatures from A20-MICRO U2 (top in blue) and A20-MICRO U3 (bottom in red). 96

5.9 Comparison of the EM signatures from PROC1A (top in blue), PROC1B (second from the top in red), PROC2A (third from the top in green), and PROC2B (bottom in magenta). 99

5.10 Comparison of the EM signatures from ETH1 (top in blue), ETH2 (middle in red), and ETH3 (bottom in green). 101

6.1 Comparison of the trapezoidal spectrum before and after τ_{r} is increased. . . 109

6.2 Comparison of the trapezoidal spectrum before and after τ_{r} is increased with a duty cycle of 49%. 110

6.3 Example of amplitude ratios from unaged and aged cases in Fig. 6.1 normalized to the unaged ratios. 112

6.4 Inverter chain example simulated in LTspice. 113

6.5 Simulated amplitude ratios as only the PMOS transistors are aged. 114

6.6 Simulated amplitude ratios as only the NMOS transistors are aged. 115

6.7 Simulated amplitude ratios as both types of transistors are aged. 116

6.8 Example of unaged and aged amplitude ratios after being projected into the new vector space. 118

6.9 The amplitude ratios of one FPGA implementing the 100% utilization shift register circuit after the equivalent of 198 days of accelerated aging. 122

6.10 ROC curves for the six shift register circuits after each round of aging. (a) 100%, (b) 50%, (c) 25%, (d) 10%, (e) 5%, and (f) 1% utilization. 124

6.11 The amplitude ratios of one FPGA implementing the AES circuit after the equivalent of 198 days aging. 125

6.12 ROC curves for the AES circuit after each day of aging. 126

SUMMARY

This dissertation is focused on two important questions related to the hardware security: (1) does the device leak confidential information and (2) are the components on the device what they claim to be. This first question and the first half of this work is related to side channels, specifically electromagnetic (EM) side channels. To help designers address and take advantage of EM side channels, two methods for locating the physical sources of EM side channels have been developed. The first method is intended to locate the low-frequency sources of the EM side channel, while the second method is intended to locate the high-frequency sources. Both methods are used to compare how the EM side-channel sources change with frequency and program activity.

The second question regards authenticating the components on the device. The same properties that make side channels such a threat, also make them useful for authentication. The second half of this dissertation introduces two methods that make use of side channels for component authentication. The first method uses EM side channels for identifying components installed on a device. Focusing on components already integrated on a device lets designers authenticate devices assembled by third parties.

While the first authentication method is effective at identifying components, it can only distinguish between components with physically different designs. It is not effective at detecting one of the most common types of counterfeits, recycled ICs. Not only do recycled ICs cost designers' money, they hurt the reliability of the devices they are integrated onto. Side channels have not been commonly used for detecting recycled ICs since most are not sensitive to the changes caused by aging. However, we take advantage of the recently defined backscattering side channel for detection. Being impedance-based, it is directly affected by the IC aging. Since the backscattering side channel requires no additional circuitry on the IC, it is low cost and more convenient than most other detection methods. The effect of aging on the side channel is investigated through simulation and experimentation.

CHAPTER 1
INTRODUCTION

1.1 Motivation

Electronic devices are integral part of modern society and we rely on them being trustworthy. Most electronic devices are comprised of a printed circuit board (PCB) with multiple components integrated on to it. Both the PCB and its components need to be secure for the device to be trustworthy. This thesis is focused on two basic questions for the hardware security of devices: (1) does the device leak confidential information and (2) are the components on the device what they claim to be?

The first question is concerned with side-channel analysis (SCA). If the device is not carefully designed, it can unintentionally leak confidential information into its environment by simply performing its normal activities. SCA takes advantage of that leakage to determine what is being done on the device. The avenue of the attack can vary; however, no matter the type of channel, these attacks are a significant threat to the security of the device. One of the most often used types of side channels are electromagnetic (EM). They can leak significant information about what is being done by the device and can be monitored using relatively cheat equipment, such as software-defined radios.

To address this leakage, designers need to be able to locate the sources on the device. Furthermore, while EM SCA is mainly thought of as a security concern, it can also be used to detect other security concerns on the device. For example, there have been several situations where the EM side channel has been used to detect to monitor the device activity and detect malware [1]–[3]. However, given how weak the side-channel signal can be, commonly it needs to be monitored as close as possible to the source. At the same time, identifying the sources of the side channel helps designers simulate the radiated emissions

generated by their design and address other security vulnerabilities, such as hardware Trojans [4]–[7].

Determining the locations of these sources can be difficult. Not only is the signal weak, but also the physical sources can change with program activity as different resources are used to execute different instructions. This dissertation introduces two methods for locating the sources of the EM side channels. The first is designed for low frequencies, while the second is for high frequencies. These methods are then used to investigate how the sources of the EM side channel change with instruction activity and frequency.

The second question regards authenticating the identity of the components on a device. Due to the significantly lower cost, designers have increasingly outsourced the manufacturing of their designs and the procurement of the integrated circuits (ICs) needed for those designs to third parties. Furthermore, the manufacturer of these ICs very rarely acts as a direct supplier to the system integrator. With so many third parties involved in manufacturing a device, there are multiple opportunities for the components to be tampered with [8].

One of the most concerning types of tampering are counterfeit ICs [9]. As the name implies, a counterfeit IC is any IC whose original design or functionality is being intentionally misrepresented to the buyer. These counterfeits can compromise the security and lifespan of the device. It is estimated that counterfeits represent 1% of all semiconductor sales [10], costing manufacturers approximately $\$100$ billion [11].

Types of counterfeits include unauthorized copies, ICs incorrectly marked or with false documentation, defective ICs, out-spec ICs, and ICs that have been tampered with (such as hardware Trojans) [9]. Even in the absence of malicious intent, one legitimate component can be (and often is) substituted with another legitimate component that the PCB manufacturer may consider to be equivalent. However, some of the properties of these devices may differ, especially when it comes to inter-operability, level of trust, reliability, bugs, and vulnerabilities.

One of the most common and difficult to detect types of counterfeits is recycled ICs.

Recycled ICs have been removed from discarded electronics and then sold as if they are new. Recycled and remarked ICs are estimated to make up 80% of all counterfeit incidents [9]. Since these ICs have been previously used, they will have a decreased lifespan, potentially causing the IC to fail prematurely. In critical applications, such as medicine, unreliable ICs can not only cost money but also jeopardize the safety of their users. Since functionally they are the same as a new IC, it is challenging to identify recycled ICs.

Given the complexity of the manufacturing process, counterfeit ICs may be impossible to avoid. Therefore, methods for detecting tampered components on an assembled device are needed. SCA provides a accurate, low-cost opportunity for detecting different types of counterfeit ICs. EM and the newly defined backscattering side channel are convenient and effective avenues for detecting different types of counterfeits. Two methods that leverage these side channels for authenticating components are discussed in the following chapters.

1.2 A Method for Efficient Localization of Magnetic Field Sources Excited by Execution of Instructions in a Processor

While EM SCA can be a major threat or boon for hardware security, studying them requires locating the sources of the emanations. Generally, the EM side-channel signals are extremely weak, at times, requiring the measurement probe to be positioned precisely and as close as possible to the source of the EM emanations. As a result, these signals are very sensitive to noise and interference. Even if the signal is strong, it can be still be difficult to locate in the frequency domain since the signal's frequency many change as different parts of the program are executed. Further complicating matters, the sources of the signal are highly software-dependent, meaning that they can shift across the device while it is active. Therefore, to assist side-channel research, a novel method for efficiently determining the low-frequency (less than 100 MHz) magnetic field sources of the EM side channel has been developed in [12]. Chapter 3 details the how this new method models the localization problem, the efficient algorithm developed to locate the sources, the simulations and

experiments used to validate the method, and the effect program activity has on the source location.

1.3 An Efficient Method for Localization of Magnetic Field Sources that Produce High-Frequency Side-Channel Emanations

The shortcoming of the method proposed in Section 1.2 is its focus on only low-frequency sources of the EM side channel. While information about these sources is important for designers, high-frequency emanations can be much more useful for attacks, since they allow for a wider bandwidth. High-frequency emanations are not always the direct result of the program activity itself. Instead, the program activity can modulate the periodic synchronizer signals already present in the device [13]. These signals are commonly used as clocks by components, such as processors, memory, and voltage regulators, and are some of the strongest signals emanated by a device [14]. As a result, the modulation increases the distance at which emanations can be potentially detected.

At low frequencies, the relatively large traces and metal connections on the PCB are the primary elements that can efficiently radiate. At higher frequencies, smaller elements located inside the IC become more important for the side channel as their radiation efficiency improves. However, locating the sources of the high-frequency signals comes with notable challenges. As at low frequencies, the location of the sources can depend on the current program activity, meaning the source location can change as different activities are being executed. On the other hand, at higher frequencies, the impact of the equipment and the surroundings on the measurement accuracy worsens. To address these challenges, a low-cost measurement system for recording and locating the sources of the high-frequency EM side channel is proposed in [15]. This system is designed for localizing EM side-channel sources at hundreds of megahertz to gigahertz frequency range. Chapter 4 provides a detailed description of the new method, the new setup used to measure the high frequency emanations, the efficient algorithm developed to locate the sources, and the simulations

and experiments used to evaluate effectiveness of the method. It also investigates the effects frequency and program activity have on the locations of the EM side-channel sources.

1.4 Leveraging EM Side Channels for Recognizing Components on a Motherboard

With the increasingly complexity of the electronic supply chain and the threat of counterfeit ICs, the ability for electronics designers to authenticate the components used in their designs is a pressing concern. Currently, industry relies on several different methods for recognition/authentication of electronic parts. However, the effectiveness of these methods can vary based on the type of counterfeit, level of intrusiveness during testing, cost, time, and other conditions [9]. Reliable, nondestructive approaches that allow for easy, precise, and cost-effective recognition/authentication of electronic components are needed.

Therefore, [16] proposes leveraging the EM side-channel signals naturally generated by the device to recognize/authenticate components integrated onto a motherboard. By focusing on components on a motherboard, this method provides an opportunity for designers and manufacturers to authenticate devices assembled by third parties. The purpose is to detect counterfeit ICs based on changes in the EM emanations that comprise the side channel. This method is intended for detecting types of counterfeiting where the physical design of the component is altered. Examples include cases where the intended IC has been replaced with a reverse engineered copy or with a lower quality component and cases where the design has been tampered with [9]. Furthermore, this method is not intended for detecting counterfeits that are physically identical to the intended IC, such as recycled ICs (these are the focus of the next section). Chapter 5 provides detailed description of this method and the experiments used to demonstrate its effectiveness.

1.5 Detection of Recycled ICs Using Backscattering Side-Channel Analysis

One of the most common and difficult to detect types of counterfeits are recycled ICs. Since these ICs have been previously used, they have a decreased lifespan, potentially

causing the device they are integrated on to fail prematurely. Unfortunately, since recycled ICs are functionally the same as the intended IC, they are difficult to detect. Most reliable detection methods are costly, requiring modifications to or the destruction of the IC.

As an alternative, [17] proposes taking advantage of the recently defined backscattering side channel. As a type of SCA, backscattering side-channel analysis (BSCA) is able to evaluate an IC without any modifications. Unlike other types of side channels, the backscattering side channel is directly impacted by the IC aging, making it well-suited for this application. The backscattering side channel is able to detect the small changes in the impedance of the IC's transistor caused by aging.

This method is intended to assist designers in checking questionable components already integrated into their designs by third-party assemblers. Unlike most other detection methods, this method can evaluate the ICs non-destructively and without directly interfacing or modifying the IC, making it low-cost and convenient to use. Chapter 6 provides detailed description of this method, the effect aging has on the backscattered side-channel signal, and simulation and experimental results to demonstrate its effectiveness.

1.6 Research Contributions

The contributions of this dissertation are:

- A method for efficiently locating low-frequency sources of the EM side channel on PCB-based devices [12].
- A localization algorithm based on Nelder-Mead simplex optimization that determines the magnetic field sources of the side channel using only measurements taken around the edge of the PCB [12].
- A demonstration that the sources of the EM side channel depend on the instructions being executed on the device and that the leakage caused by a specific instruction can occur at multiple locations on the devic [12].

- A demonstration that, at lower frequencies, the sources of the EM side channel tend to be near the decoupling capacitors and other power supply circuitry used for the program activity [12].
- A method and near field scanning setup for locating high frequency sources of the EM side channel on PCB-based devices [15].
- An updated localization algorithm for detecting the electric and magnetic field sources of the high frequency EM side channel on PCB-based devices [15].
- A new method for recognizing/authenticating components integrated onto a motherboard based on EM SCA [16].
- A singular value decomposition (SVD) based algorithm for distinguishing components using their EM side-channel signals [16].
- A BSCA-based method for detecting recycled ICs [17].
- A description of the impact aging has on the backscattered side-channel signal and identification of what parts of the frequency spectrum are the optimal locations to detect the effects of aging [17].
- An SVD-based algorithm for identifying aged ICs from backscattering side-channel measurements [17].

1.7 Outline

The reminder of this work is broken into six chapters. Chapter 2 provides background information about the topics introduced above. Chapters 3 and 4 describe the two methods developed for locating the sources of EM side channels on PCB-based devices. Chapter 5 discusses the method developed for authenticating components already integrated on a PCB using its EM side-channel signals. Chapter 6 discusses the method for detecting

recycled ICs using BSCA and models the impact aging has on the backscattering side channel. Chapter 7 summarizes the contributions of this work and discusses avenues of future research.

CHAPTER 2
BACKGROUND

The following sections discuss side channels, hardware authentication, and recycled ICs.

2.1 Side-Channel Analysis

A side channel is an unintended avenue for observing confidential information from a device, while circumventing traditional security techniques [18]–[20]. SCA is based on monitoring the effect the device has on its physical environment or how the system responds to different inputs. Types of physical side channels include EM [21], [22], power [23], acoustic [24], temperature [25], and backscattering [26]. Other types of side channels are based on how a program executes on the device, such as timing [27] and cache attacks [28]–[31].

2.2 Electromagnetic Side Channels

EM SCA is a particular concern for hardware security since it allows for non-invasive observation from a distance and can provide more information than some other types of side channels. In EM SCA, the attacker uses the EM emanations naturally generated by a device while it is in operation to determine what is being done on the device. During the attack, an antenna is used to collect EM radiation from the device. The range of the attack can vary significantly. For very weak signals, the antenna may need to be placed directly on the device. On the other hand, attacks have been performed approximately $200\ \mathrm{m}$ away, such as in [32].

After the radiation is collected, it is processed into a usable format. EM SCA do not necessarily require expensive measurement instrumentation, such as a spectrum analyzer. There are many examples of EM side-channel attacks being performed using relatively

cheap ($1000 or less) software-defined radios available to hobbyist, such as in [33].

2.2.1 Applications

In one form or another, EM SCA has been a concern for more than 70 years. The phenomena was independently discovered in the 1940s by the United States, Britain, and Germany, while developing electromechanical cipher machines [34]. However, it did not receive much attention (at least from the United States) until the 1960s under the codename TEMPEST [35]. Public attention began in the 1980s after some of the TEMPEST research was declassified [36]. A resurgence of interest occurred in 2001 with the demonstration that EM side channels could be used to attack cryptographic algorithms in [21] and [22]. Since then researchers have demonstrated attacks on multiple devices, including keyboards [37], [38], ASIC (application-specific integrated circuit) design primitives [38], monitors [39], flash drives [40], field programmable gate arrays (FPGAs) [41], smartcards [42], and desktop computers [43].

One of the most well-known applications of EM SCA is stealing confidential information, such as cryptographic keys [33]–[45]. In addition, EM SCA has been used in more benign applications such as malware detection [1]–[3], hardware Trojan detection [46], [47], and counterfeit detection [48], [49].

2.2.2 Physical Sources of EM Side Channels

The causes of an EM side-channel signal are straightforward. Activities on the device result in variations in current. These variations cause the conductive elements carrying the current to behave as ad hoc antennas and radiate. Not all the unintended emanations from a device are considered part of the EM side channel [50]. The side-channel signal is a subset of the electromagnetic interference (EMI) from the device. Physically, nothing separates the sources of EM side channels from the sources of other unwanted emanations from a device. The EM side-channel signal is distinguished only by the fact that relevant

information about the device can be gleaned from it. Therefore, the sources of the EM side channel can be viewed in the context of EMI.

The physical elements that act as sources for the side-channel signal vary according to the type of device. In keyboards, the cabling and keys themselves radiate because of abrupt changes in current caused by keystrokes [51]. In monitors, emanations can originate from the internal circuitry or the connections between the monitor and computer [52]–[54].

PCBs are one of the most important targets of attacks. For a PCB, the side channel is the result of transistors switching states during operation. This switching causes a spike in current drawn by the IC the transistors are a part of. The strength of the spike increases with the number of switching transistors. This current spike spreads through the IC as it propagates through its internal power traces. The inductance of the IC's traces and bonds cause a voltage drop between the IC and the external power supply, resulting in a ground-bounce [55]. The power supply then spreads the effects to traces and components outside the IC; however, it will be limited somewhat by nearby decoupling capacitors [56]. At higher frequencies this leakage can couple to nearby traces as well, increasing its spread throughout the device [55]. The advantage of the EM side channel over power is that, since some of the EM emanations originate from the IC itself, they can avoid being filtered by the power supply. This property gives EM SCA a higher bandwidth.

The traces and interconnects on the PCB and its components are the most obvious sources of the EM emanations; however, identifying the specific trace/element can be extremely difficult, especially on large, complex PCBs. This challenge is exacerbated by the fact that EMI from other parts of the device can obscure the useful information. Furthermore, as demonstrated in [12], the element radiating at any point in time is based on the program activity. Depending on the program, different components or even different parts of a component may be active. As a result, the locations of the sources can change over time as different parts of the device become active.

The frequency of the emanations also has a significant effect on the side channel's

sources. At low frequencies, the strongest sources of radiation are the larger structures on the PCB. While the IC's internal circuitry will radiate as well, their small size makes them poor radiators [55]. Instead, the relatively large PCBs traces and metal connections act as the main sources of emanations. As the frequency increases, the radiating efficiency of the smaller elements improves. As a result, smaller traces, pins, and the transistors inside the ICs become more noticeable sources of the side channel. In compliance testing, ICs are considered to be too small to efficiently radiate at frequencies below $10\ \mathrm{GHz}$ to be a concern [55]. However, these weak signals can still be used in EM SCA.

2.2.3 The Influence of Program Activity

Generally, switching a transistor's state results in a short (high frequency) spike in the current and a burst of EM radiation [46]. This spike is too fast to easily locate and measure directly. However, most program activity entails more than just a single change in a transistor's state. Instead the activity usually requires using transistors spread throughout the device over a longer period of time. This activity results in a time-varying waveform being superimposed onto the device's traces. This signal is easier to measure and more informative than the emanations caused by single transistor change or even the execution of a single instruction [57].

Furthermore, the program activity itself can assist in measurements. Programs commonly entail executing repeating loops of instructions. If the execution time does not change significantly between iterations of the loop, the activity will generate a periodic signal. These periodic signals are significantly easier to locate in the frequency domain than signals produced by an arbitrary program. When researching EM side channels, the program activity used to excite the device can be tailored to maximize the periodic emanations and even control the emanation frequency [57]. This strategy is used in Chapters 3–5 and is described in more detail in Section 3.2.

Not all the side-channel signals are the direct result of the program activity. The activ-

ity can also modulate the periodic synchronizer signals (the device clocks) already present in the device [13]. Both amplitude-modulated and frequency-modulated signals can be present in the side channel [58]–[60]. From this perspective, the lower frequency signals generated from the program activity can be thought of as the baseband signal. The synchronizer signals are some of the strongest signals radiated from the device, making them convenient for EM SCA [14]. By themselves, these signals are a common concern for compliance engineers since they commonly couple onto other parts on an IC, necessitating all traces to an IC being routed as high-frequency traces [55]. Given how strong the synchronizer signals tend to be, the modulated side-channel signals can greatly increase the range of the attack.

2.2.4 Locating the Sources of the Side Channel: Near Field Scanning

A better understanding of the physical sources of EM side channels greatly benefits both researchers and designers. This information can help researchers find the optimal locations to monitor the device. Knowing the sources of the EM leakage can assist designers in identifying potential vulnerabilities to EM SCA and even detect signs of tampering with the device or its components. Furthermore, this information can help designers determine the radiation properties of the device, making it easier to simulate the device's radiated emissions [4]–[7].

Near field scanning is commonly used for finding sources of EMI and to determine the emission characteristics of the device. Given that the signals in the EM side channel are a subset of EMI, these techniques can be expanded to find the sources of the side channel. During a test, an electric or magnetic field probe is scanned in a plane at a fixed height over the device. In emissions testing, the measurements are taken very close to the device, within the reactive region of the near field. The boundary of the reactive region can be estimated to be $\lambda/2\pi$ [61].

The measurements are taken at multiple points to obtain the field over an area. The

position of the probe is usually adjusted using an automated plotter. The size, range, and positioning accuracy of the plotter vary based on the measurements' requirements. Reflections from the plotter, leakage from the cables, and interference from other devices can reduce the accuracy of the measurements [62]. Furthermore, the amount of radiation scattered by the plotter and other surrounding objects increases as the measurement frequency increases since the size of plotter becomes a larger fraction of the signal's wavelength. At the same time, the leakage from the cables and equipment used in the measurements increases with frequency. As a result of these factors, the measurement error can increase with measurement frequency.

The error caused by these factors can be limited by keeping the plotter and other equipment as far from the device under test as possible. Adding absorbing anechoic material to the measurement setup can further decrease reflections and interference. This improvement is limited unless the absorbing material can be configured correctly (as in an anechoic chamber). However, such a configuration usually requires a large space, a requirement that is not always possible. Wrapping the cables connecting the probe to the measuring instrument in absorbing material can also limit the leakage by decreasing the parasitic current flowing on the outer conductor [55].

Generally, only the magnitude of the field can be measured by scanning a probe over the device. The phase measured at each position will vary with measurement time, making it impossible to determine how the phase changes with position. A common method for determining the actual phase of the field with respect to position is to use a second probe. The second probe measures the field at a fixed location at the same time the moving probe's measurements are recorded. The phase of the field at a point can then be calculated by subtracting the phase measured by the moving probe at that point from the phase measured by the stationary probe [63].

2.3 Backscattering Side Channels

The backscattering side channel was originally defined in [64] as an avenue for detecting hardware Trojans. Unlike power and EM, the BSCA is impedance based [46]. BSCA involves transmitting a high frequency sinusoidal signal at the active IC and the monitoring the backscattered signal. The concept behind the backscattering side channel is similar to passive RFID tags. When the tag receives a signal, it can encode information in the backscattered signal by toggling between two impedances, modulating the signal. In BSCA, the IC's switching activity causes the modulation in similar manner. During operation, the individual transistors in the IC will switch on and off as they are used. As the transistors switch states, their impedance will toggle between a low impedance (on) and a high impedance (off). During BSCA, the carrier wave is transmitted at the IC while it is operating. The impedance mismatch at the IC will cause the signal to be backscattered. The cumulative switching activity of the transistors causes the impedance seen by the carrier to change with program activity [65]. The impedance will vary based on which transistors are on and whether they are NMOS or PMOS. As with an RFID tag, this impedance change modulates the signal as it is backscattered.

Monitoring the backscattered waveform provides information about how the impedance is changing on the IC. Since the backscattering side channel is impedance-based, it is sensitive to physical alterations that impact the IC's impedance, such as adding or removing circuitry or changing the impedances of the transistors. These alterations impact the frequency content of the modulated waveform.

The advantages are that BSCA has a high bandwidth, can be tuned to a part of the spectrum without interference by changing the carrier frequency, and the received power can be increased by increasing the transmitted power (as long as it remains too low to cause a fault in the IC). Furthermore, unlike current-based SCA, BSCA is affected by any change in the impedance of the IC, even in parts where the current is not flowing.

The main limitation of BSCA is that the strength of the backscattered signal is related to the size of the targeted circuit. Smaller circuits are likely to have a weaker backscattered signal, making the signal more susceptible to noise. However, BSCA's effectiveness at detecting hardware Trojans demonstrates that it can be sensitive to very small impedance changes [46]. Furthermore, a weak signal can be somewhat compensated for by increasing the carrier power.

BSCA is much more sensitive to changes at lower frequencies than current-based side channels, such as EM and power. When a transistor switches state, its impedance will change once and then remain constant throughout the clock cycle. As a result, the impedance change will occur in discrete steps, similar to a square wave. On the other hand, in CMOS-based devices, current flows through the transistors in short bursts, immediately after the transistors switch states. Therefore, the strength of the low frequency harmonics of the current-based side channels is much weaker.

2.4 Hardware Authentication

Another important concern for hardware security is determining whether electronic devices and the components on the device are what they claim to be. Designers have increasingly lost control over the manufacturing process. With the cost rapidly decreasing, many designers have outsourced the manufacturing of their designs and the procurement of the components needed for their designs to third parties. This outsourcing has left devices and their components vulnerable to being tampered with during assembly [66].

In general, any unauthorized IC can be considered a counterfeit IC. Types of counterfeits include unauthorized copies, incorrectly labeled/re-marked ICs, out-of-spec ICs, hardware Trojans, and old ICs that have been recycled [9]. Recycled ICs are discussed in more detail in Section 2.5. Counterfeits do not need to be malicious to be damaging. Device assemblers commonly replace components selected by a designer with components they consider equivalent. Usually this replacement is not a concern. However, in some

cases, even if the replacement is functionally compatible, it may differ in other properties, such as reliability, environmental tolerances, inter-operability with software and other components, level of trust, and other bugs and vulnerabilities. To overcome this problem, it is important to correctly recognize/authenticate components on a PCB or in a system, so that the appropriate software patches and workarounds can be applied, and so that tracking and mitigation of reliability and inter-operability issues can be correctly implemented.

2.4.1 Common Component Authentication Approaches

Currently, industry relies on several methods to authenticate electronic components. The most common can be classified as either physical or electrical inspection [67]. Types of physical inspection include visually examining the interior and exterior of the component and analyzing its material composition. Types of electrical inspection include testing the component's electrical characteristics, performance, and durability through burn-in tests [67]. However, the effectiveness of these methods can vary based on factors such as the type of counterfeit and level of intrusiveness during testing [9]. For example, parametric tests are useful for identifying recycled components, but are not effective at identifying unauthorized copies [9]. Furthermore, many of these tests are expensive and time consuming, and, at times, the components are destroyed in the process. One way of improving testing is by including additional circuitry into the design to aid in the authentication process [68]. However, this addition requires physical modifications to the component before it is manufactured, increasing the complexity and price of the component.

To combat IC counterfeiting, reliable inexpensive methods for recognizing/authenticating electronic components are still needed. One notable approach intended to address many of the previous challenges is the Supply Chain Hardware Integrity for Electronics Defense (SHIELD) [69]. SHIELD installs on a passive sensor called a "dielet" into the component's package during manufacturing. This sensor records unauthorized attempts to access or to modify the component's hardware. The dielet's recordings can then be accessed wirelessly

during authentication. While this approach may be effective for detecting access or modifications, it requires additional hardware to be added to the component, increasing its cost. Furthermore, the dielet itself is vulnerable to being tampered with.

Identification using EM Side-Channel Signals

EM side channels are a promising alternative for identifying both ICs and the devices they are integrated onto. The EM emanations generated by a component/device are directly related to the program activity and the physical properties. If the program activity is kept constant, the emanations can be used as a signature for identification. The most important factor in this process is carefully selecting the features in the signature to use for identification.

EM side channels have already been used for identifying devices such as automobiles, desktops, phone chargers, cell phones, toys, and microcontrollers [70]–[77]. However, in most of these situations, the identification accuracy was likely assisted by the fact radically different devices were being compared.

For component authentication, the side-channel signal can be used to distinguish between similar types of ICs or detect unauthorized physical changes to a component. This signal is affected by the physical properties of the component. For example, the physical properties can affect the emanations' frequency, magnitude, and directivity. Changes to the component can be detected by comparing the measured emanations to the expected signature. One of the few attempts to identify components in this manner is discussed in [48], [77], and [78]. That work, however, is narrow. It focused only on distinguishing individual devices of the same type from each other and is limited to two types of FPGAs and one type of microcontroller.

2.5 Recycled ICs

One of the most common types of counterfeits are recycled ICs. Recycled ICs have been removed from discarded electronics and then sold as if they are new. Assuming the counterfeiter competently removed signs of previous use from the package, the only differences between a recycled and a new IC are the slight degradation in the properties of the transistors. As a result, it is challenging to identify recycled ICs (until they prematurely fail).

As with general component authentication, detection methods rely on either physical or electrical inspection. Physical inspection involves examining the IC for signs of previous use. Examples include checking the ICs' pins for solder, removing layers of the package to check for remarking, and searching defects in the IC [9]. The drawbacks of these methods are that they have poor accuracy when the recycler is careful about repacking and remarking the IC, are time consuming, and may destroy the IC [9], [48], [79].

Electrical inspection focuses on the detecting degradation of the IC's electrical properties that are related to the age. It involves taking parametric measurements on questionable ICs and comparing them to the parameters expected if the IC is unaged. Statistical tests, such as machine learning, are used to determine whether the questionable IC differs significantly [80]–[82]. One of the most common parameters tested is path delay. Over time, the impedance of a transistor increases, causing its switching speed to decrease. This slowdown is commonly measured by integrating part of the circuit into a ring oscillator and measuring the frequency. As the IC ages, the path delay increases, causing the ring oscillator's frequency to decrease [83]–[88]. The drawback is that a ring oscillator can only test a single path on the IC. Since the effect of aging is unlikely to be consistent over the entire IC, multiple paths usually need to be tested, with each path requiring a ring oscillator. The extra resources needed to construct and interface with each ring oscillator will increase the cost and requires ICs already on the market to be redesigned. Some of this cost can be avoided on FPGAs since the oscillators can be easily configured using the FPGA's

resources [89]–[91]. However, even in this situation, extra resources are needed to communicate with the FPGA, and it can be time consuming to test enough paths to accurately measure the age. Furthermore, ring oscillators are sensitive to other environmental factors, such as temperature, hence decreasing their accuracy.

2.5.1 Mechanisms of IC Aging

There are four main aging mechanisms for MOSFET transistors: Bias temperature instability (BTI), Hot Carrier Injection (HCI), time-dependent dielectric breakdown (TDDB), and electromigration [92], [93]. In nanometer technologies, the biggest concerns are BTI and HCI [94], [95].

BTI has the biggest impact on the lifespan of a smaller MOSFETs [93], [96]. There are two types BTI, negative BTI (NBTI) and positive BTI (PBTI). NBTI impacts PMOS transistors, while PBTI impacts NMOS transistors. The process behind both is challenging to accurately model. It depends heavily on the size and material the MOSFET is made of. In general, NBTI is caused by static negative gate voltage and high temperature [93]. The electric field resulting from the voltage causes the hydrogen ions in the interface of the channel and gate dielectric to defuse towards the gate, leaving defects behind [92]–[96]. PBTI is the reverse. A positive gate voltage and high temperature causes electrons to become trapped in pre-existing defects in the gate dielectric [98], [99]. The effects of both types of BTI will somewhat reverse once the gate voltage is no longer applied [94]. However, the partial-recovery will not completely reverse the degradation.

In larger/older CMOS technologies, the degradation from NBTI is noticeably worse than PBTI [90]. However, with smaller technologies and the introduction of high-k gate dielectrics, PBTI in NMOS transistors has become a similar level of concern [99]–[101].

HCI is the result of “hot” carriers from the MOSFET’s channel becoming injected in the interface between the channel and gate dielectric [93]. The injected carrier will either become embedded in the dielectric or become gate current [80]. Since HCI is gener-

ally caused by current flowing through the MOSFET's channel, it occurs during transistor switching. This is the opposite of BTI, which is caused by static voltages [92]. Similar to BTI, as the impact of HCI increases in NMOS and PMOS transistors as their size decreases [93].

The defects caused by both BTI and HCI result in an increase in the magnitude of the threshold voltage, V_{th}, of the MOSFET (becoming more negative in PMOS transistors and more positive in NMOS transistors). Among other things, this increase results in an increase in the impedance and a decrease in the MOSFET's switching speed [90], [92], [97]. In other words, as the transistors age, the rise and fall times of the signal propagating through the transistor increase. Eventually, the degradation results in timing failures. While this effect is detrimental to the operation of the device, it can be used to track the age of the IC.

CHAPTER 3

A METHOD FOR EFFICIENT LOCALIZATION OF MAGNETIC FIELD SOURCES EXCITED BY EXECUTION OF INSTRUCTIONS IN A PROCESSOR

3.1 Overview

As discussed in Chapter 1, one of the main challenges for EM SCA is identifying the sources of the signal. To address this problem, we have developed a method that can efficiently locate the *instruction-dependent* magnetic field sources on a PCB. This method can significantly reduce the number of measurement points and the time needed to identify the sources by only taking measurements along the edges of the PCB.

When testing a device, it is excited in a controlled manner using the benchmark introduced in [102]. This benchmark makes it possible to relate the sources to the specific instruction that causes the radiation. It also makes is possible to generate an artificial leakage signal at a specific frequency that is directly related to processor instructions.

After the measurements are taken, they are used to solve a forward-backward optimization problem to identify the locations of the magnetic field sources, the magnitudes of their magnetic moments, and their orientations. The sources are assumed to be electrically small, time-harmonic quasi-stationary magnetic loops, where the number of sources can be either known (provided to the model) or unknown (chosen by the model itself). The equations describing the magnetic-flux density generated by sources are referred to as the forward model. Two optimization techniques, either a Nelder-Mead simplex method [103] or a particle swarm optimization (PSO) [104], are used to optimize the parameters of the sources in the forward model to generate a magnetic field that matches the measured field.

The accuracy of the method is verified against a variety of simulations in AWAS [105] and against measurements on an FPGA and an Internet-of-Things (IoT) development board.

The results demonstrate that the proposed algorithm can accurately identify those sources, regardless of the instructions executed. Furthermore, the experimental results show that the number of strong magnetic field sources on a PCB depends on the instructions being executed. This is an interesting result because it indicates that some instructions cause emanations from multiple sources, making them potentially easier to exploit for SCA.

The rest of this chapter is organized as follows. Section 3.2 describes the benchmark used to generate the program activity at controlled frequencies, Section 3.3 describes proposed algorithm for localizing magnetic field sources on PCBs, Section 3.4 tests robustness of the algorithm, Section 3.5 describes the measurement setup and presents experimental results, and Section 3.6 concludes the chapter.

3.2 Creating System Activity at Controlled Frequencies

The benchmark described in [102] provides a way of producing artificial leakage from a PCB that is directly related to the processor instructions. This benchmark makes it possible to generate strong, repeatable instruction-dependent emanations at a specific frequency and for a controllable amount of time. This control simplifies finding the side-channel signal in the frequency domain, since in a real program, the frequency of the emanations may change as different parts of the code are executed. Controllable side-channel emanations are generated by executing a pattern of two instructions repeatedly. Any difference in the current drawn when executing the instructions results in a periodic current being superimposed onto the power and signal interconnects in the IC and on the PCB.

The excitation program is shown in Fig. 3.1. It is comprised of two loops, one that repeatedly performs some activity X (line 2), while the other performs some other activity Y (line 8). These loops are enclosed in an outer loop (line 1) that causes the program to repeatedly alternate between activities X and Y. This alternation creates periodically-changing signal whose period equals the execution time of one iteration of the outer loop. This alternation period, T_{alt}, is the inverse of the frequency, $f_{\mathrm{alt}} = 1/T_{\mathrm{alt}}$. Changing the

activities in the benchmark excites different parts of the processor and memory circuitry on the PCB.

```
1  while(true){
2    // Execute the X activity
3    for(i=0;i<inst_x_count;i++){
4      ptr1=(ptr1&~mask1)|((ptr1+offset)&mask1);
5      // The X-instruction, e.g., a load from L2
6      value=*ptr1;
7    }
8    // Execute the Y activity
9    for(i=0;i<inst_y_count;i++){
10     ptr2=(ptr2&~mask2)|((ptr2+offset)&mask2);
11     // The Y-instruction, e.g a store from L2
12     *ptr2=value;
13   }
14 }
```

Fig. 3.1. Pseudo-code for generating the X/Y alternation activity.

While the effect of a single event (i.e., execution of a single memory access or processor instruction) on the side-channel signal is unknown, as long as there is some difference between the X and Y activities, there will be a signal generated at the frequency f_{alt} and some of its harmonics ($2f_{alt}, 3f_{alt}, ...$). In the measurements, only the magnitude of the first harmonic, f_{alt}, is recorded. We can choose the frequency of emanations to select a part of the spectrum that has minimal interference and investigate how the emanations change with frequency.

Finally, note that the excitation signal also amplitude modulates other periodic signals generated by the PCB [13]. The synchronizer clocks of the components used for executing the program activity can act as carriers for the modulated waveform. During normal operations, the clocks produce periodic currents at the clock frequency f_{C} along the signal and power traces. When the excitation program is executed, the periodic current from the clock is modulated by the current drawn when executing the program activity. As a result, the modulated waveform can then be observed at $f_{\mathrm{C}} \pm f_{\mathrm{alt}}$ in the spectrum.

3.3 Localization of Magnetic Field Sources

This section describes the proposed method for the locating the magnetic field sources of the side-channel signal.

3.3.1 Problem Statement

Here, an inverse electromagnetic problem of identifying sources of a time-harmonic quasi-stationary magnetic field is considered. The operating frequency f, and therefore the angular frequency $\omega = 2\pi f$, is known. The sources are located in a parallelepiped whose dimensions are $2a_s$, $2b_s$, and h_s, as shown in Fig. 3.2. This parallelepiped is the source space (ν) and is assumed to be known. The source space is sitting on an infinite, perfectly conducting (PEC) ground plane.

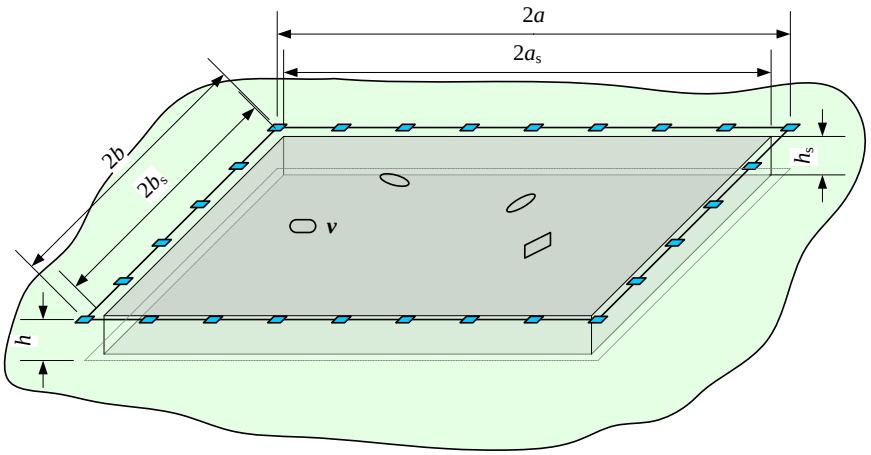

Fig. 3.2. Source space (ν) with several current-carrying loops and test points (the light blue squares).

Within ν, there can be one or more sources. Each source is a small loop, and can be considered as a magnetic dipole of an unknown magnetic moment **m**. Here, it is assumed the number of dipoles is known. However, in Section 3.4.2, the method is evaluated by assuming an unknown number of dipoles. In both cases, the magnitudes of the moments and

their orientations are unknown. All moments are assumed to be linearly polarized. While in practice the moments seem to be in-phase across the whole PCB, when determining these moments, one moment is selected as the reference (whose initial phase is 0). The phases of the other sources are determined with respect to the reference.

It is assumed that measurements of the magnetic field are performed at a known set of P points (referred to as the test points) located around the source space (i.e., the PCB). These points are located along a rectangle, whose sides are $2a$ and $2b$, positioned at an elevation h above the ground plane, as shown in Fig. 3.2. The test points are equally spaced along each side of the rectangle, with n_a and n_b spacings along a and b, respectively. The total number of the test points is thus $P = 2n_a + 2n_b$. At each point, the magnitudes of the three Cartesian components of the magnetic-flux density (magnetic induction) vector (B_x, B_y, and B_z) are measured. Therefore, $3P$ scalar quantities are known for each test point.

The objective is to estimate the magnetic moment vectors of the sources. The forward model described in the following section was developed to estimate these vectors. Simplex or the particle swarm optimization (PSO) is used to estimate the moments of these dipoles in an attempt to make the forward model produce magnetic fields at the test points as close as possible to the measured fields.

3.3.2 Forward Model

The forward electromagnetic model is used to calculate the magnetic field of the magnetic dipoles sources. The number of dipoles and their phasor magnetic moments, locations, and orientations are known. The resulting magnetic field (magnetic-flux density) vector is evaluated at a set of test points. The field is assumed to be quasi-stationary. The currents induced in the PEC plane are taken into account using images of the magnetic moments.

Figure 3.3 (left) shows the local Cartesian and spherical coordinate systems attached to a magnetic dipole. The dipole is located at the coordinate origin and its momentum is oriented along the z-axis. The index "l" indicates that this is the local coordinate system

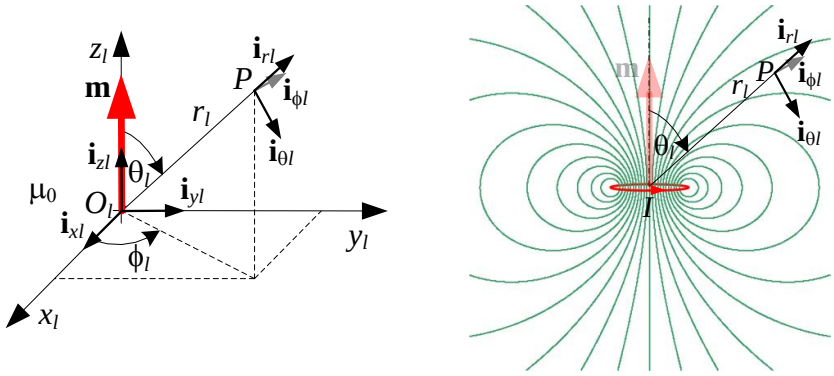

Fig. 3.3. (left) Local Cartesian coordinate system attached to a magnetic dipole with the corresponding spherical system; (right) magnetic field lines of a small loop.

and is different from the global system. Figure 3.3 (right) shows a small current-carrying loop, which represents a magnetic dipole, with the lines of the magnetic field and the local spherical coordinate system.

Assuming the magnetic dipole to be located in a vacuum, its magnetic-flux density at the field point P is given by

$$\begin{aligned} \mathbf{B} &= -\frac{\mu_0}{4\pi}\mathrm{grad}\left(\frac{\mathbf{m}\cdot\mathbf{r}_{0l}}{r_l^2}\right) \\ &= \frac{\mu_0}{4\pi}\frac{m_{zl}}{r_l^3}\left(2\mathbf{i}_{rl}\cos\theta_l + \mathbf{i}_{\theta l}\sin\theta_l\right), \end{aligned} \tag{3.1}$$

where $\mathbf{m} = m_{zl}\mathbf{i}_{zl}$ is the magnetic moment of the dipole, $\mathbf{i}_{\mathrm{x}l}$, $\mathbf{i}_{\mathrm{y}l}$, and $\mathbf{i}_{\mathrm{z}l}$ are the unit vectors of the local Cartesian coordinate system, whereas $\mathbf{r}_{0l} = \mathbf{i}_{\mathrm{r}l}$, $\mathbf{i}_{\theta l}$, and $\mathbf{i}_{\phi l}$ are the unit vectors of the local spherical coordinate system attached to the dipole. Alternatively, the local Cartesian components of the vector $\mathbf{B}$ can be found by noting that $\mathbf{m}\cdot\mathbf{r}_{0l} = zm_{zl}/r$, as

$$\begin{aligned} \mathbf{B} &= -\frac{\mu_0}{4\pi}m_{zl}\mathrm{grad}\left(\frac{z_l}{r_l^3}\right) \\ &= \frac{\mu_0}{4\pi}\frac{m_{zl}}{r_l^3}\left(\frac{3z_l}{r_l^2}(x_l\mathbf{i}_{xl} + y_l\mathbf{i}_{yl} + z_l\mathbf{i}_{zl}) - \mathbf{i}_{zl}\right) \\ &= \frac{\mu_0}{4\pi}\frac{m_{zl}}{r_l^3}\left(\frac{3z_l}{r_l}\mathbf{r}_{0l} - \mathbf{i}_{zl}\right). \end{aligned} \tag{3.2}$$

If the magnetic moment **m** has an arbitrary orientation in the global system, its Cartesian components can be found. For each component, we can design a local coordinate system as in Fig. 3.3, find the components of the vector **B** in that system, and then translate them to the Cartesian components in the global Cartesian system. The resulting vector **B**, due to all magnetic dipoles, is obtained by summing the global Cartesian components of the field due to the individual dipoles and their images.

3.3.3 Optimization Procedure

In the optimization process, it is assumed that the number of magnetic dipoles is known, but can be equal to, smaller than, or greater than the actual number of loops. For each dipole, except for the last one, seven optimization variables are used with the following restrictions:

- the magnitude of the dipole moment, *m*;
- the Cartesian *x*-coordinate of the dipole center with respect to the global system, $-a_{\mathrm{s}} \leqslant x \leqslant a_{\mathrm{s}}$;
- the Cartesian *y*-coordinate of the dipole center with respect to the global system, $-b_{\mathrm{s}} \leqslant y \leqslant b_{\mathrm{s}}$;
- the Cartesian *z*-coordinate of the dipole center with respect to the global system, $0 \leqslant z \leqslant h_{\mathrm{s}}$;
- the spherical θ-coordinate of the dipole moment vector, $-2\pi \leqslant \theta \leqslant 2\pi$;
- the spherical ϕ-coordinate of the dipole moment vector, $-2\pi \leqslant \phi \leqslant 2\pi$;
- the initial phase of the dipole moment, $-2\pi \leqslant \alpha \leqslant 2\pi$.

Since the initial phase for the last dipole is assumed to be 0, it is not included. Therefore, the total number of optimization variables is $N = 7n - 1$, where $N \leqslant P$ and n is the

total number of dipoles. The restrictions for the angles θ, ϕ, and α are extended beyond their respective basic ranges necessary for their definition in order to enable a flexible optimization. Otherwise, the optimization process may dwell around the boundaries of the basic ranges. The cost-function used in the optimization algorithms is

$$F = \frac{4\pi}{\mu_0 P} \left(||B_{x\mathrm{f}}| - |B_{x\mathrm{m}}|| + ||B_{y\mathrm{f}}| - |B_{y\mathrm{m}}|| + ||B_{z\mathrm{f}}| - |B_{z\mathrm{m}}|| \right), \tag{3.3}$$

where the index f denotes the values from the forward model, and the index m denotes the measured values.

A simplex optimization or particle-swarm optimization (PSO) is used to solve the optimization problem. For the simplex algorithm, the cost function is set to a large value (100) if a boundary for any variable is exceeded. Additionally, the simplex algorithm is combined with a random search. We implemented these algorithms were implemented in FORTRAN and C.

The lengthiest algorithm randomly selects a point within the allowed space for the optimization variables, launches a simplex search from the point and then records the results. The algorithm repeats this procedure multiple times before selecting the optimal solution. This approach produces thousands or even millions of evaluations of the cost function; however, it has an increased probability of finding the global optimum. Most runs follow this approach, and they are used to prove the concept and tailor the cost function. A more sophisticated approach is to run a random search, find the optimal point, and then launch a simplex search from the optimal point. Finally, the PSO algorithm is launched within the allowed space for the optimization variables.

3.4 Model Verification and Robustness

In this section, the proposed localization model is first tested on simulation results from AWAS [105]. The second step in the validation is to add white Gaussian noise and check

the accuracy of the localization model. Furthermore, how the model behaves if the number of sources is unknown (which is typical for practical applications) is investigated.

3.4.1 Verification of Localization Algorithm via Simulated Data

The first step in validating the localization model is to use simulated data as a proxy for real measurements. AWAS is used to compute the magnetic-flux density of a known set of sources at the test points.

The operating frequency in the simulations is $f = 154.5$ kHz. The test points are located along a rectangle with dimensions $a = 50$ mm and $b = 40$ mm, and an elevation above the ground plane of $h = 30$ mm. The numbers of test points along a and b are $n_{\mathrm{a}} = 10$ and $n_{\mathrm{b}} = 8$. Therefore, the total number of test points is $P = 36$. The dimensions of the source space are defined by $a_{\mathrm{s}} = 0.975a$, $b_{\mathrm{s}} = 0.975b$, and $h_{\mathrm{s}} = h$.

Simulations with one, two, and three magnetic loops (i.e., dipoles) are used for verification. Each dipole is a small square, whose sides are $s = 1$ mm, centered at a point (within the source space), and whose Cartesian coordinates are (x_i, y_i, z_i), where $i = 1, 2, 3$. Each loop is fed by an ideal current generator, whose RMS current is I_i. The magnetic moment of a loop is $m_i = I_i S^2$, where $S = 1$ mm^2 is the surface area of the loops. The currents of the loops are $I_1 = 1$ mA, $I_2 = 2$ mA, and $I_3 = 0.25$ mA. The reference directions of the vectors $\mathbf{m}_{\mathrm{i}}$ are different for the three loops: $\mathbf{m}_1$ is directed along the y-axis, $\mathbf{m}_2$ along the x-axis, and $\mathbf{m}_3$ along the z-axis. The data for the three loops are summarized in Table 3.1, where the magnetic moments are defined as phasors in terms of their Cartesian components. Alternatively, the magnetic moment vector can be defined in terms of spherical coordinates as (m, θ, ϕ) , where m is the RMS phasor, θ is the zenith angle (the angle with respect to the z-axis, $0 \leqslant \theta \leqslant \pi$), and ϕ is the azimuth angle (the angle between the projection of the vector onto the Oxy-plane and the x-axis, $-\pi \leqslant \phi \leqslant \pi$).

In the first model, only loop 1 from Table 3.1 is present. The second model comprises loops 1 and 2, while the third model has all three loops. Note that the electromagnetic sys-

TABLE 3.1
Parameters of simulated loop sources.

i	x_i [mm]	y_i [mm]	z_i [mm]	m_x [nA m^2]	m_y [nA m^2]	m_z [nA m^2]	m [nA m^2]	θ [rad]	ϕ [rad]
1	30.5	20	5.5	0	–1	0	1	$\frac{\pi}{2}$	$-\frac{\pi}{2}$
2	0	–10.5	2.5	–2	0	0	2	$\frac{\pi}{2}$	π
3	–40.5	20.5	15	0	0	0.25	0.25	0	0

tem is quasi-stationary: the dimensions of the system are much smaller than the wavelength (which is close to $2\ \mathrm{km}$). Hence, the retardation of the electromagnetic fields is negligible, and the field (test) points are in the near-field zone of the sources. In AWAS, the near-field is computed at a uniform rectangular grid of $(n_x + 1)(n_y + 1) = 99$ points, spanned by the test points. This grid is parallel to the Oxy-plane at the height h above this plane. The computed field is used to create a visualization of the intensity of the magnetic field at this height. Also, the data for the field at the test points are extracted and used to estimate the field sources.

Figure 3.4 shows contour plots for the total field (more precisely, the intensity of the magnetic-flux density vector normalized to its maximal value on the grid) for 1, 2, and 3 loops, obtained using AWAS. Dark blue square markers denote the positions of the magnetic dipoles. White lines represent the moments of the dipoles. The moments are plotted in a single direction to stress their initial phases. The moment of the third dipole (positioned in the upper-left corner) is perpendicular to the plot and directed outwards. Light blue square markers denote positions of the test points.

Note that the local maxima of the field need not be exactly above each loop. There are two reasons for this. First, for an arbitrarily positioned loop, the maximum may be shifted because the intensity of the magnetic field, at a fixed distance from the loop, varies with spherical angles defined with respect to the local coordinate system attached to the

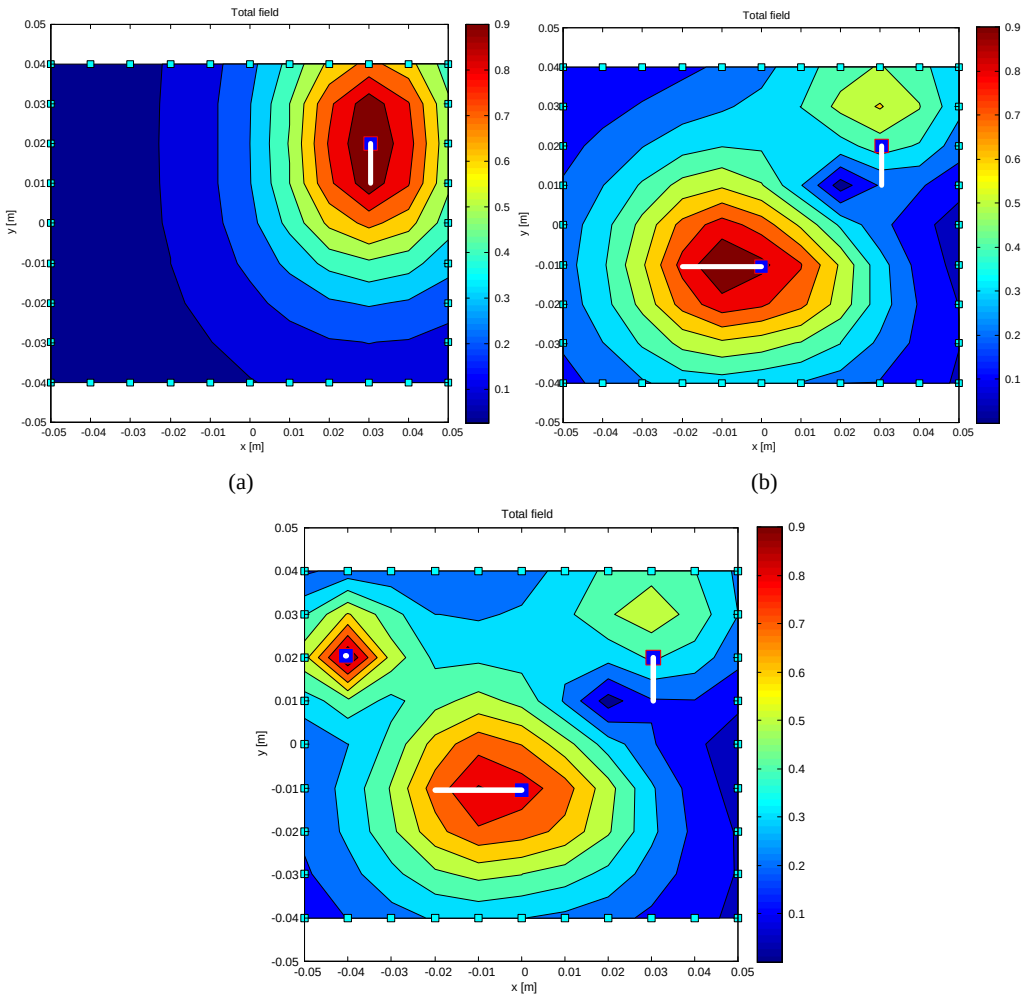

Fig. 3.4. Normalized intensity of the magnetic field for (a) one loop, (b) two loops, and (c) three loops.

loop. Second, the fields due to the loops are synchronized; hence, they interfere, creating a complicated pattern in the plane of the plot. Also note that the magnetic fields of the loops are strongly affected by the presence of the ground plane: far away from the loops, the fields of the first two loops are enhanced, whereas the field of the third loop is reduced.

Simplex optimization is run with $10,000$ randomly selected starting points, an initial step of 0.05, a maximal number of simplex iterations 2000, and the prescribed accuracy for function values and optimization variables set to 10^{-12}. Unless stated otherwise, the same settings are used for the other examples. The results are shown in Fig. 3.5.

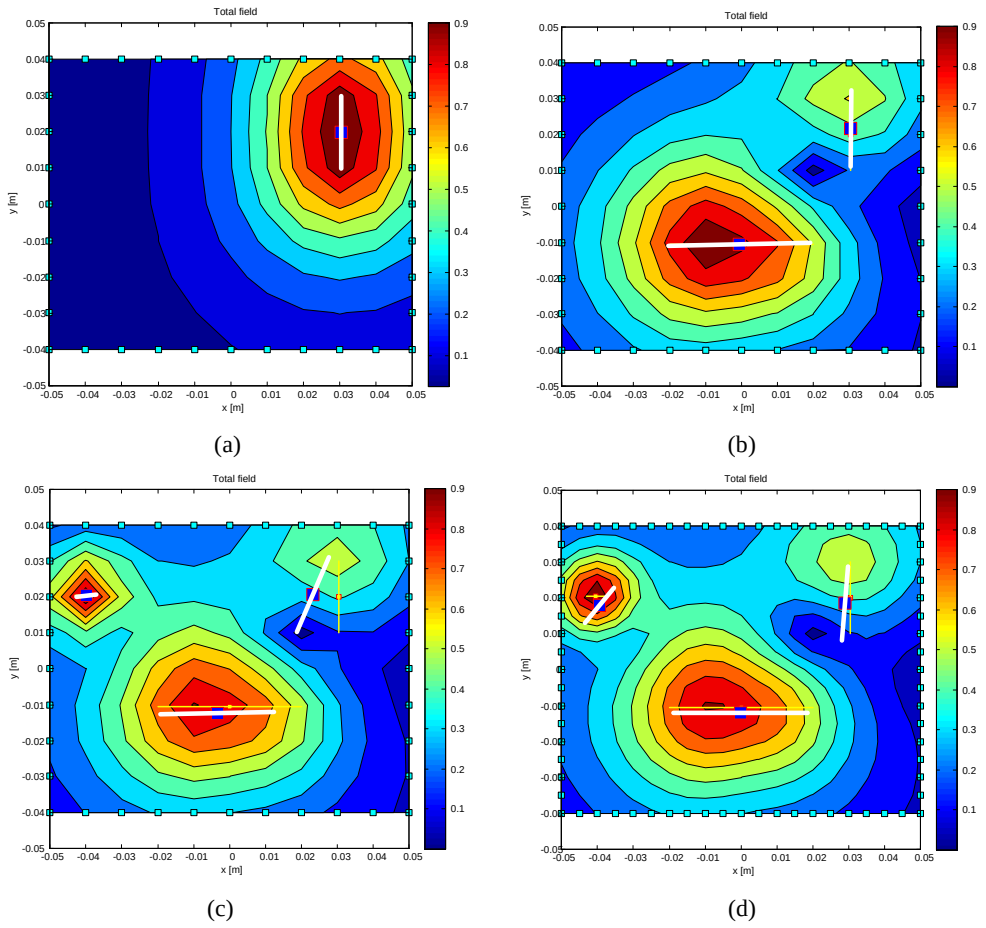

Fig. 3.5. Results of simplex optimization launched from randomly selected points for (a) one loop, (b) two loops, (c) three loops, and (d) three loops with an increased number of test points.

In Fig. 3.5, the exact solutions for the magnetic dipoles are represented by smaller square markers and yellow lines. These lines go in both directions to highlight that we are no longer interested in the phases. The solutions of the optimization procedure are indicated by larger square markers and white lines. The lengths of these lines are proportional to the estimated moments. The directions of these lines correspond to the estimated azimuth angle, but the estimated zenith angle is not reflected in the plots.

As Fig. 3.5 demonstrates, with one loop, the problem is solved almost perfectly. However, the accuracy of the solution deteriorates with increasing the number of loops. The accuracy of the solution for the system with three loops is acceptable, but far from being

perfect. In complicated systems, the improvement of the optimization function is relatively small, only one order of magnitude. The optimization function is multimodal, i.e., it has a multiple minima. Therefore, the simplex optimization can terminate prematurely at a local minimum and the global minimum (optimum) is very rarely found. The problem is considered to be is ill-posed. Increasing the number of test points (P) may improve the quality of the solution, as illustrated on Fig. 3.5 (d). However, note that this is not the major limitation for the applications such as SCA or hardware Trojan detection where one reliable location to collect EM signals is sufficient.

3.4.2 Localization with Unknown Number of Sources

This subsection illustrates what happens if the number of sources is unknown and instead set arbitrarily. This situation is realistic since often the number of sources is not known beforehand. Fig. 3.6 (a) presents results for the case of one loop when it is assumed that there are three dipoles present. The algorithm finds two additional, parasitic sources, indicating the ill-posedness of the problem, which is typical for inverse problems such as the one solved here.

Figure 3.6 (b) presents results for the system of three loops when it is assumed that there is one magnetic dipole, while Fig. 3.6 (c) presents the same system when it is assumed that there are five dipoles. When only one magnetic dipole is assumed, the algorithm finds the dominant source. When it is assumed that there are more dipoles than the actual number of loops, parasitic solutions exist. However, the parasitic dipoles have small moments, indicating that that they are parasitic solutions and can be neglected.

3.4.3 Localization Using Simulated Data with Noise

The second step for verifying the localization algorithm accuracy is to test it on simulated data with added white Gaussian noise. In the simulations, the signal-to-noise ratio (SNR) is calculated as the ratio of the intensity of the strongest field at the test points and the standard

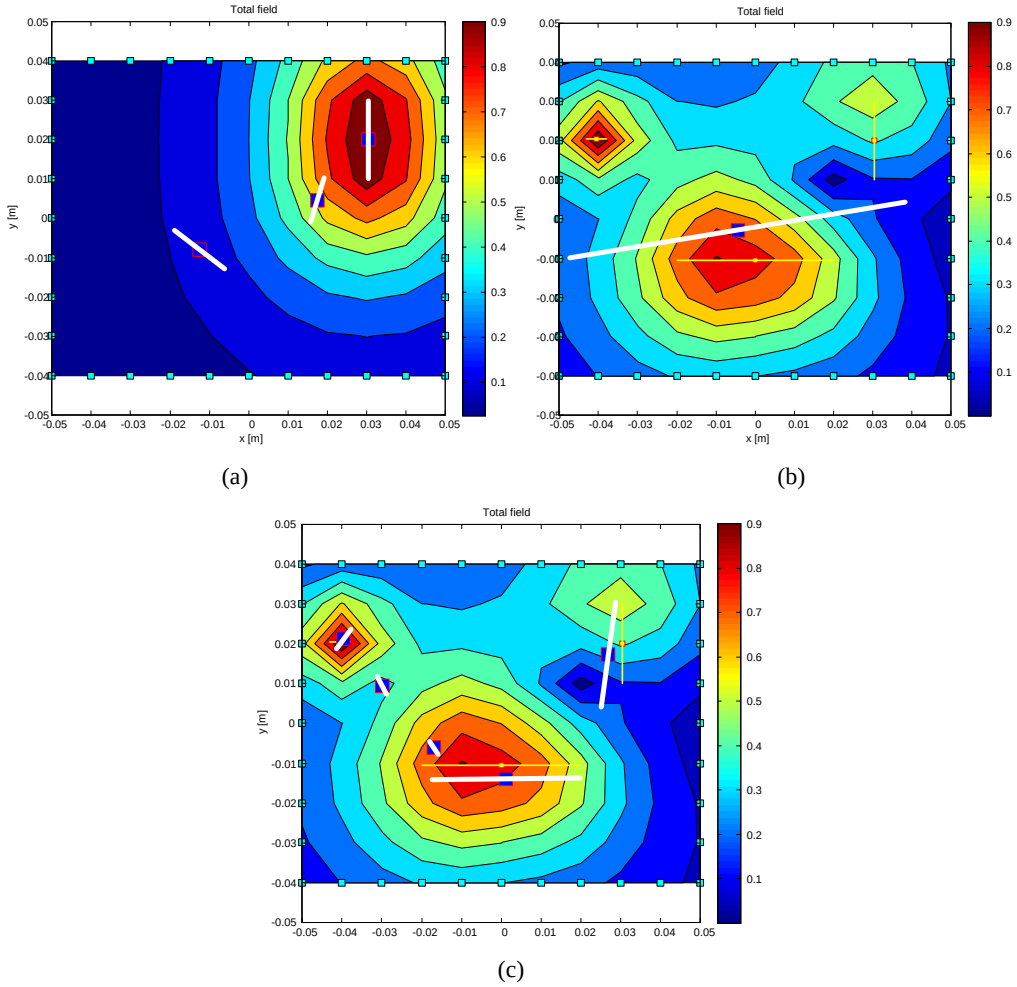

Fig. 3.6. Results of simplex optimization when the number of assumed dipoles does not match the actual number of loops: (a) one loop and three dipoles, (b) three loops and one dipole, and (c) three loops and five dipoles.

deviation of the Gaussian noise added to each field component at these points. We consider the case of three loops when there is no added noise, when the SNR is 30 dB, 20 dB, and 10 dB. The results are shown in Fig. 3.7 and should be compared with the results of localization of three loops without added noise in Fig. 3.5 (c). The results demonstrate that the algorithm is robust against noise: the results without noise and for $\mathrm{SNR} = 30$ dB do not differ significantly. Even for $\mathrm{SNR} = 20$ dB, the results are acceptable. The quality of the results is significantly deteriorated only for $\mathrm{SNR} = 10$ dB.

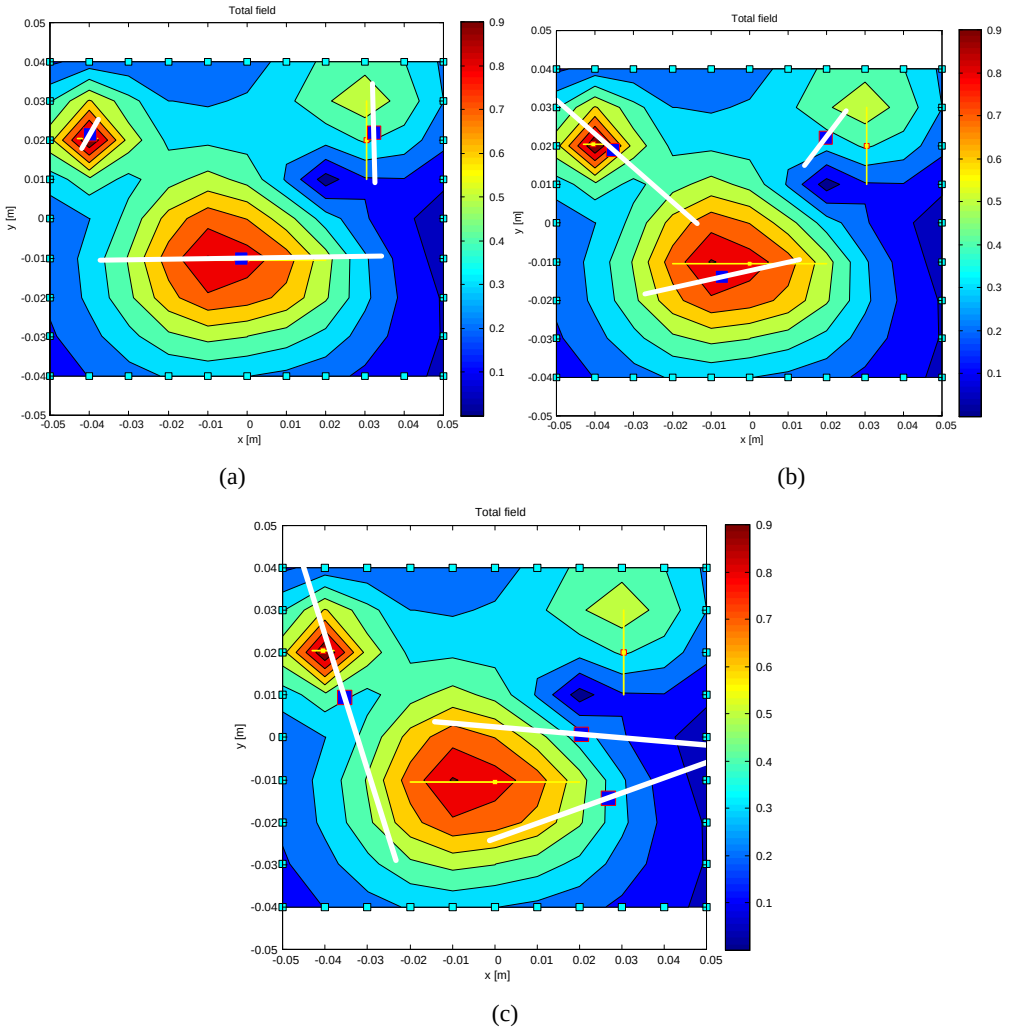

Fig. 3.7. Optimization for the case of three loops when the SNR is (a) 30 dB, (b) 20 dB, and (c) 10 dB.

3.5 Experimental Results

In the following subsections, we first describe a measurement setup that has been developed for accurately and consistently measuring the magnetic fields generated by an electronic device and devices used the algorithm is tested on. Then, the proposed localization model is tested by comparing modeled with measured data from the real devices. Finally, we present localization results obtained for several different pairs of instructions on different devices.

3.5.1 Measurement Setup

The measurement setup is shown in Fig. 3.8. To measure the magnetic field around the device under test, a MakeBlock XYPlotter Robot Kit v2.0 (which has an x-y accuracy of 0.1 mm) is used for positioning. This setup has a significantly lower cost compared to professional 2-D scanners. A hand-made 33-turn coil magnetic field probe with a radius of 5 mm is connected to the plotter for collecting measurements. At all times, the probe remains 3 cm above the PCB being tested. The magnitude of the power across the loop probe is measured by a MXA N9020A Spectrum Analyzer from Keysight. Any cables connected to the device are shielded (i.e., wrapped by copper tape) to minimize the influence on the measurements, as shown in Fig. 3.8.

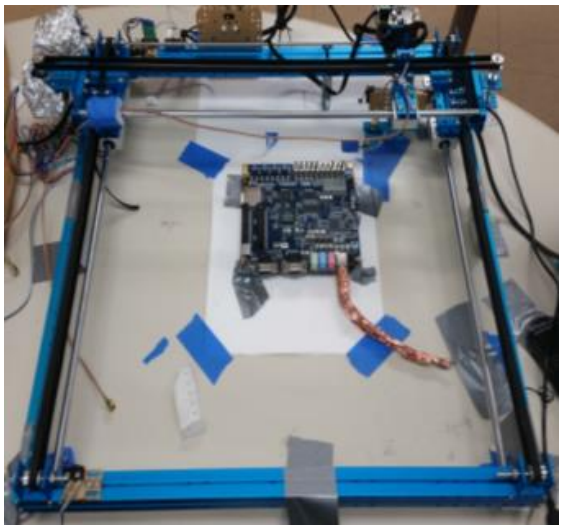

Fig. 3.8. The measurement setup.

To approximate an isotropic receiver, three sets of measurements are taken for each pair of instructions. In each set, the probe is oriented parallel to the x-axis, y-axis, or z-axis, respectively. Taking separate measurements with the probe in the three different orientations is equivalent to taking measurements using a set of collocated loops. During the tests, the magnetic probe moves with a step size of 0.5 cm across the entire area of the device. Using this setup, it takes 15 hours to measure the magnetic field over the entire area

of the larger device (the DE1). However, the model only needs the measurements around the edge of the device. The rest of the measurements are for demonstrating the accuracy of the model. It takes only an hour to measure around the edges of the DE1.

3.5.2 Devices Under Test

In the following subsections, measurements are collected for two devices: a Cyclone II FPGA DE1 development board from Altera and Terasic and an A13-OLinuXino-MICRO IoT Linux computer board from Olimex. For simplicity, the devices are referred to as the DE1 and the A13-MICRO for the rest of this chapter. The DE1 implements a Nios-II soft processor, while the A13-MICRO has an ARM A13 Cortex-A8 processor. The accuracy of the model is first confirmed on the DE1 and then used on the A13-MICRO to demonstrate its effectiveness on more complicated devices. These devices are convenient for testing and representative of a wide variety of embedded and IoT computer systems, giving confidence that the same measurement process can be applied to most programmable devices.

During the measurements, the device executes an alternating pair of instructions, as described in Section 3.2. These instructions can be classified into two categories: *on-chip* and *off-chip*. On-chip instructions are those executed exclusively on the processor chip, without interacting with other ICs, such as the system memory. In the experiments, addition, subtraction, division, multiplication, on-chip load, and on-chip store are examples of such instructions. These instructions are referred to as ADD, SUB, DIV, MUL, LDL1, and STL1 respectively. For consistency, the same operands are used for all the instructions, and the same registers are used for storage.

The off-chip instructions used in the experiments are off-chip load and off-chip store (LDM and STM). To ensure the off-chip instructions require the processor to interact with the external memory, the addresses used for the LDM and STM are specified to be located on the external memory. For consistency, the same addresses are used for all off-chip instructions. In cases where the memory address cannot be explicitly defined, the size of

the operand is set to be too large to fit on the cache, forcing the use of the external memory.

3.5.3 Localization of Sources Created by On-Chip and Off-Chip Instructions on the DE1

In this subsection, the localization results for the DE1 when it executes DIV/ADD, MUL/ADD, and LDM/DIV are presented. The alternating frequency for all measurements in this and the following subsections is 156 kHz. DIV/ADD and MUL/ADD use only on-chip instructions, while LDM/DIV uses one off-chip and one on-chip instruction.

The last step in the verification of the proposed localization algorithm is to verify if the estimated locations of sources are physically meaningful. While the magnetic dipoles identified by the technique are only approximations of the actual time-varying components of currents, the equivalent dipoles and the actual currents need to produce similar magnetic fields not only along the perimeter of the PCB (where the test points are located), but also anywhere above the PCB (except extremely close to the affected chips, their pins, and decoupling capacitors).

To demonstrate that the sources determined by the algorithm are good approximations of actual sources, we compare the measured fields when the FPGA executes DIV/ADD with simulated magnetic fields (obtained in Ansys Maxwell and confirmed by AWAS) when small loops are positioned in the locations found by the localization algorithm.

The total measured magnetic field of DIV/ADD is shown in Fig. 3.9 (left), and the physical locations of the sources on the DE1 are shown in Fig. 3.10 (right).

The locations of the sources are represented by the dark blue markers, while the white lines represent the orientation of the dipoles' moments. The light blue squares indicate the positions of the measurement points used by the algorithm. Before discussing the sources themselves, it should be noted that the magnetic field in Fig. 3.9 is strongest around the FPGA. This result supports the assumption that the instruction-dependent current will be limited to the area around the chip or chips executing the instructions. Similar results can be seen in the measurements taken for the other pairs of instructions.

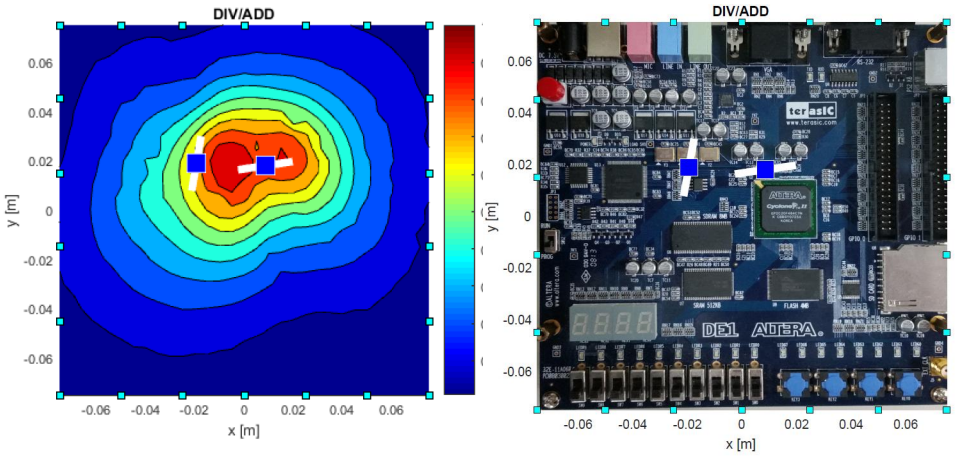

Fig. 3.9. (left) Total measured magnetic field for DIV/ADD with locations of the sources; (right) locations of the sources for DIV/ADD on the DE1.

As Fig. 3.9 illustrates, the algorithm determined two sources for DIV/ADD. Both sources are located near the FPGA's decoupling capacitors. These sources are likely caused by variations in the current being drawn by the FPGA as it switches between executing DIV and ADD. The results for DIV/SUB are nearly identical to the results for DIV/ADD (two sources located at the same locations). This similarity is unsurprising since, on the Cyclone II, SUB is a pseudo-instruction for ADD. To demonstrate the accuracy of the results for DIV/ADD, magnetic field of the sources located by the algorithm is simulated using the Ansys Maxwell simulation suite. If the algorithm is accurate, the simulated magnetic fields will be similar to the measured fields. In the simulation, the sources are represented by small loops positioned at the locations and in the orientations indicated by the algorithm. For simplicity, the PCB is represented by a finite metallic foil having the same dimensions as the board. To match the measurements, the magnetic field is simulated 3 cm above the board. The total measured magnetic field and the measured x, y, and z components of the magnetic field are shown in Fig. 3.10, while the simulated total magnetic field and its components are shown in Fig. 3.11. Comparing the simulated and measured fields illustrates that shape and magnitudes of the fields are very similar. This similarity indicates that the sources found by the algorithm produce a similar magnetic field over the entire area of the

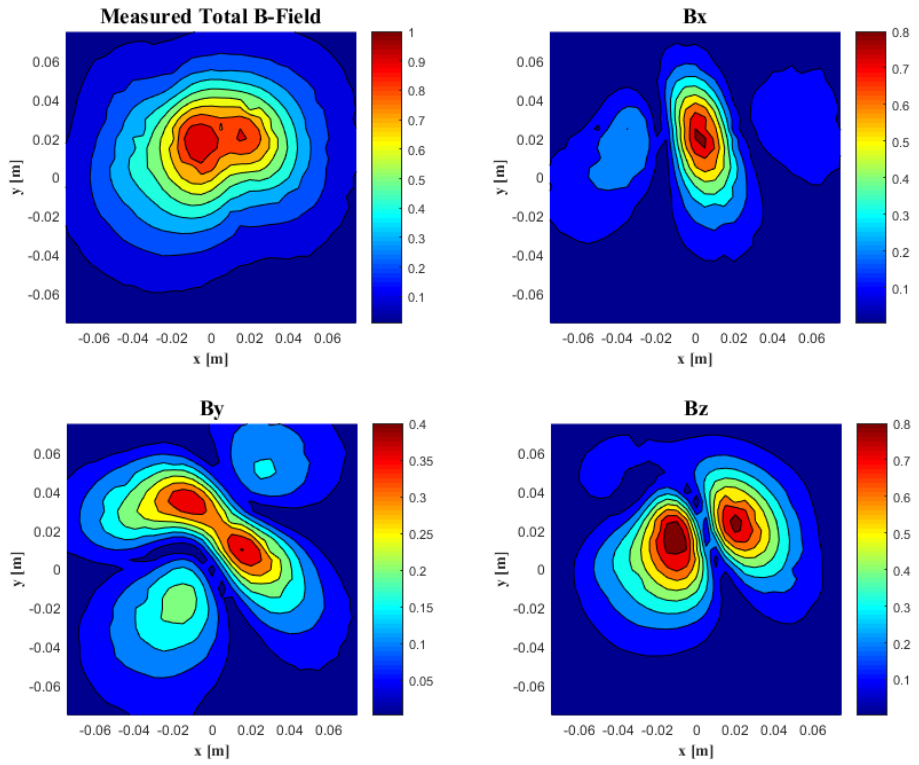

Fig. 3.10. The measured total magnetic field and the x, y, and z components of the magnetic field generated by DIV/ADD on the DE1.

PCB and not only at the points used by the algorithm.

Next, the measurements taken when the FPGA executes MUL/ADD are evaluated. As shown in Fig 3.12 (left), the algorithm identified one dominant source. This result contrasts with the two sources found for DIV/ADD, but it is not surprising: [43] demonstrated that the DIV instruction has much higher power consumption compared to ADD, SUB, and MUL. Fig 3.12 (right) shows that the physical location of the source found for MUL/ADD is near one of the FPGA's decoupling capacitors.

Next, the measured magnetic fields and the sources determined by the algorithm for LDM/DIV are shown in Fig. 3.13. Comparing LDM/DIV to DIV/ADD highlights both the similarities and differences between the on- and off-chip instructions. As with the findings for the on-chip pairs of instructions, the localization algorithm identified two sources near the decoupling capacitors at the upper left corner of the FPGA. This result is expected

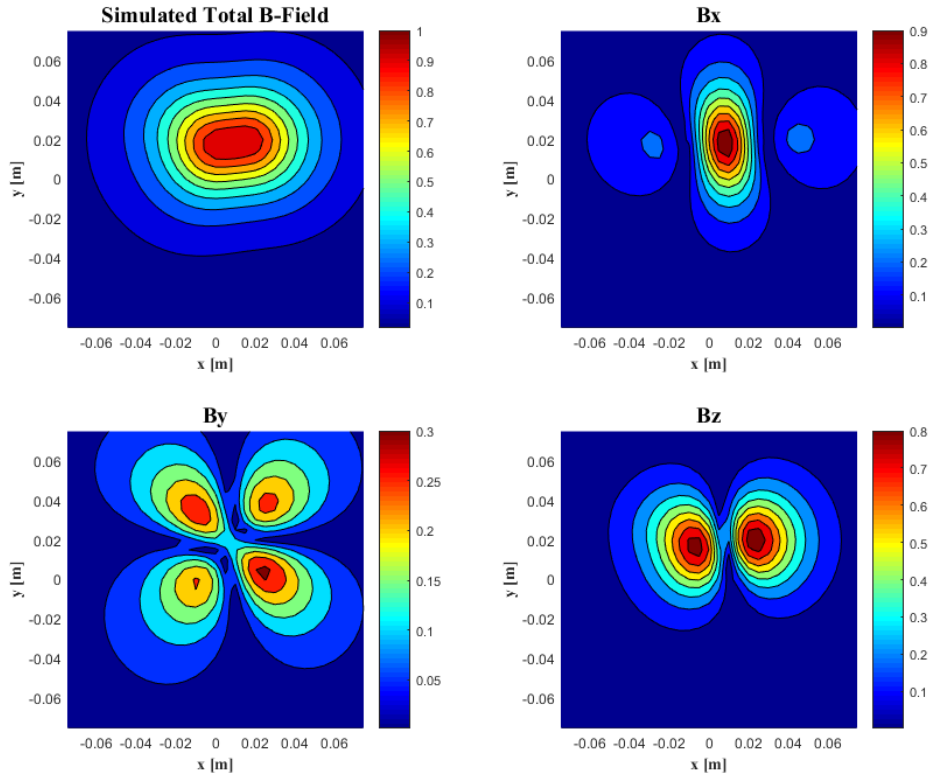

Fig. 3.11. The simulated total magnetic field and the x, y, and z components of the magnetic field generated by DIV/ADD on the DE1 using the optimal location of the sources found by the algorithm.

since DIV is still being executed on the FPGA. However, unlike findings for the on-chip instructions, the algorithm found a third source located much further away from the FPGA, near the SDRAM chip that serves as the FPGA's external memory. This source is likely caused by the FPGA communicating with the SDRAM when executing the off-chip store.

3.5.4 Localization of Sources Created by Instructions Modulated onto FPGA Processor Clock

As discussed in Section 3.2, the processor clock signal can be modulated by periodic activity in an executed program. To generate the modulated emanations, the FPGA is again programmed to execute an alternating pair of instructions. However, instead of measuring the emanations at the alternation frequency, measurements are taken near the FPGA's clock

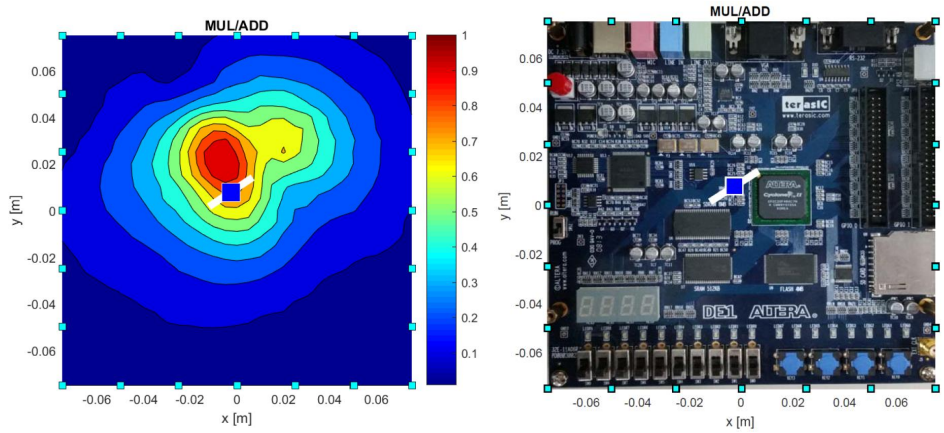

Fig. 3.12. (left) Total measured magnetic field for MUL/ADD with the location of the source; (right) the location of the source for MUL/ADD on the DE1.

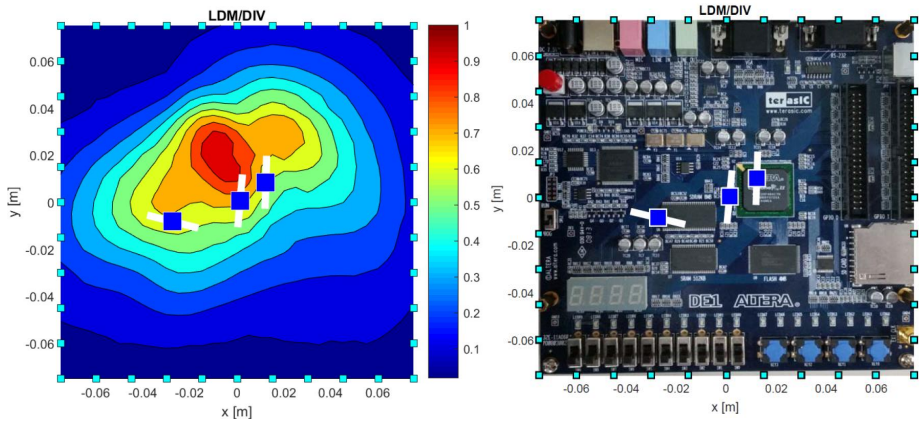

Fig. 3.13. (left) Total measured magnetic field for LDM/DIV with locations of the sources; (right) locations of the sources for LDM/DIV on the DE1.

frequency. Since the clock is modulated, there are upper and lower sidebands at $+f_{alt}$ and $-f_{alt}$ away from the clock frequency. Measurements of the sidebands are used by the algorithm to determine the source locations. For all the following measurements, the FPGA's processor clock operates at 50 MHz, and f_{alt} is 20 kHz.

Note that for the following subsection the algorithm did not use the measurements taken along the outer edge of the board to determine the source locations. Instead the algorithm used points located along a smaller 10 cm by 10 cm square contained within the area of the FPGA. This change is made because the magnitude of the magnetic field along the edges are too low to distinguish from noise. As a result, points closer to the sources need to be

used.

The first pair of instructions evaluated is DIV/ADD. The measured magnetic fields and source locations are shown in Fig. 3.14 (left) and Fig. 3.14 (right). From the measurements, the algorithm determined that DIV/ADD has two sources at the decoupling capacitors along the top edge of the FPGA.

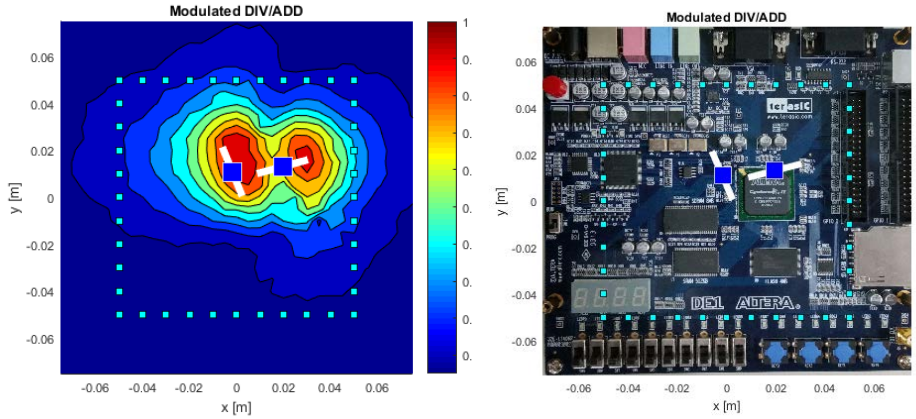

Fig. 3.14. (left) Total measured magnetic field for modulated DIV/ADD with locations of the sources; (right) locations of the sources for modulated DIV/ADD on the DE1.

Fig. 3.15 shows the measured magnetic field and the source locations for LDM/LDL1. The algorithm has identified two sources, one at the decoupling capacitors near the right edge of the SDRAM and the other at the decoupling capacitors in between the FPGA and SDRAM.

3.5.5 Localization of Sources on the A13-MICRO

Next, the localization results when the A13-MICRO executes MUL/ADD and MUL/SUB are presented. The DIV instruction is not used since the device does not have a separate division instruction. The alternation frequency for all measurements is again 156 kHz. Here, we illustrate how the localization algorithm works on a more complex device that not only runs the excitation benchmark, but also has an operating system active. The results again demonstrate that the model can locate the sources of EM side-channel leakage.

The measured total magnetic field of MUL/ADD and the sources found by the local-

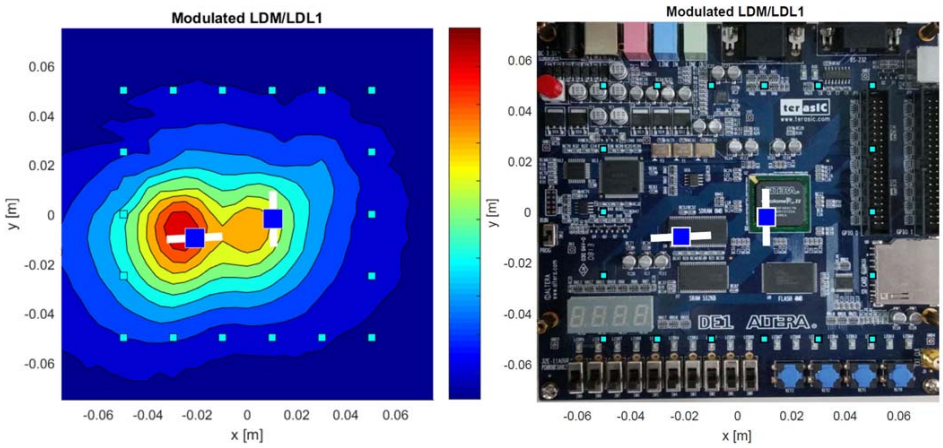

Fig. 3.15. (left) Total measured magnetic field for modulated LDM/LDL1 with locations of the sources; (right) locations of the sources for modulated LDM/LDL1 on the DE1.

ization algorithm are shown in Fig. 3.16 (left). Fig. 3.16 (right) shows that one source is located at the bottom left part of the board, near the power supply circuitry. A second source is located near the center of the board, between the processor and its decoupling capacitors.

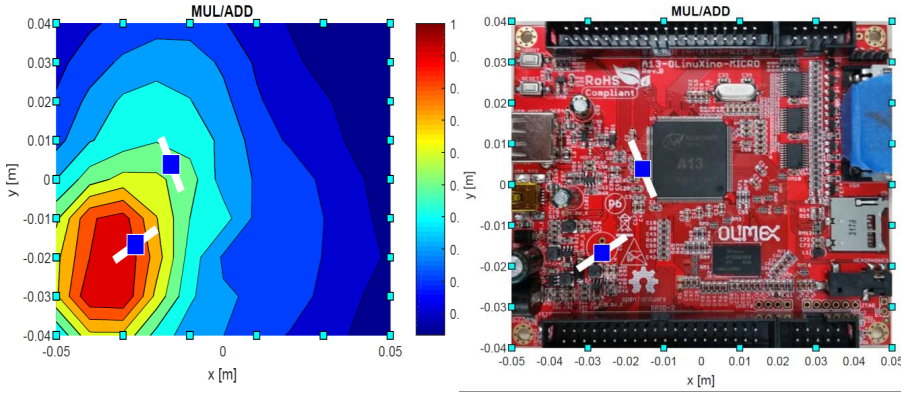

Fig. 3.16. (left) Total measured magnetic field for MUL/ADD with locations of the sources; (right) locations of the sources for MUL/ADD on the A13-MICRO.

Next, the magnetic field when the A13-MICRO executes MUL/SUB is shown in Fig. 3.17 (left). The sources determined from the measurements are shown in Fig 3.17 (right). Comparing the source found in Fig. 3.17 (right) to Fig 3.16 (right) demonstrates that the locations of the sources for MUL/SUB are close to the sources for MUL/ADD, near the processor's decoupling capacitors and the power supply circuitry.

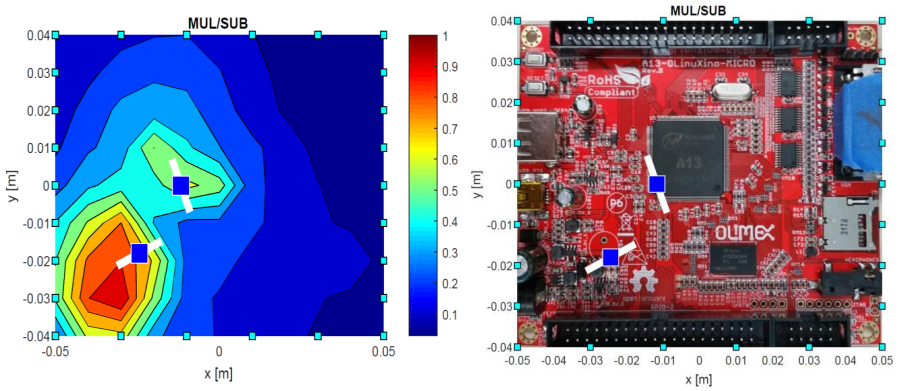

Fig. 3.17. (left) Total measured magnetic field for MUL/SUB with locations of the sources; (right) locations of the sources for MUL/SUB on the A13-MICRO.

3.6 Conclusions

This chapter proposes a method for efficiently locating the instruction-dependent sources (the sources of the EM side channel) on a PCB using a limited number of measurements. The localization results are first verified using simulations, then tested when noise is added to the simulation results, and finally verified against measurements on FPGA and IoT development boards. The results show that the number and location of strong magnetic field sources on a board depends on the instructions used to excite the board. Furthermore, the results demonstrate that the proposed localization algorithm can accurately identify those sources, regardless of the frequency at which the measurements are conducted and the instruction pairs that are executed. Finally, compared to scanning over the entire face of the PCB, this method achieves a 15-fold decrease of the overall measurement time.

CHAPTER 4
AN EFFICIENT METHOD FOR LOCALIZATION OF MAGNETIC FIELD SOURCES THAT PRODUCE HIGH-FREQUENCY SIDE-CHANNEL EMANATIONS

4.1 Overview

While the method presented in the previous chapter is effective, it is only intended for localizing low-frequency instruction dependent sources, i.e., the low frequency sources of the EM side channel. Information about low-frequency leakage is important; however, high-frequency side channels can be more useful for attacks. The high-frequency emanations are not always the direct result of the program activity itself. Instead, they can be caused by the program activity modulating the periodic synchronizer signals (such as the device clock) already present in the device [13]. Furthermore, high-frequency measurements allow for wider bandwidth and can potentially improve the measurement distance (if the signal is a modulated synchronizer).

Given these advantages, high-frequency emanations can be more important for hardware security than the low-frequency emanations. To assist researchers studying side channels and vulnerabilities of PCBs, this chapter outlines a new low-cost system for locating the strongest sources of the high-frequency EM side channel on PCBs. While this system is based on the one described in Chapter 3, it has been extensively modified to address the challenges that come with working in the hundreds of megahertz-to-gigahertz frequency range. This system has been carefully designed to minimize the impact reflections and interference have on the magnetic field measurements.

As before, the potential side-channel leakage is evaluated by focusing on the basic instructions that commonly comprise programs. Furthermore, only a limited number of

measurements are needed to locate the leakage sources, drastically speeding up the measurement time. This simplification is accomplished by updating the localization algorithm originally developed in [12] to search for magnetic loop and electric-dipole sources.

The rest of this chapter is organized as follows. Section 4.2 describes the localization system used for recording the EM side-channel emanations at gigahertz frequencies. In Section 4.3, the system is validated by comparing the measurements with simulated results. Section 4.4 describes the new high-frequency tests and tested devices. In Sections 4.5 and 4.6, the setup is used to locate the high-frequency (around $1\ \mathrm{GHz}$) EM side-channel sources on two common types of devices. The newly identified sources are compared to previously identified sources on the same devices taken at lower frequencies. Finally, Section 4.7 summarizes the chapter.

4.2 Localization System

The following subsections describe how the EM side-channel signals from a device are generated, the setup used for measuring the high-frequency emanations, and the localization algorithm designed to determine the sources from the measurement data.

4.2.1 Side-Channel Signal Generation

While the EM side-channel signals are a subset of the EMI generated by the device, they can be quite challenging to address. Given their nature, the side-channel signal is dependent on the activity being performed by the device. As demonstrated in Chapter 3, during the execution of a program, the side-channel's sources can shift position across the device as different parts of the program require different hardware for executed. As a result, it can be difficult to determine the likely leakage sources if focusing on a specific piece of code.

To avoid this problem, the same benchmark described in Subsection 3.2 is used to generate side-channel leakage. Instead of trying to observe the signals generated by specific programs, the leakage is generated by executing a pattern of two instructions repeatedly.

This benchmark makes it possible to relate the side-channel sources to specific instructions. Identifying these sources highlights potential vulnerabilities for any program that uses those instructions and indicates which hardware components leaks the most.

Importantly, the excitation signal also amplitude modulates other periodic signals generated by the device [13]. As discussed earlier, the synchronizer clocks of the components used for executing the program activity can act as carriers for the modulated waveform. During normal operations, the clocks produce periodic currents at the clock frequency f_{C} along the signal and power traces. When the excitation program is executed, the periodic current from the clock is modulated by the current drawn when executing the program activity. If the excitation program has a frequency of f_{alt}, the modulated waveform can then be observed at $f_{\mathrm{C}} \pm f_{\mathrm{alt}}$ in the spectrum.

4.2.2 Near-Field Scanner Setup

Near-field scanners are commonly used for finding the general sources of EMI. However, the general EMI from a device is not representative of the EM side-channel emanations [106]. Furthermore, the time-varying nature of the EM side-channel signal can make it more challenging to find its emanation sources than other EMI sources. Despite these challenges, several scanners have been used to evaluate the EM side channel, such as the ones detailed in [50], [106], [107], [108]. However, all these examples evaluate the EM side-channel leakage as specific cryptographic programs are being executed on the device. While this approach is intuitive, the results sensitive to the how the program varies in time and are only relevant to the specific program being tested.

Fig. 4.1 shows the near-field scanner built to accurately measure the instruction-dependent magnetic field emanated by a device. The primary challenge in locating the sources at high frequencies is that the sizes of the equipment used for scanning (such as the probe, cables, and the scanning apparatus) impact the measurements, and the interference from external devices becomes more pronounced. These factors need to be carefully accounted for when

measuring the fields.

To address these challenges, the setup is comprised a 2-D plotter shielded by a hand-made anechoic box. The anechoic box is made of four RF absorber panels, three positioned at the sides, and the fourth positioned on the bottom of the plotter. These panels are SFC-4 RF absorbing panels from Cuming Microwave [109]. The fourth side is left open for accessing the device-under-test (DUT) in between the tests. During measurements, the fifth panel is positioned at the open side, leaving only the top of the setup open.

This configuration creates a roughly $61 \times 53 \times 61$ cm^3 box for testing, where the DUT is placed at the center of the bottom panel. Cables for powering and controlling the DUT are run through a small hole in the bottom panel. The need for absorber panels stems from the noticeable coupling between the plotter mechanical and electronic parts, cables, connectors, and the DUT. The positioning of the absorbers is carefully designed to minimize the influence of other electronics near the testing equipment as well as any field interference created due to the plotter, cables, and other electronics.

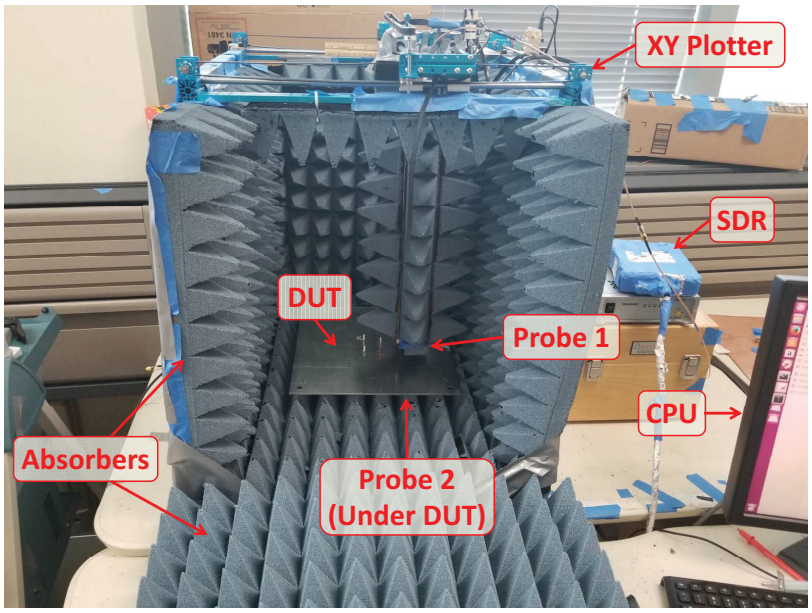

Fig. 4.1. The measurement setup. The DUT is the aluminum plate used in Section 4.3. "Probe 1" is the moving probe and "Probe 2" is the stationary probe.

Two probes, whose design is described in Subsection 4.2.3, are used for collecting the

magnetic field generated by the DUT at two locations. One probe, the moving probe, is used for measuring the emanated magnetic field in the plane above the DUT. A MakeBlock XYPlotter Robot Kit v2.0 is used for scanning the moving probe at a fixed height over the DUT. This plotter has an x-y accuracy of $0.1\ \mathrm{mm}$ for positioning. To minimize the influence of the plotter on the measurements, it is mounted on top of the box, putting it more than $40\ \mathrm{cm}$ away from the DUT. At the same time, the absorbing material is wrapped around the moving probe's cable to decrease the parasitic current flowing on the outer conductor [55].

To emulate an isotropic receiver, three sets of measurements are taken for each test case. In each set, the moving probe is oriented to record to the x-, y-, or z-components of the magnetic field, respectively. During a set of measurements, the magnitude of the field is recorded by the moving probe as it is scanned at a fixed height and with a step size of $0.5\ \mathrm{cm}$ in a plane over the device.

The second probe remains at the same position for all the measurements. This stationary probe is used for measuring the reference signal at the same time as the moving probe's measurements [110]. The stationary probe needs to be positioned where it can detect the signal of interest, not interfere with the moving probe's measurements, and remain in the same location when measuring each component of the magnetic field. For the experiments detailed below, the stationary probe is attached to the backside of the device. This location was convenient since the device's power planes limited the influence the stationary probe had on the moving probe's measurements.

The phase of the magnetic field is calculated by subtracting the phase of the moving probe signal from the phase of the stationary probe signal for each measurement position [110]. Note that the addition of the phase is a key difference from the measurement setup proposed in [12], where only the magnitude of the signal was collected. As detailed in Subsection 4.2.4, the localization algorithm needs to both the magnitude and the phase of the fields to accurately identify the high-frequency sources.

The signal induced in both probes is recorded using a two-port B210 software-defined

radio (SDR), from Ettus Research [111]. The advantage of the B210 is that it can measure two channels at the same time. During the test, a computer is used for controlling the measurement equipment and the DUT.

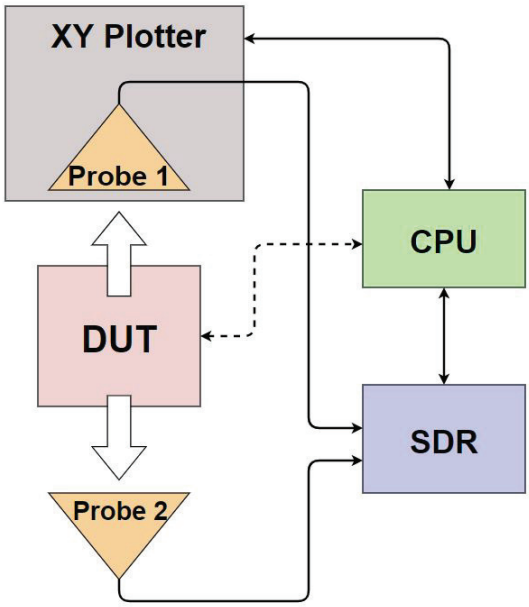

Fig. 4.2. Diagram of the measurement setup.

This system is very low cost. The plotter costs $300, the five absorbing panels cost $300, the two probes cost less than $5, and the SDR costs $1259. For comparison, a similar setup, Riscure's EM Probe Station, costs roughly $19 000, not including the cost for the software to run it or the measurement equipment. It also does not include anything to shield the measurements from external interference. A diagram of the complete measurement setup is shown in Fig. 4.2.

4.2.3 Probe Design

While the requirements for the stationary probe are not particularly stringent, the properties of the moving probe can drastically impact the accuracy of the measurements. The strictest requirement is that the probe needs to be kept several centimeters above the DUT in order to avoid hitting some of the taller components, such as electrolytic capacitors, VGA con-

nectors, and pin headers. Therefore, the moving probe is positioned $3.5\ \mathrm{cm}$ above the DUT for all measurements.

Because of the relatively large test distance, a highly sensitive probe is needed to adequately measure the field. Therefore, the shielded-loop probe described in [112] is used, shown in Fig. 4.3. The advantage of a shielded-loop probe is that it partly suppresses the influence of the electric field when measuring the magnetic field [113].

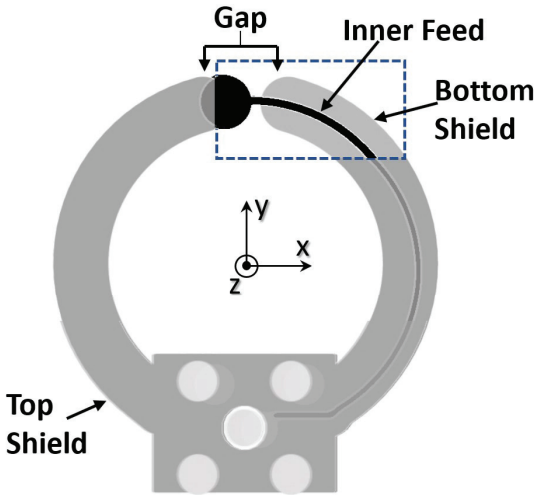

Fig. 4.3. The shield-loop probe used for measurements.

The drawback of the probe is that its size makes it sensitive not only to the magnetic fields, but to the electric fields as well [114]. However, the y-component of the electric field is suppressed by the shield and the z-component is minimized given that the probe is only a few millimeters thick. On the other hand, the electric field in the x-direction still affects the measurements. Its impact is removed by repeating the measurements with the probe rotated 180° around the z-axis [114], [115]. The probe will still measure the same magnetic field and electric field component at both orientations; however, the sign of the signal (in phasor form) induced by the electric field is reversed. The magnetic field can then be found by adding the measurements (to cancel the influence of the electric field) and dividing the result by two.

4.2.4 Localization Algorithm and EM Source Model

In [12], an algorithm for efficiently locating the sources of the magnetic field created by the program activity on a PCB was developed. This algorithm is based on Nelder-Mead simplex optimization [103]. When the number of sources is unknown, the algorithm will run several times, assuming different numbers of sources. The location, intensity, and orientation of each source are the optimization variables. Within the optimization process, the values of these variables are randomly selected from within an allowable range (for example, the source must be located on the DUT). These values are used as an initial solution to the problem. To avoid being biased to a specific location, the algorithm generates several hundred initial solutions. Initial simplexes are created from the individual initial solutions, where each simplex has a nonzero volume. Each simplex is optimized over a limited number of iterations. Afterwards, the solution which generates a magnetic field that best matches the measured field is reported.

In the algorithm, the sources of the measured fields are approximated by elementary dipoles. Representing the sources as simplified elements decreases the complexity of the model and makes it possible to simulate complex designs [116]. This approach is commonly used with near-field scanning when modeling the EMI from electronic devices [116]–[118]. The specifics of the approach vary based on the application and desired accuracy.

The advantage the localization algorithm is that it requires only a small number of measurements, thus drastically decreasing the measurement time. Generally it is more practical to use the measurements taken around the edge of DUT. Using the edges decreases the chance the probe will collide with the taller components and limits the interference from components on the DUT. Furthermore, concentrating the measurements around the edges of the DUT makes them more sensitive to smaller changes in the field compared to measuring the same number of points over the entire face of the DUT. As an example of the amount of time the algorithm saves, in the experiments in Section 4.5, it took roughly

20 times longer to measure the entire face of a device than it took to only measure the outer edge of the DUT. The algorithm makes this simplification by identifying only the locations of the strongest sources of emanations.

The types of approximated sources and their relationship with the measured field are based on the frequency. In [12], the sources and the measurement distance were small compared to the wavelength, making it possible to assume that *the magnetic field was quasistatic*. Therefore, the sources were approximated by small magnetic loops. The influence of electric-dipole elements was ignored since, at low frequencies, the magnetic field generated by electric dipoles is negligible compared to the field from loops.

However, the quasistatic approximation no longer holds when localizing high-frequency sources. As the frequency increases, the magnetic field emanated from *electric*-dipole structures is no longer negligible. In order to accurately represent an arbitrary source, both magnetic loops and electric dipoles are needed. Therefore, for high-frequency side-channel emanations, the quasistatic equation for a loop is replaced by the full equations for the magnetic fields from a small loop and small electric dipole. The magnetic field of a loop is

$$\mathbf{H} = \frac{(\mathrm{j}\beta)^2}{4\pi r}\left(\left(1+\frac{3}{\mathrm{j}\beta r}+\frac{3}{(\mathrm{j}\beta r)^2}\right)(\mathbf{m}\cdot\hat{\mathbf{r}})\hat{\mathbf{r}} - \left(1+\frac{1}{\mathrm{j}\beta r}+\frac{1}{(\mathrm{j}\beta r)^2}\right)\mathbf{m}\right)\mathrm{e}^{-\mathrm{j}\beta r}, \quad (4.1)$$

where β is the wavenumber, $\mathbf{m}$ is the magnetic moment vector, the vector $\mathbf{r}$ defines the location of the observation point with respect to the source, r is the radial distance between the source and the observation point, and $\hat{\mathbf{r}}$ is the unit vector of $\mathbf{r}$. Both types of sources are still considered to be small compared to the wavelength and the measurement distance, r. The magnetic field for an electric dipole is

$$\mathbf{H} = \frac{c(\mathrm{j}\beta)^2}{4\pi r}(\mathbf{p}\times\hat{\mathbf{r}})\left(1+\frac{1}{\mathrm{j}\beta r}\right)\mathrm{e}^{-\mathrm{j}\beta r}, \quad (4.2)$$

where **p** is the electric moment of the dipole and c is the speed of light in a vacuum. When attempting to locate the sources, the algorithm now tests multiple combinations of magnetic loops and electric dipoles before selecting the solution that best matches the measurements.

Furthermore, PCBs commonly have one or more ground planes. By default, the algorithm assumes that any identified sources are situated above this ground plane. Therefore, the magnetic field generated by an image source located below the predicted loop or electric-dipole source is included in the calculations. The only exception is in the setup tests in following section. Since electric monopoles are used for the test sources, the algorithm is modified to take into account their images in the ground plane, resulting in a set of symmetrical dipoles without the ground plane.

4.3 Validation of Measurement Setup

To validate the accuracy of the measurement setup, several controlled cases are tested. In the tests, one or more SMA connectors are used as the magnetic-field sources [119]. These connectors are mounted to the bottom of a large, $3\ \mathrm{mm}$ thick aluminum plate. The connector's inner conductor and its Teflon insulation are run through a hole in the plate. The conductor has a diameter of $1.79\ \mathrm{mm}$ and a length of $17.87\ \mathrm{mm}$. Mounting the connector to the plate ensures that the outer conductor of the connector and the plate are at the same potential. In this configuration, the inner conductor of the SMA connector acts as a monopole antenna.

During the measurement, the SMA connector is excited at $1.008\ \mathrm{GHz}$, using an Agilent MXG N5183A function generator, and the magnetic field is recorded using the setup in Subsection 4.2.2. The SMA connectors are used as test sources since they can be modeled precisely as monopoles. As discussed in Section 4.2.4, the magnetic fields generated by electric dipoles (or monopoles) cannot be ignored at high frequencies. Therefore, the SMA connectors are appropriate and important for testing the setup.

Field measurements are taken in $5\ \mathrm{mm}$ steps in a rectangular grid spanning a $10 \times 8\ \mathrm{cm}^2$

plane over the aluminum plate, where the center of the rectangle is considered the origin of the coordinate system. While the algorithm needs only the measurements around the edge of the rectangular area, the inner measurements are used for evaluating the measurement accuracy. After the tests are complete, the measured fields are compared to the fields simulated using the Ansys HFSS simulation suite [120]. The localization algorithm from Subsection 4.2.4 is then run on the measurement and simulation data, and the results are compared.

The measured and simulated magnetic fields for three test cases are shown in Figs. 4.4–4.6. In each figure, the magnitudes of the total magnetic field and the x-, y-, and z-components of the magnetic field are shown separately. In the first case, a single monopole is placed at the center of the measurement area, while in the second case, a single monopole is placed at (1.5 cm, -1.4 cm). The third test case includes two monopoles, one placed at (-1.7 cm, 2.3 cm) and the other at (2.5 cm, 1.8 cm). In the two-source case, a splitter is used to supply an excitation signal with same magnitude and phase to each monopole. In each figure, the correct location of each monopole source is indicated using a yellow star, while the location of the source found by the localization algorithm is indicated using a white triangle.

As Figs. 4.4–4.6 demonstrate, not only do the measured and simulated fields visually match, the localization algorithm is able to accurately determine the location of the monopoles. In the first case, the position error is 2.2 mm for the measurements and 1.6 mm for the simulations. In the second case, the position error is 3.2 mm for the measurements and 3.7 mm for the simulations. For the two-source case, the average position error is 5.6 mm for the measurements and 2.3 mm for the simulations. Hence, here we have demonstrated that the algorithm is able to correctly locate the sources using high-frequency measurements or simulations.

Additionally, Fig. 4.7 shows the measurements for the first test case prior to the corrections discussed in Subsection 4.2.3. In addition to the measured fields not visually match-

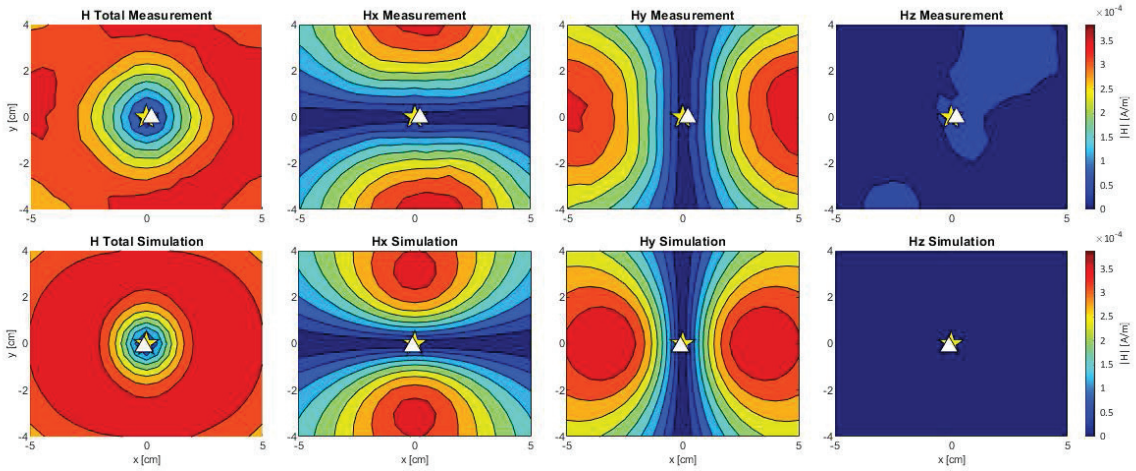

Fig. 4.4. Comparison of the components of the measured (top) and simulated (bottom) magnetic fields for the first test case.

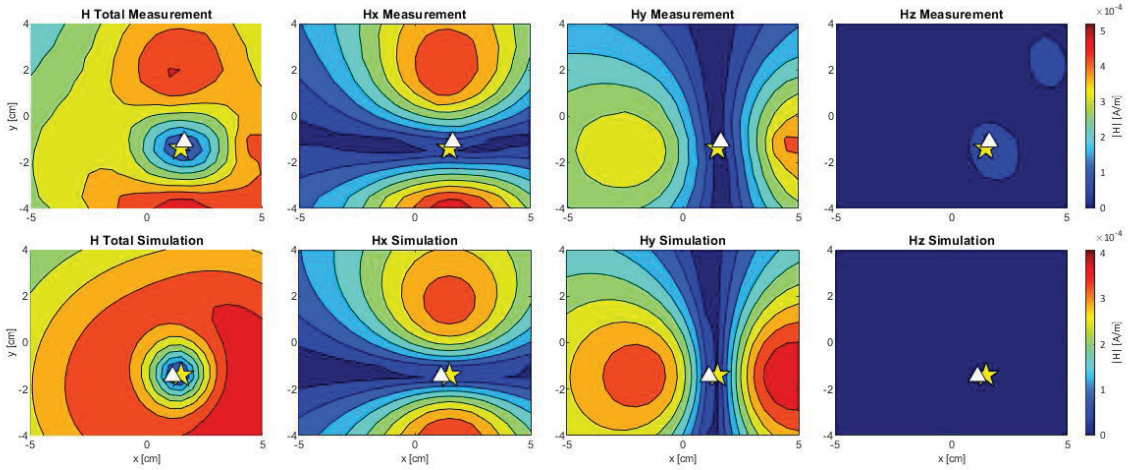

Fig. 4.5. Comparison of measured (top) to simulated (bottom) magnetic fields for the second test case.

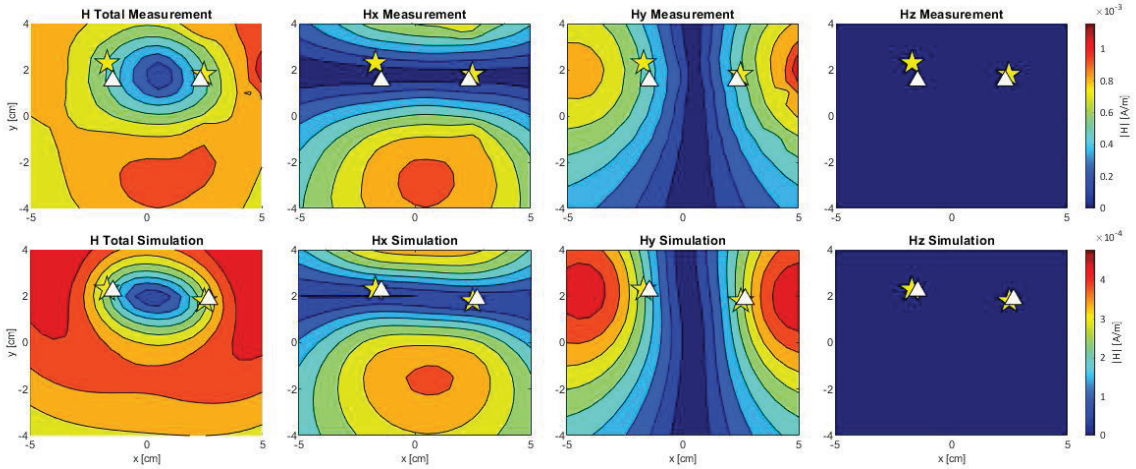

Fig. 4.6. Comparison of measured (top) to simulated (bottom) magnetic fields for the two-source test case.

ing the simulated fields, the source determined by the algorithm has a relative error of $34.2\ \mathrm{mm}$. This inaccuracy is likely the result of the electric field's influence on the probe's measurements, demonstrating the importance of the corrections. Similar errors were seen in other measurements taken without the modifications to the setup described in previous section. These errors caused the algorithm to be unable to accurately identify the location and number of sources.

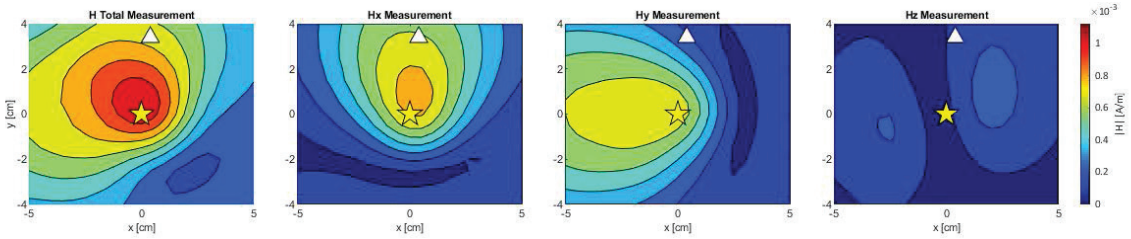

Fig. 4.7. The measured magnetic field in the first test case before correction. (Compare with Fig. 4.4.)

4.4 Localization of Sources of High-Frequency Emanations on PCBs - Experimental Results

In the following sections, we demonstrate that the proposed localization system can identify sources of high-frequency EM side-channel emanations and show that these sources differ from those observed from low-frequency EM side channels. The low-frequency baseband fields and sources originally measured in [12] are compared to the new measurements of high-frequency fields. Two devices are tested: a Cyclone II DE1 FPGA development board from Altera and Terasic [121] and an A13-OLinuXino-MICRO IoT Linux computer board from Olimex [122]. For simplicity, the devices are referred to as the DE1 and the A13-MICRO for the rest of this chapter. Pictures of the two devices with relevant components labeled are shown in Fig. 4.8. These devices are good representatives of available embedded and IoT devices that are used for a multitude of applications. However, their security is a major concern given the speed at which they enter the market and their limited resources [123].

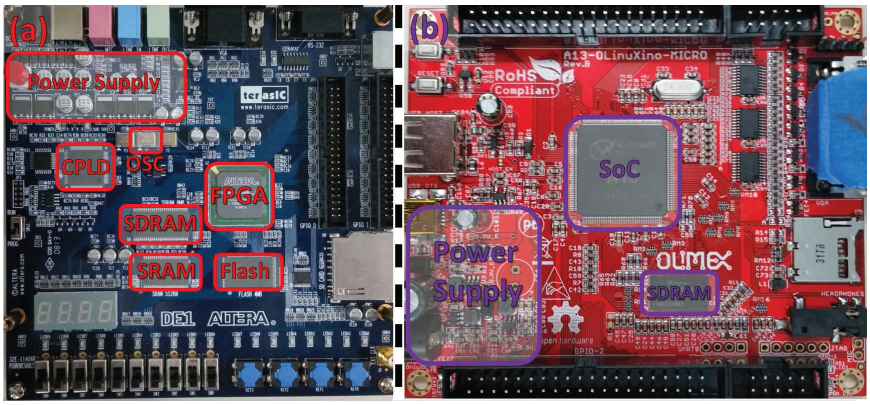

Fig. 4.8. The (a) DE1 development board and (b) A13-MICRO with relevant components labeled.

The same measurement procedure described in Section 4.2 is used for all the 1 GHz measurements. During the measurements, the DUT is excited using the excitation program detailed in Subsection 4.2.1. As discussed in Subsection 4.2.4, only the measurements taken around the edge of the DUT are used by the algorithm to locate the sources. In the experiments, the fields from several different instruction pairs are recorded and compared. Only a small number of measurement results are included to illustrate the proposed localization system. For measurements taken below 100 MHz, the localization algorithm identifies the locations of quasistatic loop sources. For the high-frequency measurements, the algorithm identifies both loop and electric-dipole sources using (4.1) and (4.2).

The excitation frequency for the baseband measurements is 156 kHz, while the excitation frequency for the modulated measurements is 20 kHz. These frequencies are selected to avoid interference during measurements. For the modulated emanations, the carrier and the first harmonics of the upper sidebands are shown.

4.5 DE1 Experiments

The processor and memory components are the most active components during excitation. On the DE1, the Cyclone II FPGA serves as the processor. The Cyclone II implements a Nios-II soft processor. It uses an 8 MB SDRAM as the external memory. The clock for

the FPGA is provided by a 50 MHz crystal oscillator. The FPGA then provides a 50 MHz clock to the SDRAM. Other notable components include the power supply circuitry, flash memory, SRAM, and the complex programmable logic device (CPLD) used in the USB circuitry. These components are labeled in Fig. 4.8(a). In the experiments, the fields generated by the DE1 are recorded in the baseband, modulating the 50 MHz clock frequency, and modulating the 20th harmonic of the clock (1 GHz), while the board is excited using different instruction pairs.

4.5.1 Baseband DE1 Measurements

Examples of the baseband DE1 measurements are shown in Fig. 4.9. The fields recorded while the device executes LDM/LDL1, STM/DIV, DIV/ADD, and MUL/SUB are included. In the figure, the white squares represent the quasistatic loop sources. The measurements demonstrate that the emanation sources and the fields vary noticeably based on the instruction being executed. The positions of the loop sources change because different parts of the FPGA are used for executing different instructions. For example, the fields for LDM/LDL1 are concentrated around the top left side of the FPGA and the SDRAM. The algorithm identified three loop sources based on the measured field, one near the left side of the SDRAM, one between the FPGA and SDRAM, and one near the top right corner of the FPGA. The latter two sources are located near the decoupling capacitors and their traces for the FPGA. These capacitors are connected to the 3.3 V power supply. The 3.3 V supplies power to multiple components, including the FPGA, SDRAM, and the 50 MHz oscillator.

On the other hand, the fields and sources for STM/DIV are concentrated around the top left corner of the FPGA. The right-most source is again near the FPGA's decoupling capacitors, while the other source is several millimeters to the left of another set of decoupling capacitors. Despite the SDRAM being active, any emanations from it are being overshadowed.

In the case of the strictly on-chip instructions, DIV/ADD and MUL/SUB, the sources

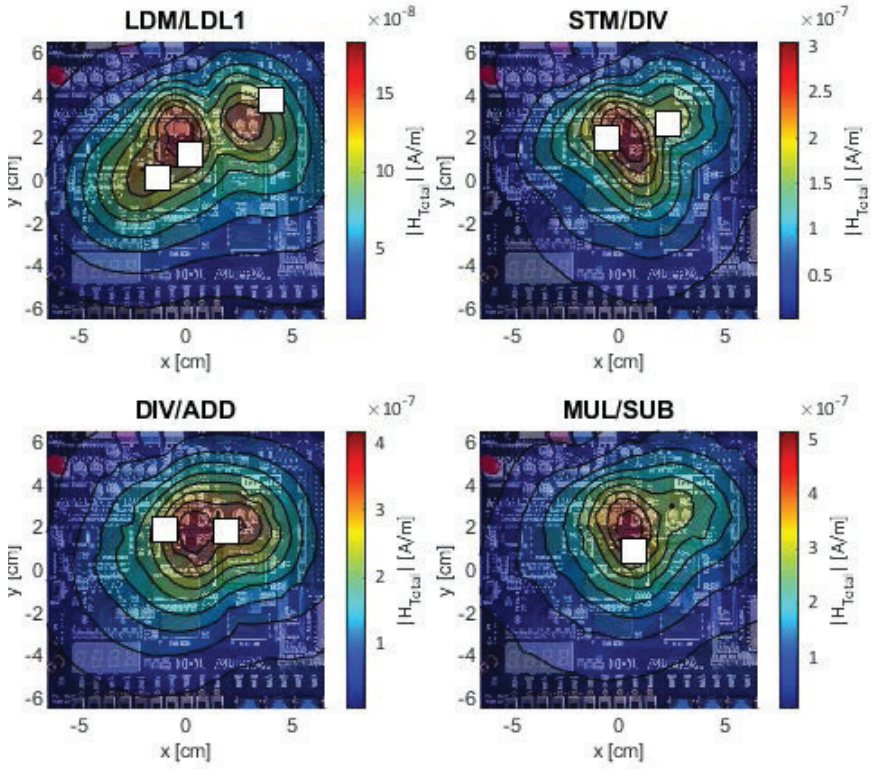

Fig. 4.9. The magnetic fields measured at the baseband (156 kHz) for LDM/LDL1, STM/-DIV, DIV/ADD, and MUL/SUB.

are constrained to the area around the top left corner of the FPGA. The algorithm located two sources for DIV/ADD and one for MUL/SUB. The right-most source for DIV/ADD and the source for MUL/SUB are located near the FPGA's decoupling capacitors. The other source for DIV/ADD is located several millimeters to the left of the decoupling capacitors.

Overall, these measurements demonstrate that, at the baseband frequency, the emanations mostly originate from the decoupling capacitors connected to the FPGA and SDRAM and the 3.3 V power supply. These emanations from the decoupling capacitors are commonly the result of a ground-bounce caused by variations in the current drawn as transistors inside the components switch state. Any emanations from the ICs themselves are being overshadowed by the emanations from outside the IC. At such low frequencies, potential emanation sources inside the IC (such as dipoles or loops formed by the IC's pins and inner traces) are too small to radiate efficiently.

4.5.2 50 MHz DE1 Measurements

Examples of the magnetic field measured at the sideband for the $50\ \mathrm{MHz}$ clock are shown in Fig. 4.10. The figure includes the sidebands measured while the device executes the same instruction pairs as in the previous subsection. Like the baseband measurements, the field distribution varies with the instruction pair. For cases involving off-chip memory instructions, LDM/LDL1 and STM/DIV, the fields are concentrated near the bottom left corner of the FPGA and the left side of the SDRAM. For both the baseband and modulated measurements, LDM/LDL1 has one quasistatic loop source near the left of the SDRAM and one between the FPGA and the SDRAM. However, unlike in the baseband, LDM/LDL1 does not have any sources near the top of the FPGA. On the other hand, the measurements for STM/DIV vary significantly between the baseband and the modulated measurements. The fields shift from being concentrated around the FPGA to being concentrated around the SDRAM.

The field and sources for the strictly on-chip instruction pairs are nearly identical. Furthermore, the differences between the baseband and modulated fields for both instruction pairs are not as pronounced as the off-chip instructions. As in the baseband, the fields are mainly concentrated around the top left corner of the FPGA. The primary difference between the baseband and modulated measurements is that in the modulated measurements for MUL/SUB, the algorithm identified a second source near the top of the FPGA.

As with the baseband measurements, the sources for the modulated emanations are located near the decoupling capacitors for the FPGA and SDRAM. While the locations of the loop sources for the on-chip instruction pairs are similar to those for the baseband, the sources for the off-chip instruction pairs have shifted to be predominately around the SDRAM. This change likely occurs because the leakage signal is now the modulated $50\ \mathrm{MHz}$ clock. As a result, instead of being directly caused by the program activity, the emanation sources are related to the parts of the device the clock interacts with after being modulated. Therefore, the emanations at the sidebands share some of the emanation

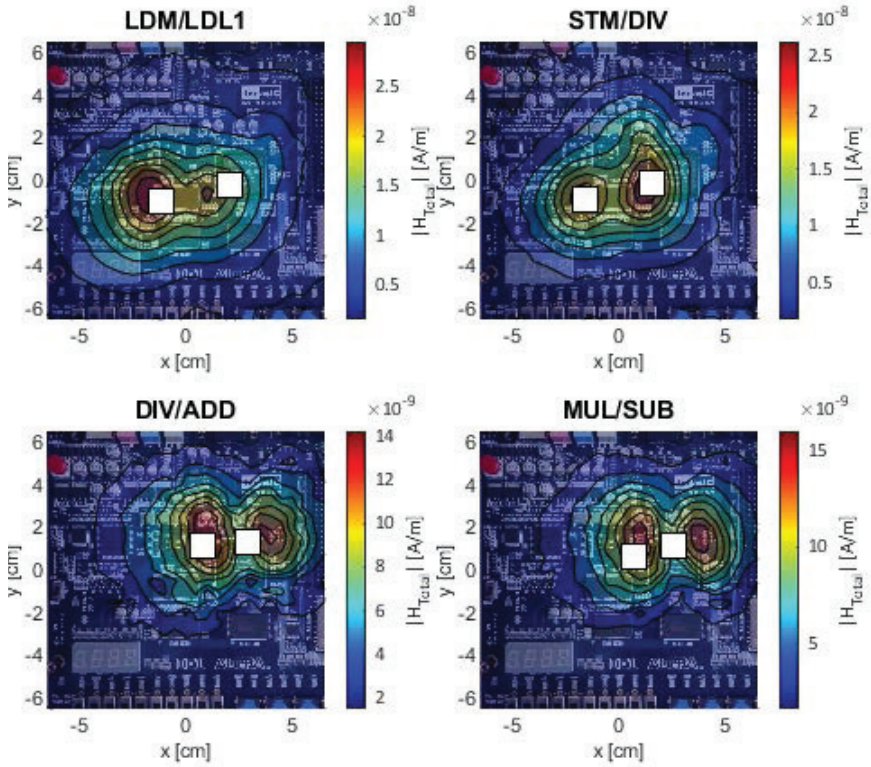

Fig. 4.10. Magnetic fields measured at the upper sideband ($+20$ kHz) of the 50 MHz FPGA clock.

sources for the clock. However, not all the clock's emanation sources will act as sources for the side channel. To demonstrate this, an example of the fields measured at the 50 MHz carrier are plotted in Fig. 4.11. Despite the field distribution at the sidebands varying with the instruction pair, the distribution at the carrier frequency remains concentrated at the same location for all instruction pairs. The source for the carrier is located between the FPGA and the 50 MHz crystal oscillator. The fact this source is not present in the previous measurements in Fig. 4.10 indicates that it is isolated from the effects of the program activity. The emanations from this source overshadow the emanations from the instruction-dependent sources. The source can be considered a source of EMI, but not of the EM side channel.

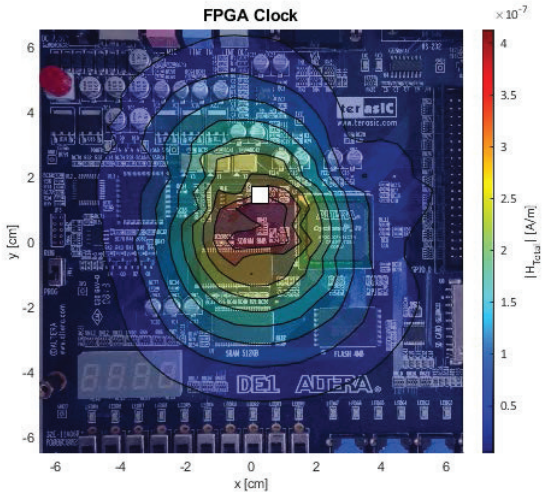

Fig. 4.11. Magnetic field measured at the 50 MHz FPGA clock.

4.5.3 1 GHz DE1 Measurements

Examples of the modulated emanations at the 1 GHz harmonic of the clock are shown in Fig. 4.12. The figure includes plots of the modulated fields measured while the device executed the same instruction pairs as in the previous subsections. As the figure demonstrates, the fields at 1 GHz vary significantly from the low-frequency measurements. This variance is partially due to the influence of electric dipoles of the magnetic fields; however, the loop sources have also somewhat changed. In the figure, the white squares represent the loop sources, while the white triangles represent the electric-dipole sources.

Similar to the 50 MHz measurements, the fields for LDM/LDL1 are concentrated near the area between the FPGA and the SDRAM. The algorithm determined that the fields had two loop and two electric-dipole sources. The first loop is located at the left edge of the SDRAM, at nearly the same location as in the baseband and 50 MHz measurements. On the other hand, a second, weaker loop source is located near the top edge of the flash. This source does not correspond to any of the sources found in the lower frequency measurements. The first electric dipole is located between the SDRAM and the FPGA, close to the FPGA's decoupling capacitors, and the loop sources found in the previous measurements. The final electric-dipole source is relatively weak and is located near the left edge of the

CPLD. This source is a trace connecting the JTAG interface of the FPGA to the CPLD. This trace is also connected to a set of pull-down resistors that are located near the other electric-dipole source for LDM/LDL1.

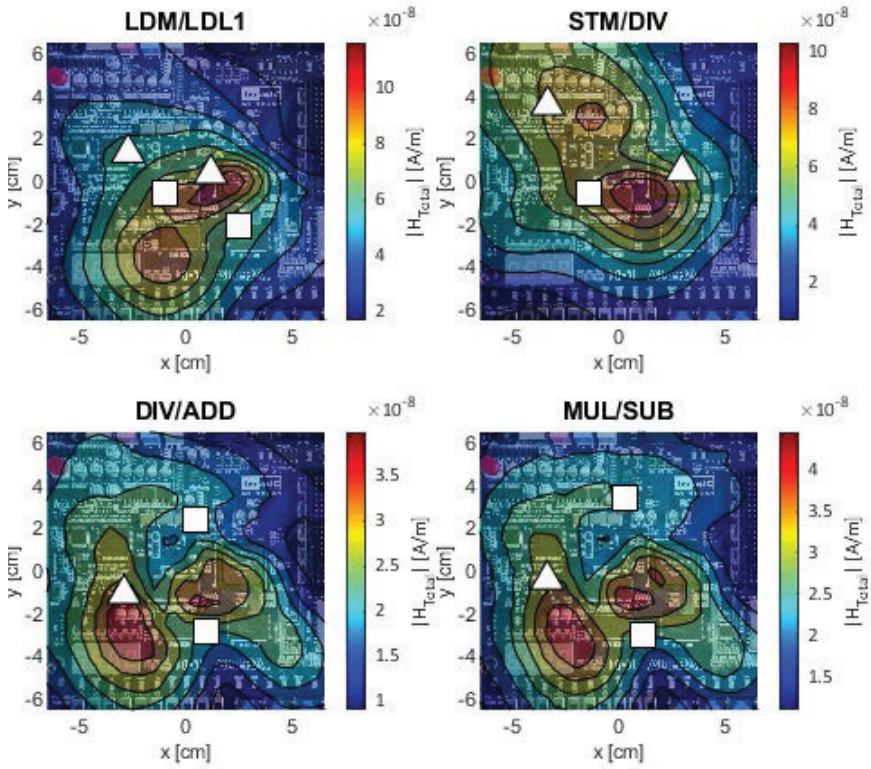

Fig. 4.12. Magnetic fields measured at the upper sideband (+20 kHz) of the 1 GHz harmonic of the FPGA clock.

As with the low-frequency measurements, a part of the 1 GHz STM/DIV fields is concentrated between the FGPA and the SDRAM. However, unlike the low-frequency measurements, the fields extend past the FPGA, near the 50 MHz oscillator and the power supply. Based on the measurements, the algorithm identified one loop source and two electric-dipole sources. The loop source is located at the left edge of the SDRAM at the same point as one of the loops found for LDM/LDL1. This source was also present in the 50 MHz STM/DIV measurement. The first electric dipole is located near the top left of the board at the power supply circuitry. The final electric dipole is located inside the FPGA itself.

Like the 50 MHz measurements, the fields and sources for DIV/ADD and MUL/SUB are nearly identical. However, the field distributions differ significantly from the 50 MHz measurements. At 1 GHz, the fields are concentrated past the bottom left corner of the FPGA, closer to the SDRAM. While not represented in the figure, the magnitude of the fields is roughly half the magnitude of the fields from LDM/LDL1. For both situations, the algorithm identified two loops and one electric dipole. The first loop source is located near the top left corner of the FPGA, near its electrolytic capacitors. The second loop is located at the right side of the SRAM. The electric-dipole source is located between the left side of the SDRAM and the bottom of the CPLD. At this location, a set of transistors connects the main 3.3V power plane to the CPLD power pins.

The results illustrate that, despite being a harmonic of the 50 MHz clock, the fields measured at 1 GHz are significantly different. This is likely the result of the radiation properties of the various structures on the board changing with frequency. At higher frequencies, the radiation efficiency of smaller structures improves, giving them a stronger impact on the emitted fields. At the same time, the magnetic field from electric-dipole elements becomes a concern. Furthermore, the higher frequency increases the coupling between nearby traces. This allows the leakage signal to spread past the power supply traces and to areas not used for executing the instruction pairs (such as the CPLD and flash).

Finally, the field measured at the carrier frequency is shown in Fig. 4.13. Like at 50 MHz, the field distribution of the 1 GHz carrier does not change with the instruction pair. Furthermore, like the sidebands, the field at 1 GHz carrier is significantly different from the 50 MHz measurements. The field has spread further from the FPGA and oscillator to the SDRAM. The magnitude is also significantly stronger than the sidebands, at least 20 times larger. Based on the measurements, the algorithm identified one loop source and two electric dipoles. The loop source is between the left side of the SDRAM and the bottom of the CPLD. The first electric dipole is at the bottom left corner of the SDRAM, near the 3.3V power supply pins. The second electric dipole is located at the middle left side of the

FPGA near traces to the SDRAM and decoupling capacitors. This source is similar to one found for LDM/LDL1.

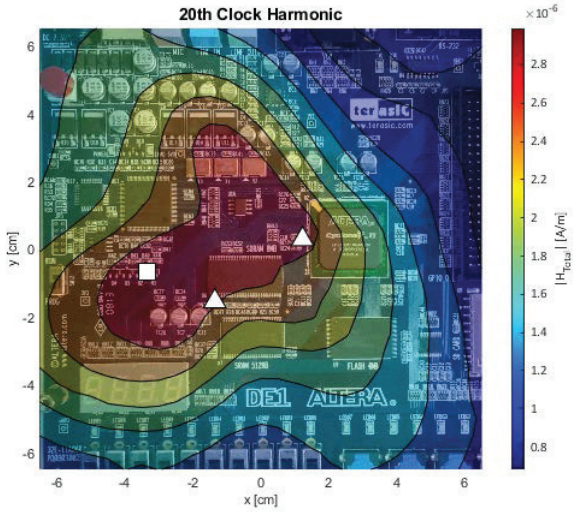

Fig. 4.13. Magnetic field measured at the 20th harmonic (1 GHz) of the FPGA clock.

4.6 A13-MICRO Experiments

As with the DE1, the fields generated by the A13-MICRO are recorded at the baseband and at the modulated clock frequencies while the board executes different instruction pairs. The A13-MICRO's ARM Cortex-A8 processor is integrated into an A13-Allwinner system-on-chip (SoC). The A13-MICRO has a Linux operating system provided by its manufacturer and remains active during the measurements. The processor clock is set to 1.008 GHz. The SDRAM controller is also integrated into the SoC. Unlike the FPGA, the clocks for the SoC/processor and memory operated at different frequencies. The SDRAM controller provides a 408 MHz clock to an external 2 Gb DDR3. A picture of the device with the important areas labeled (the SoC, SDRAM, and power supply circuitry) is provided in Fig. 4.8(b). For the experiments, measurements are taken at the baseband, around the 1.008 GHz processor clock, and the second harmonic of the memory clock (816 MHz).

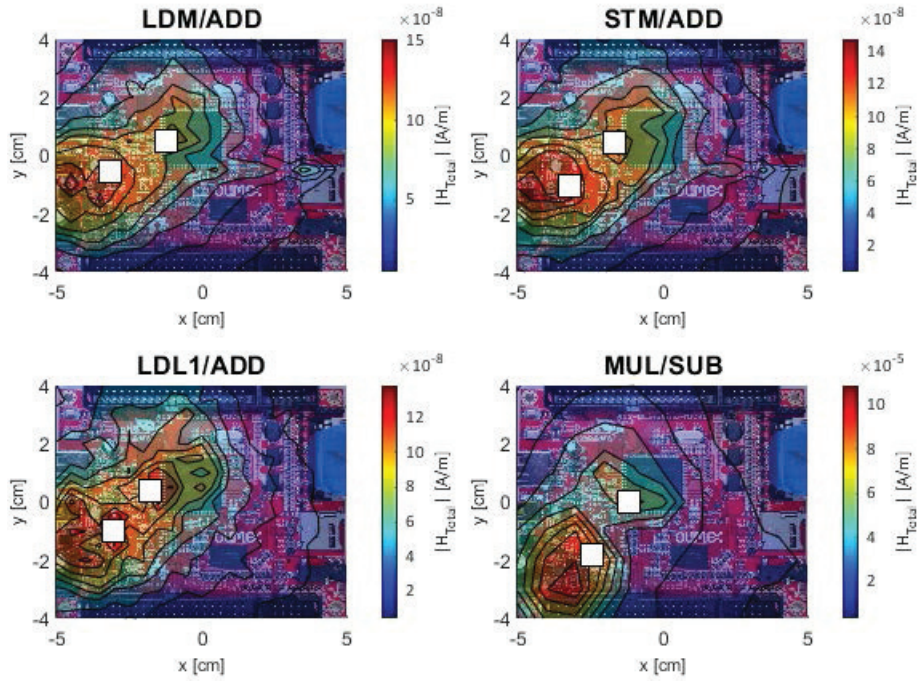

Fig. 4.14. Example of the baseband (156 kHz) A13-MICRO measurements for LDM/ADD, LDL1/ADD, STM/ADD, and MUL/SUB.

4.6.1 Baseband A13-MICRO Measurements

The baseband fields and sources for LDM/ADD, LDL1/ADD, STM/ADD, and MUL/SUB are shown in Fig. 4.14. Unlike the DE1, the baseband fields are similar. For example, the shapes of the load and store fields are nearly identical. Furthermore, while the field distribution for MUL/SUB does not match the others exactly, the fields for all instruction pairs are concentrated at similar locations. As a result, the algorithm determined that all the baseband measurements have two quasistatic loop sources, located at similar positions. The strongest source for each instruction pair is located near the bottom left corner of the board, near the power supply circuitry. Each measurement has a second weaker source located at the decoupling capacitors near the left side of the processor. The location of this second source shifts slightly based on the instruction pair. Based on the fields, the second source is significantly weaker than the first for the MUL/SUB pair.

Interestingly, even the fields generated by LDL1/ADD and LDM/ADD are similar, with

the primary difference being that LDL1/ADD has more noise due to the emanations being weaker. Despite the external memory being active during both LDM/ADD and STM/ADD, their fields are concentrated on the other side of the device from the external SDRAM, just as the on-chip instruction pairs (LDL1/ADD and MUL/SUB). This result indicates that the emanations from the external SDRAM and its traces are being overshadowed by the emanations from the power supply. While the traces connecting the SoC to the SDRAM are similar in size to the power supply traces, their emanations are much weaker, likely being suppressed due to the design of the PCB layout. The clock is provided to the SDRAM on a pair of differential traces. One trace provides the memory clock, while the second trace provides the return path. The traces are kept as close together as possible and the same length to ensure that the magnetic flux from the two traces is canceled out.

Overall, these measurements demonstrate that, at the baseband, the emanations are mostly coming from the power supply circuitry and the decoupling capacitors connected to the SoC. The emanations are again likely the result of variations in the current draw as the processor executes the instructions. Any emanations from inside the components are being overshadowed by the emanations from outside the components.

4.6.2 A13-MICRO Processor Measurements

Next, examples of the measurements taken at the upper sideband of the modulated $1.008\ \mathrm{GHz}$ processor clock are provided in Fig. 4.15. In the figure, the fields and sources for same instruction pairs as in the previous subsection are shown. The fields for each instruction pair are similar, with only slight variations in the shape. For each instruction set, the field is concentrated near the device's power supply circuitry. Based on the measurements, the algorithm determined that each instruction pair has a two electric dipole and one loop source. The physical locations of the three sources vary slightly between instruction pairs (less than $5\ \mathrm{mm}$); however, they are at the same general locations for all measurements.

The loop and one of the electric-dipole sources are located near a via for the SoC's

processor power supply. At this location the via connects a trace on the top layer of the board to a trace from an inner layer of the board. This inner trace is connected to the SoC. These sources are located near the strongest source determined for the baseband measurements. However, the contribution of other sources closer to the SoC decreased at the higher frequency.

The second electric-dipole source is located on the opposite side of the board, between the bottom right corner of the SoC and the top of the SDRAM. It is present even when the SDRAM is not active. It is also noticeably weaker than the other sources.

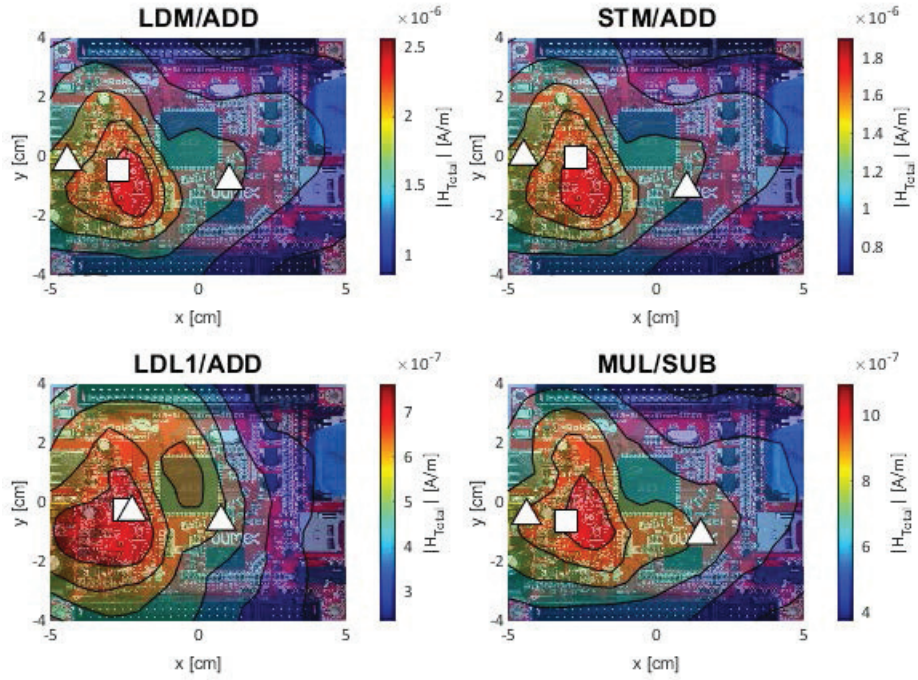

Fig. 4.15. Examples of the upper sideband (+20 kHz) of the 1.008 GHz processor clock for different instruction pairs.

The measured fields taken at the carrier frequency are shown in Fig. 4.16. As with the DE1, the field distribution of the carrier does not change with the instruction set. However, unlike the DE1, the field distribution and source for the carrier match the sidebands. This result indicates that the carrier has already been modulated by the time it reaches the strongest source.

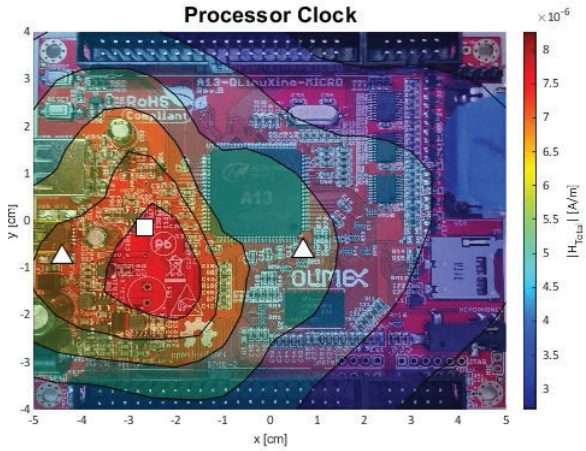

Fig. 4.16. Magnetic field measured at the 1.008 GHz processor clock.

4.6.3 A13-MICRO Memory Measurements

An overview of the modulated memory clock measurements is shown in Fig. 4.17. Since the external memory is active only during off-chip instructions, LDM/ADD, LDM/MUL, STM/ADD, and STM/MUL are shown in the figures. The fields generated by the different instruction pairs are nearly identical, with a slight difference related to whether load or store are used. This difference is likely caused by slightly different parts of the SoC being active when executing the instruction pairs. Based on the measurements, the algorithm determined that the instruction pairs share a single loop source and have no electric-dipole sources. This loop source is located inside the SoC, near the middle of its bottom edge. It is near the pins for the SoC's SDRAM data buses and the SDRAM output clock. The source location is unsurprising given that the instruction pairs involve the external memory and the SDRAM clock acts as the carrier for the emanations.

As with the baseband and processor measurements, the emanations from the external memory are being overshadowed. Importantly, unlike the previous measurements, the fields from inside the SoC are significantly stronger than the fields emanated from anything outside the component. Despite the frequency being only 816 MHz, the emanations from inside the chip are strong enough to overshadow the emanations from the rest of the device.

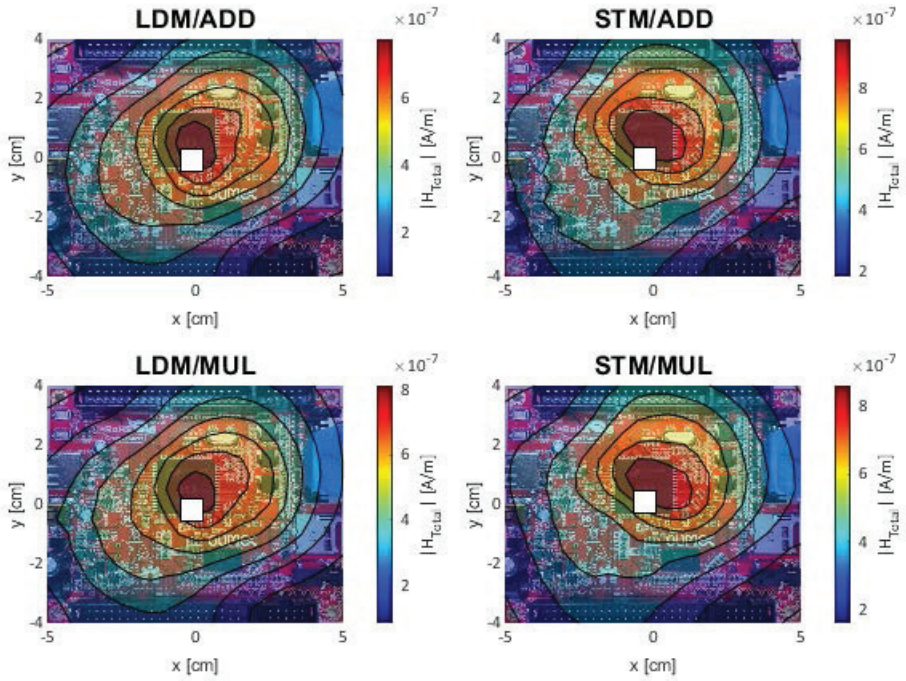

Fig. 4.17. Magnetic fields measured at the upper sideband (+20 kHz) of the 816 MHz memory clock for different instruction pairs.

The magnetic field measured at the memory clock frequency is shown in Fig. 4.18. As with the processor measurements, the fields at the clock frequency are nearly identical to the sidebands. This result indicates that the memory clock has already been modulated by the time it is emanated.

4.7 Conclusions

This chapter presents a new method for measuring and locating sources of high-frequency EM side channels on PCBs. This method is well suited for evaluating the leakage from embedded and IoT devices. Given their ubiquity and vulnerabilities, the hardware security of those devices is an important concern for designers and researchers.

This method is low cost and time efficient. It can measure the intensity of the EM side-channel leakage while limiting the impact of interference and reflections. Furthermore, unlike similar systems, it relates the sources to the basic instructions commonly used on the device. To be able to focus on these instructions, the user needs to be able to implement

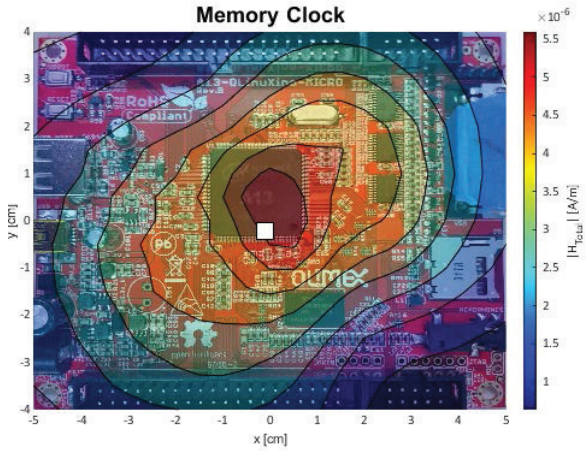

Fig. 4.18. Magnetic field measured at the 816 MHz memory clock.

the excitation program described in Section 4.2.1 on the device. However, given that this method is intended for researchers and designers, it is reasonable to assume the user would have this access.

The accuracy of the measurement setup was verified by comparing the measured with simulated results. As a demonstration of its capabilities, this system was used to identify the side-channel sources at 1 GHz of two devices. The results were compared to the measurements taken at the low frequencies to determine the effect the frequency has on the sources. For the DE1, the results demonstrate that the sources can vary significantly with the frequency, with the source spreading over the device as the frequency increases. For the A13-MICRO, the sources varied noticeably less with the frequency. Instead, they remained concentrated near the SoC and the power supply circuitry.

CHAPTER 5
LEVERAGING EM SIDE CHANNELS FOR RECOGNIZING COMPONENTS ON A MOTHERBOARD

5.1 Overview

As discussed in Chapter 1, with the increasingly complexity of the electronic supply chain, the ability for electronics designers to authenticate the ICs used in their designs is a pressing concern. To address this problem, we propose using EM SCA to recognize/authenticate components integrated onto a motherboard. By focusing on components on a PCB, designers can authenticate devices assembled by third parties. This method is intended for detecting types of counterfeiting where the physical design of the component differs from the expected component. Changes in the IC will impact the emanations that comprise the EM side channel. Examples include cases where the intended IC has been replaced with a reverse engineered copy or with a lower quality component and cases where the design has been tampered with [9].

The proposed method authenticates components based on the modulated signals emanated while the component is active. These signals are generated by exciting the component in a controlled manner, and the emanated spectrum is recorded. During training, the emanations from several trustworthy examples of each component are recorded. The spectrums recorded during the training are used as signatures for the components of interest. When testing an unknown component, its spectrum is recorded and then compared to the previously recorded training signatures. To improve the efficiency, the size of the spectrum is first reduced by projecting it into a vector space generated from the training signatures. The identity of the tested component is then determined using a k-NN algorithm [124].

More complicated and expensive ICs, such as memories and processors, are the focus,

since they are some of the most important and expensive components the device. Counterfeits of passive components on the motherboard, such as capacitors or inductors, cannot be detected using this method; however, they are less critical to the device. Unlike other authentication methods, our method requires no additional hardware, nor does it damage the component during testing, hence reducing its cost. At the same time, the equipment for detecting counterfeits is always under the control of the user, unlike for SHIELD, where the dielet itself could be tampered with.

The proposed authentication method has successfully classified external memory, processor, and Ethernet transceiver components integrated on seven types of IoT devices. Nine to ten different instances of each device are used in the experiments. Cross-type testing of boards is conducted as well. Since manufacturers commonly use the same components in multiple designs, being able to collect the training signatures on one motherboard and test components from different motherboards significantly speeds up the process and decreases the cost. The method achieves a classification accuracy greater than 96% for all tested components. These results demonstrate that this method can recognize components based on their EM emanations even if they are integrated on different a type of motherboard.

The rest of this chapter is organized as follows. Section 5.2 describes the approach used to excite components and maximize the presence of EM side-channel features. Section 5.3 describes how the measurements are processed to improve the efficiency of the identification. Section 5.4 details the process for training on and testing components. Section 5.5 describes the experiments performed to evaluate the proposed method. Finally, Section 5.6 presents the conclusions.

5.2 Signals Carrying Information About Device Side-Channel Signatures

This section discusses what spectral features are found to be relevant for component recognition and how we excite electronic components to maximize the presence of the EM side-channel features.

5.2.1 Spectral Features

The proposed method relies specifically on the modulated signals emanated by the component for identification. As mentioned in the previous chapters, these are some of the strongest signals available in the side channel [13]. They are caused by the device's program activity unintentionally modulating periodic synchronizer signals, such as the clock, already present in a device [13]. Program activity results in the superposition of a time-varying current on the traces inside and connecting the components used for the execution. The magnitude of the emanations depends on the change in power when executing the activity, while modulation frequency is related to the time it takes to execute the repetitive behavior. While multiple types of modulation can occur in a device, this work relies only on amplitude modulated (AM) signals generated by a component for identification [14].

Any change in the component or program activity affects the properties of the emitted AM signal. For example, the shape and spread of the sidebands in the frequency spectrum are related to the time it takes to execute parts of the program activity. If the execution time varies, the sideband's shape will contort and spread in frequency. By keeping the program activity consistent for all tests, the spectral features of the emanations can be used as a signature for identifying the components. Experimentally we have determined that these features include: 1) the overall modulation frequency, 2) the sideband shape and spread, 3) the relative strength of the sideband's fundamental frequency and the higher harmonics, 4) the carrier frequency, and 5) the carrier spread. Fig. 5.1 in the following subsection provides an example of a signal with these properties highlighted.

5.2.2 Device Excitation

When selecting spectral features to use for identification, the device needs to be in a known and repeatable state. In this state, the component of interest needs to be active; otherwise, there will be nothing in the spectrum to use for identification. For example, leaving the device in standby is an obvious option for a measurement state. However, the device will

not be very active, making it difficult to find useful features for identification (assuming the component is active at all).

Even if the component is being excited, it can still be difficult to locate useful spectral features. Depending on the program activity, the spectrum can change rapidly in time. As a result, locating useful spectral features, in both time and frequency, can be challenging. Selecting a measurement state where the device's spectrum is relatively constant makes identification easier.

To address these challenges, the excitation program described in Section 3.2 is used to excite the device during testing. As demonstrated in Chapters 3 and 4, different instructions will excite different components (such as the processor and RAM), depending on what is being used for execution. This effect limits the influence other components have during the measurements.

The excitation program can be easily implemented on a variety of devices, making it ideal for authenticating components integrated on multiple types of motherboards. It has already been implemented on several different types of laptops, desktops, cell phones, computer boards, and FPGAs [12], [57], [125], [126]. In these experiments, the devices' OS included several versions of Windows, multiple different implementations of Linux, and the embedded operating systems unique to the FPGA.

The duty cycle of the excitation is based on the amount of time spent executing each instruction. For all of the experiments in Section 5.5, the duty cycle is set to 50%; therefore, the modulating baseband signal should be a square wave. While a square wave should have only odd harmonics, there will likely be weak even harmonics present in the generated spectrum. Since it takes a fixed amount of time to execute each instruction, it is difficult to tune the waveform to have a precise duty cycle. This distortion can be used as a factor in identifying the component since the execution time of each instruction is dependent on the components used for the execution.

An example of the spectrum (in decibels) generated using the excitation program is

shown in Fig. 5.1. To normalize the signal, it has been divided by its mean. In the figure, the carrier and the modulated sidebands caused by the excitation program are labeled. The excitation has a fundamental frequency of 10 kHz with higher harmonics at ± 20 and ± 30 kHz from the carrier. Since the duty cycle is set to 50%, the odd harmonics are much stronger than the even. However, the presence of the weak even harmonics at ± 20 kHz indicates the actual duty cycle is not exactly 50%. The shape and spread of the sidebands are caused by variations in the execution time of the excitation program, while the spread of the carrier is caused by the instability of its source. Finally, the power of the sidebands relative to each other and to the carrier are based on three factors: the physical properties of the component, the excitation program, and the location of the measurement probe. *Keeping the excitation program and probe position fixed allows us to observe the physical properties of the components.*

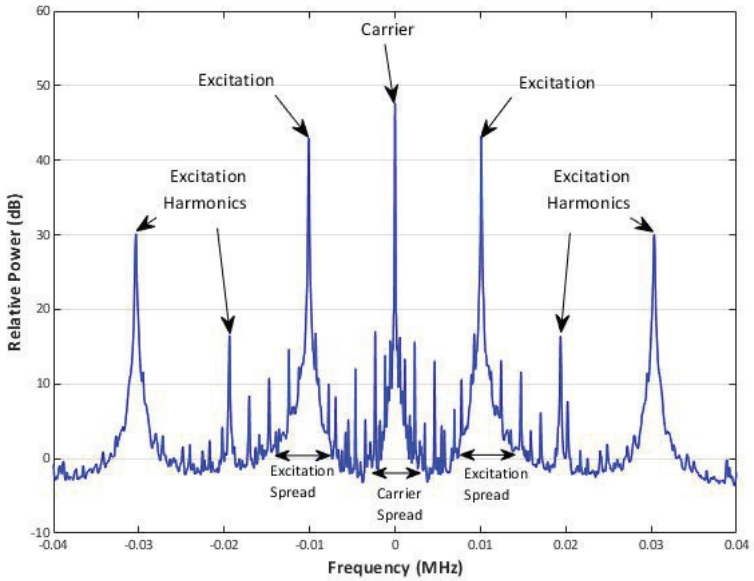

Fig. 5.1. Example of the spectrum produced by the excitation program.

5.3 Signal Compression and Processing

After the component has been excited and its emanations have been recorded, the measured signal is processed to make it easier to locate important features for identification. To eliminate the influence of time variations and noise on the measurements and to improve the efficiency of identification, an updated version of the approach described in [127] is used. This new approach is described below.

During a test, the EM emanations from the excited component are recorded for a period of time, T. The number of samples, M, is equal to the measurement time multiplied by the sample frequency, f_{s}. To better emphasize the elements of the signal used for identification, the signal is converted to the frequency domain. However, instead of converting the entire signal at once, the measurement is first broken into several segments, N_{R}. A short-time discrete Fourier transform (STFT) is then applied to each segment. A flattop window is used to improve the accuracy of the relative amplitudes of different frequency components.

Each STFT operation for the $k\mathrm{th}$ frequency component is calculated by

$$Y_h[k] = \sum_{n=1}^{N} y[n + (h-1)N_{\mathrm{s}}]w[n-k]\exp(-\mathrm{i}2\pi kn/N) \tag{5.1}$$

where y is the original measurement, h is the STFT operation number, w is the flattop window function, N_{s} is the number of nonoverlapping samples, and N is the window size for the STFT operations. The nonoverlap time is the number of samples, N_{s}, shifted between STFT segments.

The number of elements in each segment, N, affects the resolution of the window. The resolution needs to be high enough for features of the spectrum to be distinguishable. However, if the size of N is too large, it will increase the measurement and processing times. After M, N, and N_{s} have been defined, the number of STFT operations performed for the data can be calculated by

$$N_{\mathrm{R}} = \mathrm{floor}\left(\frac{M-N}{N_{\mathrm{s}}} - 1\right). \tag{5.2}$$

To reduce the impact of noise and other time variations in the signal on the spectrum, these N_{R} STFT operations are averaged together as

$$\overline{\mathrm{Y}}[k] = \frac{1}{N_{\mathrm{R}}} \sum_{h=1}^{N_{\mathrm{R}}} |Y_h[k]| \tag{5.3}$$

where $\overline{\mathrm{Y}}$ is a row vector containing the averaged frequency magnitudes. Here, to eliminate the impact of the starting time on the data, the phase is removed by taking the magnitudes of the STFT operations. Afterwards, the strongest component in the spectrum is downconverted to 0 Hz. This component is assumed to be the carrier. Since the exact frequency of the carrier is influenced by factors such as manufacturing variability and temperature, the carrier frequencies of two identical components on separate devices will be slightly different. Shifting the carrier to the center of the spectrum ensures that this difference does not influence the identification. Afterwards, a band-pass filter is applied to the downconverted signal by removing 10% of the bandwidth. This filtering ensures that the carrier is located at the center of the spectrum.

Next the measurements are converted from a linear scale to a decibels scale (dB) using

$$\overline{\mathrm{Y}}_{h_{\mathrm{dB}}}[k] = 20\log_{10}(\overline{\mathrm{Y}}_h[k]). \tag{5.4}$$

This conversion reduces the influence of the strong signal components, while increasing the influence of weaker components in the spectrum. Since the modulation is not intentional, the carrier tends to be significantly stronger than the sideband components, usually tens to hundreds of times stronger. Without this conversion, the carrier has a disproportionate influence on the data. Afterwards, the measurements are standardized by subtracting their mean and dividing by their standard deviation.

Once all the measurements are processed, they are combined in the HxN matrix

$$\mathbf{Y} = \begin{bmatrix} \overline{\mathrm{Y}}_{1_{\mathrm{dB}}} \\ \overline{\mathrm{Y}}_{2_{\mathrm{dB}}} \\ . \\ . \\ . \\ \overline{\mathrm{Y}}_{H_{\mathrm{dB}}} \end{bmatrix}, \tag{5.5}$$

where H is the number of measurements. In other words, each row represents the averaged frequency components of a measurement. However, the matrix generally contains redundant information because H and N are not equal. To reduce the size of the data and improve efficiency, SVD (singular value decomposition) is applied to the data matrix. As a result, the data matrix is decomposed into the following form:

$$\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \tag{5.6}$$

where $\mathbf{U}$ is the left singular vectors matrix, $\mathbf{\Sigma}$ is the singular values matrix, and $\mathbf{V}^T$ is the transpose of the right singular vectors matrix. The measurement data is then projected into a new vector space, $\mathbf{Z} \in \Re^{H \times K}$, by

$$\mathbf{Z} = \mathbf{Y}\mathbf{V}_K, \tag{5.7}$$

where $\mathbf{V}_K$ has been reduced to $K \times K$ matrix. The size of K is significantly smaller than the length of the original measurements and corresponds to the K largest singular values in $\mathbf{\Sigma}$. These K vectors represent the signal components of the data matrix that contains the directions corresponding to largest K singular values. As a result, the original $H \times N$ dimensional feature space of the measurements has been reduced to $H \times K$ dimensions, without much loss in information about the data.

5.4 Training and Testing Process

The identification process can be broken into training and testing phases. In the training phase, EM signatures from example components are recorded and processed. Here, the identities of the components are already known. If the training components are used on multiple types of devices, only one of the devices needs to be used for training. Each training component is excited using the excitation program, and its emanations are recorded for T seconds. The modulation frequency used for the excitation program is set beforehand and kept constant for all measurements. The carrier frequency of the component is located using either a method such as the ones detailed in [13] and [14] or manually. The measurements are processed into averaged STFTs and stored in a matrix. If multiple types of components are tested, such as memory and processor, the signatures can be divided into multiple matrices based on the component's function or the carrier frequency. The training matrix or matrices are then decomposed using SVD. The first K columns from the resulting matrix $\mathbf{V}$ are selected. Next, the matrix $\mathbf{V}_K$ is used to project the training data into the new vector space. In other words, assuming the projected training data set is $\mathcal{Y}$, we generate a model ($\mathcal{Y}$, $\mathbf{V}$) that is used in the testing phase.

In the testing phase, the EM emanations generated by one or more unknown components are recorded. As in the training phase, the tested component is excited using the excitation program. The measurements are then processed into averaged STFTs. The resulting signatures are projected into a new vector space using the matrix created from the training data. Afterwards a k-NN algorithm is applied to the projected training and testing data by using the model ($\mathcal{Y}$, $\mathbf{V}$). The k-NN algorithm determines the identity of the tested component based on the standardized Euclidean distance between the projected measurements and the training measurements.

Although k-NN is a fairly simple type of clustering method, it is a practical tool for classifying components. As the experiments in the following section demonstrate, it can

accurately determine the identities of several types of components.

5.5 Experimental Validation

To demonstrate its effectiveness, the new identification method is applied on components from seven types of IoT devices from Olimex and two FPGA developement boards. This method is not limited to such devices; we use them here only as examples. All the IoT devices are from the same manufacturer (Olimex) because it presents the most difficult case for identifying components. As demonstrated in Subsection 5.5.3, contrasts between devices from different manufacturers are even larger, making them easier to identify.

For the experiments, the external memory, processor, and Ethernet transceiver components from each device are tested. In the following subsections, the measurement setup, the tested components, and the experimental results are discussed. The experimental results are broken into four subsections. Subsection 5.5.3 demonstrates the effect of projecting the measurement data in the new feature space, while the other three subsections (5.5.4 to 5.5.6) describe the results of testing on different types of components. The measurement settings described in the measurement setup applies to these three subsections. Subsection 5.5.3 uses a simpler combination of settings.

5.5.1 Measurement Setup

Fig. 5.2 shows the measurement setup used for the experiments. A small hand-made circular coil probe with a $1\ \mathrm{mm}$ radius is used to measure the magnetic field emanated by the component being tested. This probe was selected so that its small size would limit the influence of emanations from other nearby components on measurements. The probe is connected to a Keysight M9391A PXIe Vector Signal Analyzer (VSA) for recording the signal. During the measurements, the probe is placed directly on top of the monitored component, where the emanations are the strongest. To ensure consistency between measurements, the probe is positioned using the EM Probe Station Motorized XYZ Table from

Riscure [128]. The table is controlled through a USB port using a laptop. The laptop is also used to control the DUT and the VSA through their Ethernet ports. In cases where the DUT does not have an Ethernet port, a USB-to-Ethernet adapter is used. A diagram of the test setup is shown in Fig. 5.3.

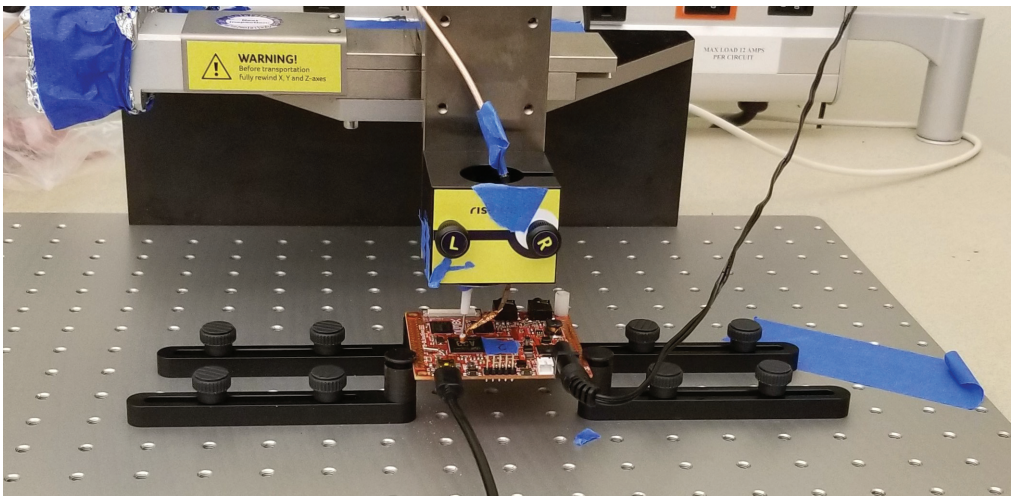

Fig. 5.2. Measurement setup used for the experiments.

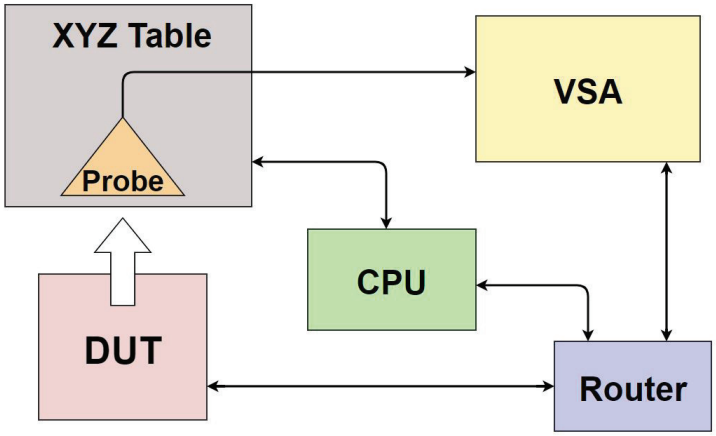

Fig. 5.3. Diagram of the measurement setup.

During a measurement, all the components are excited using the benchmark outlined in Subsection 5.2.2. The program generates a 10 kHz excitation signal with a 50% duty cycle by executing an alternating pattern of addition and load instructions. When measuring the external memory components, the array size of the load instruction is set to be much larger

($8.4\ \mathrm{MB}$) than the processor's cache to ensure that the external memory is active during the measurements. When exciting the processors, the array size of the load instructions is kept small ($8.2\ \mathrm{kB}$) to ensure that the load execution is mostly confined to the processor's cache. This small array size minimizes the influence the external memory has on the signal. When testing multicore processors, the excitation program is executed only on the first core, while the rest are left idle.

While a component is excited, its emanations are recorded for $1\ \mathrm{s}$ (which is T in Section 5.3). In Subsections 5.5.4 to 5.5.6, each of the seven types of IoT devices had 10 individual units; however, one unit (an A33-MAIN) is removed from the results since it is damaged. Therefore, a total of 69 boards are tested. Furthermore, cross-type testing is conducted because some of the device types have the same components. In these situations, measurements from only one type of devices are used for training. Similarly, in situations where the same component is used multiple times on the same device, measurements from only one instance of the component are used for training.

To account for the limited number of boards, a k-fold cross-validation scheme with five-folds is run for $10\,000$ iterations. Such schemes are commonly used in cases where the sample size is small since their results have a relatively low bias and variance compared to other cross-validation approaches [129]. The accuracy of correctly identifying each type of component and the overall accuracy of correctly identifying all the components are calculated by averaging the results for each iteration.

Several of the component measurements have interrupts caused by their device's operating systems [130]. To improve consistency and ensure the spectrum is primarily a result of the excitation program, these interrupts are removed from the measurements. The measurement bandwidth is $220\ \mathrm{kHz}$ (reduced to $200\ \mathrm{kHz}$ after processing). A small bandwidth is used to reduce the amount of interference present in the measurement. This interference may help distinguish different components from each other but can also make the same components on different device types appear dissimilar.

During processing, each measurement is processed into 43, N_{R}, STFTs with a length, N, of $20\,480$ before being averaged. The number of non-overlapping samples, N_{s}, is $4\,000$, which corresponds to $20\ \mathrm{ms}$. After the measurements are processed, the training measurements are used to generate a new vector space for evaluation. Afterwards, the k-NN algorithm is applied to the testing data by using the model ($\mathcal{Y}$, $\mathbf{V}$). For the following sections, only the first four dimensions of the new vector space are used for evaluation, i.e. $K = 4$. As mentioned previously, the k-NN algorithm classifies a measurement based on the standardized Euclidean distance between the measurement and each training signature. The differences between the four coordinates of the test and training points are scaled by dividing by 1, 1, 2, and 3, respectively, for the memory and processor and by 1, 1, 3, and 3 for the Ethernet transceivers. The number of dimensions and the scaling factors can be tuned to improve the classification accuracy for a specific set of components.

5.5.2 Test Devices

In the experiments, seven types of IoT devices from Olimex are tested. These devices are the A10-OLinuXino-LIME [131], A13-OLinuXino [122], A13-OLinuXino-MICRO [132], A20-OLinuXino-LIME [133], A20-OLinuXino-LIME2 [134], A20-OLinuXino-MICRO [135], and A33-OLinuXino [136]. For simplicity, these devices will be referred to as A10-LIME, A13-MAIN, A13-MICRO, A20-LIME1, A20-LIME2, A20-MICRO, and A33-MAIN for the rest of this work. All the devices run Linux OS provided by Olimex. Instead of being installed on the device itself, the OS are saved to SD cards. The only change made to the devices is installing the excitation program.

Pictures of tested IoT devices are shown in Fig. 5.4. All the IoT devices are part of Olimex's OLinuXino open source hardware product line. They are convenient options for these experiments since Olimex has provided detailed information about each device (such as the schematic, parts list, and PCB layout) on their website [137]. In Subsection 5.5.3, two extra memory components (MEM5 and MEM6) are tested to demonstrate the impact

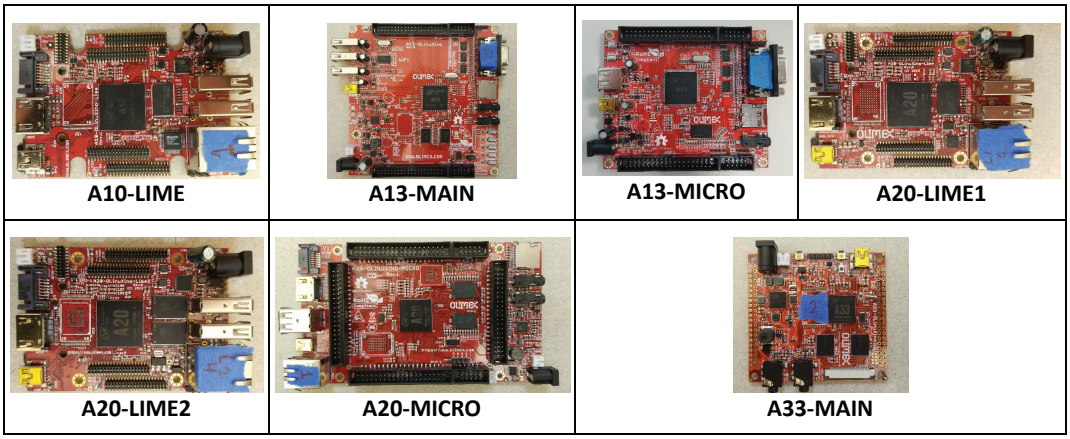

Fig. 5.4. The seven IoT devices (not to scale).

of projecting the measurements into the new feature space. These extra components are integrated into DE0-CV Cyclone V [138] and DE1 Cyclone II development boards [139]. Pictures of the devices are shown in Fig. 5.5. These components and their motherboards are significantly different from the Olimex devices, both in functionality and physical properties. The differences result in the projected data from the development boards being easily distinguishable from the IoT measurements. Therefore, these components are not included in later subsections. Correctly classifying similar, yet physically different components is more challenging, especially if the components are integrated onto similar motherboards from the same manufacturers.

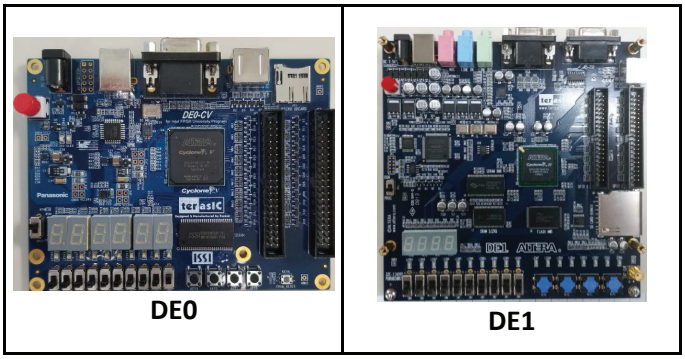

Fig. 5.5. The DE0-CV Cyclone V and DE1 Cyclone II development boards (not to scale).

The reason motherboards from the same manufacturer increase the difficulty is that manufacturers commonly use the same components and PCB layouts in multiple designs to save time and money. These similarities can influence the parts of the spectrum not related to the component of interest. The A10-LIME and A20-LIME1 are examples of reused PCB layouts. The PCB for the A10-LIME is an older revision of the PCB for the A20-LIME1. In this situation, the traces on both motherboards will have similar emanation properties since they are almost identical in shape and composition. If one motherboard is made by a different manufacturer, the emanation properties change since the physical configurations of the traces and the material properties of the PCB would change. However, in either case, the signal properties of the emanations from the traces and the components still depend on the component and the program activity.

The following experiments focus on identifying the external memory, processor, and Ethernet transceiver components present on each device. Some of these components are not present on all the devices, while, in other cases, some devices use the same components as others. A complete list of the components, the devices they are present on, and the measurement frequency are provided in Table 5.1. For simplicity, the components will be referred to by their label in the table. More information about the devices and their components can be found on the component manufacturer's websites (referenced in the table).

5.5.3 Measurement Projection

Before discussing the method's performance, the impact of projecting the measurement data into the new feature space needs to be explored. As an example, three measurements from all six types of memory components are projected into a new 3-D feature space (shown in Fig. 5.6). All measurements are recorded for 0.2 s with a bandwidth of 1 MHz and an excitation frequency of 100 kHz. For simplicity, the example measurements are used for generating the feature space before being projected into it.

TABLE 5.1
List of Tested Components

Label	IC Name	IC Type	Devices	Carrier Frequency (MHz)	Source
MEM1	K4B4G1646Q-HYK0	4Gb DDR3 SDRAM	A10-LIME	384	[140]
MEM2	H5TQ2G83FFR	2Gb DDR3 SDRAM	A13-MAIN U1 and U2	408	[141]
MEM3	H5TQ2G63BFR	2Gb DDR3 SDRAM	A13-MICRO	408	[142]
MEM4	MT41K256M16HA-125:E	4Gb DDR3 SDRAM	A20-LIME1, A20-LIME2 U2 and U3, A20-MICRO U2 and U3	384	[143]
MEM5	IS42S16320D-7TL	512Mb SDR SDRAM	DE0	100	[144]
MEM6	A2V64S40CTP-G7	64Mb SDR SDRAM	DE1	50	[145]
PROC1A	Allwinner A10 (Cortex-A8)	SoC (Processor)	A10-LIME	1008	[146]
PROC1B	Allwinner A13 (Cortex-A8)	SoC (Processor)	A13-MAIN, A13-MICRO	1008	[147]
PROC2A	Allwinner A20 (Cortex-A7)	SoC (Processor)	A20-LIME1, A20-LIME2, A20-MICRO	960	[148]
PROC2B	Allwinner A33 (Cortex-A7)	SoC (Processor)	A33-MAIN	960	[149]
ETH1	RTL8201CP	Ethernet Transceiver	A10-LIME	25	[150]
ETH2	LAN8710A-EZC-TR	Ethernet Transceiver	A20-LIME1, A20-MICRO	25	[151]
ETH3	KSZ9031RNXCC-TR	Ethernet Transceiver	A20-LIME2	25	[152]

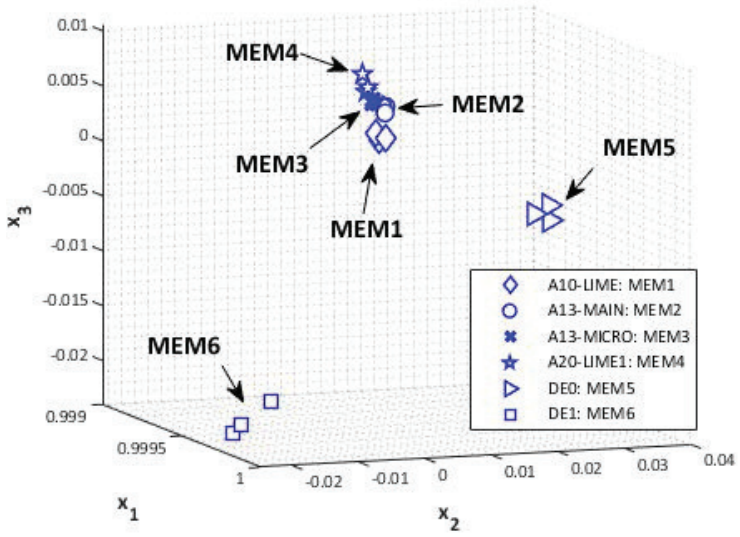

Fig. 5.6. Example of the memory measurements projected into the new feature space.

In Fig. 5.6, the points from MEM5 and MEM6 are relatively far away from the other components. Their isolation is the result of their measurements having significantly different spectral features. When the measurements are projected, the differences in the spectral features translate into different coordinates in the feature space. These differences are the result of the first four memory components being radically different from the last two. While all six are SDRAM, the first four are DDR3 with operating frequencies greater than 300 MHz, while the last two are SDR (single data rate) with operating frequency below 200 MHz.

On the other hand, points from MEM1 through MEM4 are clustered close enough together that they are difficult to visually distinguish from one another. The reason for this clustering is that the components are similar in functionality to one another. Furthermore, the fact the motherboards are from the same manufacturer likely helped their similarities. However, the closer the projected points from different types of components are to one another, the greater the risk of them being misclassified.

More detail about the relationships between the measurements can be gained by comparing the separation distances between the projected data. The average distances between data from two different types of components and the average distance between data from

the same type of component can be calculated. For simplicity, the distance between points from two different types of components is called the *class-distance*, and the distance between points belonging to the same type of component is called the *self-distance*. The class-distance is the result of the measurements from the different components having different spectral features. The larger the distance, the greater the difference. The self-distance is caused by differences in the spectrums measured from two of the same type. These differences can be the result of changes in how the measurements are taken (such as probe type and probe position), changes in the program execution over time, and manufacturing variation. By the design, the test process described in Section 5.4 minimizes the first two factors. However, the influence of manufacturing variation cannot be removed.

These manufacturing variations are caused by random fluctuation during the manufacturing process of an IC. These fluctuations result in small physical distinctions between individual ICs of the same type and are the basis of physically unclonable functions (PUFs) [153]. These small distinctions can impact the EM emanations generated by the IC, causing slight disparities in the measurements taken on individuals of the same type. While the impact of manufacturing variation cannot be removed, the likelihood of the manufacturing variation causing a misclassification can be evaluated by comparing the class-distances and self-distances of the data. If the self-distance for a component is much smaller than the component's class-distances, it can be concluded that the effect of the manufacturing variation is outweighed by the dissimilarities in the spectral features between the types of components. Therefore, the chance of manufacturing variation causing misclassification is small.

The average distances between the measurements for each type of memory are shown below in Table 5.2. The diagonal values (in bold) are the self-distances for each component, while the rest are the average class-distances between different types of components. For readability, the values here and in later sections have been multiplied 100. The values themselves are unitless and their only significance is their size relative to one another.

TABLE 5.2
Average Distances Between Select Memory Components

	MEM1	**MEM2**	**MEM3**	**MEM4**	**MEM5**	**MEM6**
MEM1	**1.8**	5.6	7.3	9.7	80.4	59.9
MEM2	5.6	**1.1**	3.0	4.8	82.1	64.9
MEM3	7.3	3.0	**1.2**	2.7	85.1	65.0
MEM4	9.7	4.8	2.7	**1.9**	86.1	67.4
MEM5	80.4	82.1	85.1	86.1	**15.4**	107.8
MEM6	59.9	64.9	65.0	67.4	107.8	**6.1**

Reflecting the results in Fig. 5.6, the class-distances between the first four memories and MEM5/MEM6 are significantly higher than the class-distances between the first four memory components only. For example, the class-distance between MEM1 and MEM4 is 9.7 while the class-distance between MEM1 and MEM6 is 59.9.

Furthermore, the table demonstrates that the class-distances between each component are larger than self-distances. This indicates that the impact of the manufacturing variation is outweighed by the differences between different types of components. Therefore, the risk of the manufacturing variation between the tested components causing misclassification is small.

5.5.4 Recognition of Memory Components

Next the devices with the first four memory components are evaluated. Only one MEM1, MEM3, and MEM4 component is integrated on the A10-MAIN and A13-MICRO, and A20-LIME1, respectively. However, A13-MAIN has two MEM2s, and the A20-LIME2 and A20-MICRO have two MEM4s. For devices with more than one of the same type of memory component, each component is measured and evaluated separately. The external memory from the A33-MAIN is not included since it uses a spread spectrum memory clock.

Fig. 5.7 shows a comparison between the spectrums measured from MEM1 on a A10-LIME, MEM2 on a A13-MICRO, MEM3 on a A13-MAIN, and MEM4 on a A20-LIME1. For example, MEM1's carrier has a much stronger and wider spread than the other three components. Furthermore, the relative strength of the odd sidebands for MEM1 is lower

than the other memory components, while the even harmonics are much stronger.

At the same time, the signatures for MEM2 and MEM3 are similar, an unsurprising result given that the components are from the same product line. (The part numbers for MEM2 and MEM3 are H5TQ2G83FFR and H5TQ2G63BFR respectively). However, noticeable differences can be identified between the two components. For example, the harmonics of the MEM3 are more spread out compared to those of MEM2.

The sidebands for MEM4 are the strongest among the memory. Furthermore, the spread of the sidebands for MEM4 is much larger than the spread of the other components. This spread is especially noticeable at the first harmonics, where it is nearly $10\ \mathrm{kHz}$.

Example measurements from A10-LIME, A13-MAIN, A13-MICRO, and A20-LIME1 are used for training. These devices are selected because they have the cleanest spectrums. Since there are two instances of MEM2 on A13-MAIN, only the components designated U1 on the motherboard are used for training. While all the A20 devices use the same memory components, only A20-LIME1 are used for training. The overall classification accuracy for the memory components after cross-validation is 100%. Since there were no classification errors, a confusion matrix for the results is not provided. The algorithm had no difficulty correctly classifying all the memory components, even MEM2 and MEM3, since the distinguishing spectral features are prominent and unique for each memory component.

Furthermore, the average distances for the memory components are shown in Table 5.3. The distances are calculated for each iteration of the cross-validation process before being averaged. As the table demonstrates, the class-distances for all four memory components are significantly larger than the self-distances. Therefore, there is little risk of manufacturing variations causing misclassification for the memories.

As a side note, the memory measurements provide an opportunity to examine the impact other factors have on the measured signal. For example, measurements taken on the two MEM4 components on a A20-LIME2 are shown in Fig. 5.8. The only difference between the two components is their physical location on the device; however, the spectrums

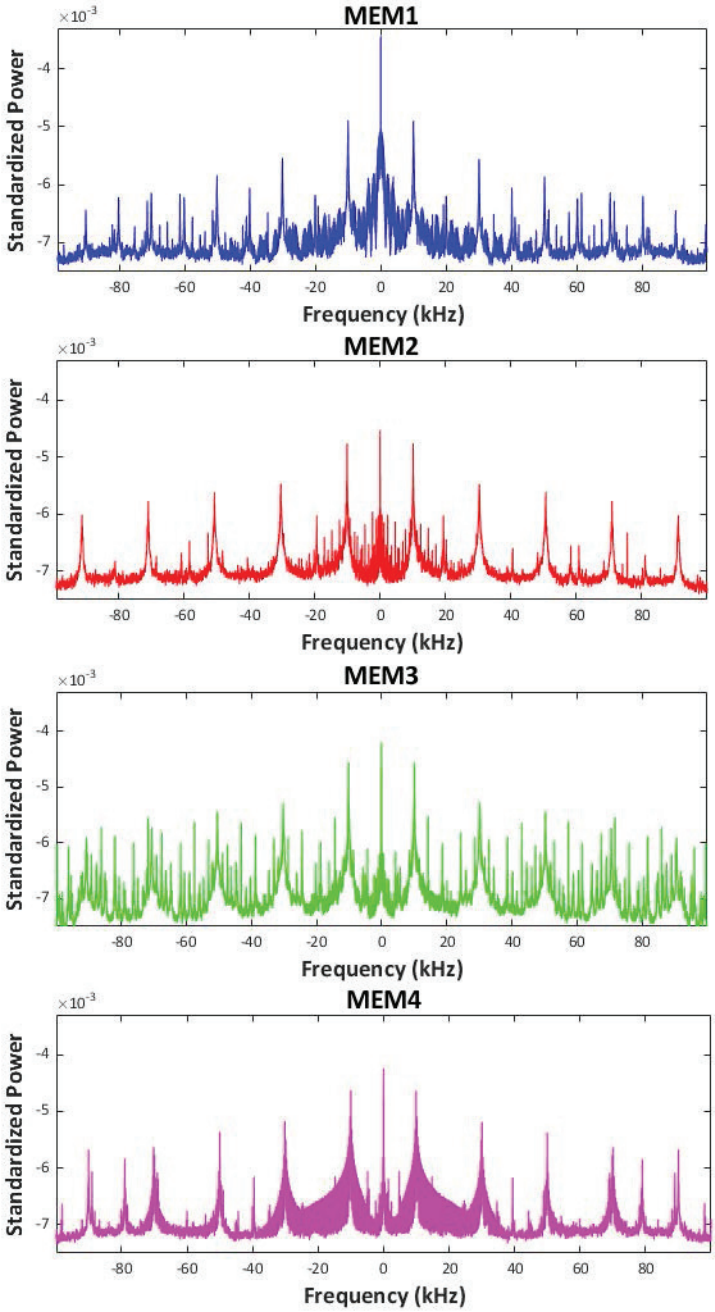

Fig. 5.7. Comparison of the EM signatures from MEM1 (top in blue), MEM2 (second from the top in red), MEM3 (third from the top in green), and MEM4 (bottom in magenta).

differ noticeably. Visually, it can be difficult to determine whether the spectrums belong to the same family of component. The sidebands generated by the excitation program are much weaker relative to the carrier at U3 compared to U2. Furthermore, the other activity

TABLE 5.3
Average Distances Between Memory Components

	MEM1	**MEM2**	**MEM3**	**MEM3**
MEM1	**1.0**	34.7	50.6	32.1
MEM2	34.7	**1.9**	19.6	13.6
MEM3	50.6	19.6	**1.0**	27.8
MEM4	32.1	13.6	27.8	**4.2**

modulating the carrier is stronger at U3. These differences are likely due to differences in PCB traces connected to the components and the relative position of the measurement probe. Despite their differences, both spectrums are correctly identified as being from MEM4. This identification is possible because most of the differences in the two spectrums are minimized after the measurements are projected into the new feature space and the dimensions are reduced to $K = 4$. Since the strongest variation in the data is represented by the first four singular vectors, smaller variations between measurements are lost when decreasing the dimensions.

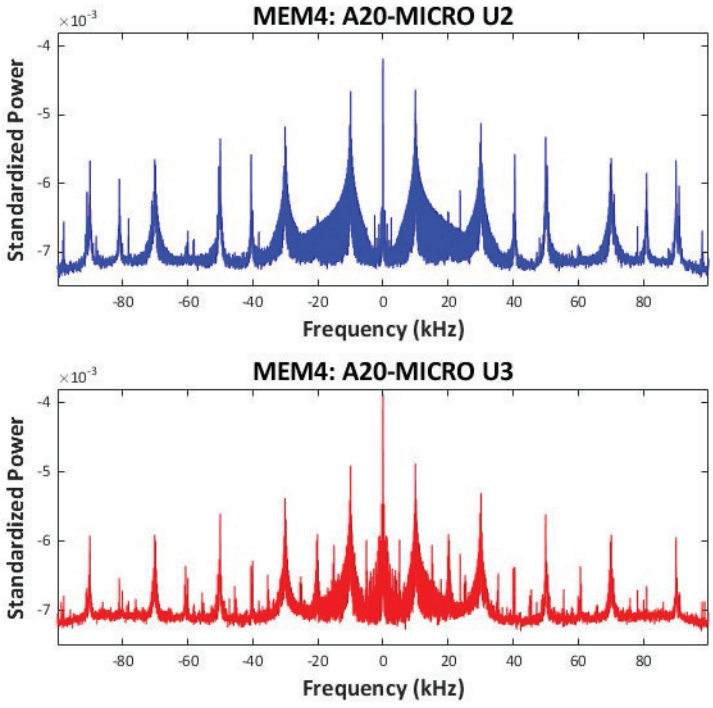

Fig. 5.8. Comparison of the MEM4 EM signatures from A20-MICRO U2 (top in blue) and A20-MICRO U3 (bottom in red).

5.5.5 Recognition of Processor Components

All the IoT devices have their processors integrated into a SoC from Allwinner Technology. The specific SoC is identified by the beginning of the device name (Allwinner A10, Allwinner A13, Allwinner A20, and Allwinner A33). Allwinner licensed designs for the processors from ARM Holdings and integrated the design in their SoCs. The A10 and A13 SoCs have a single Cortex-A8 CPU-core [154]. The A20 SoCs have a dual Cortex-A7 CPU-core [155]. The A33 Allwinner has a quad Cortex-A7 CPU-core.

While A10 and A13 have the same type of processor and A20 and A33 have the same type of processor, they still need to be classified as separate components. In both cases, the processors are based on the same design from ARM; however, the processors are not supplied as discrete components directly from ARM. Instead, Allwinner has to implement the design on each SoC. While the functionality may be the same, there will be slight differences in the processor layout based on the other electronics integrated into the SoC and the layout choices of the designer. From the prospective of classification, the differences between how a processor is implemented on difference types of SoC is similar to reversed engineered or tampered components. Therefore, each type of SoC needs to be classified as a unique group. Since the main focus of this section is identifying the processors, the Allwinner A10 is classified as PROC1A, the Allwinner A13 as PROC1B, Allwinner A20 as PROC2A, and Allwinner A33 as PROC2B.

Fig. 5.9 shows an example of the spectrums measured from an example of each type of processor while they are being excited. The top spectrum is an example of PROC1A from a A10-LIME, the second is an example of PROC1B from an A13-MAIN, the third is example of PROC2A from a A20-LIME1, and the bottom is an example of PROC2B from A33-MAIN. The measurements have been standardized.

All four spectrums have a strong carrier, with harmonics caused by the excitation program at 20 kHz intervals (the even harmonics are too weak to see), giving them the same general shape. Furthermore, the similarities are strongest between SoCs that share the

same processor design (i.e. PROC1A and PROC1B are very similar and PROC2A and PROC2B are very similar). However, there are slight differences in properties discussed in Section 5.2.1. For example, the sidebands and carrier for PROC1A are stronger and have a larger spread than the others. This and other differences are magnified after projecting the measurements into the new feature space generated from the training data.

During classification, measurements from A10-LIME, A13-MAIN, A20-LIME1 and A33-MAIN are used as training data for the processors. The overall classification accuracy after cross-validation is 99.5%. A breakdown of the classification results for each device type is shown in Table 5.4. In the table, the rows correspond to the measured devices and their correct classification, while the columns correspond to the classification determined by the algorithm. The percentage the device is correctly classified appears in bold, while the percentage the device is incorrectly classified appears in red.

As the table demonstrates, the individual classification accuracies for all devices are 98% or higher. Importantly, the algorithm is able to accurately classify each processor regardless of the motherboard it is integrated into. Despite using only examples from A13-MAIN and A20-LIME1 for training, the algorithm correctly classified both A13 devices as having a PROC1B and all three A20 devices as having a PROC2A. At the same time, the algorithm is able to correctly distinguish the A10-LIME from the A13 devices and the A20 devices from the A33-MAIN despite having the similar processors. The differences in how the processor is implemented on the SoC are enough to distinguish them. Furthermore, the algorithm correctly differentiated the processors on the A10-LIME and A20-LIME1 despite the A10-LIME's PCB being an older revision of the A20-LIME1 and the A20 Allwinner being pin-to-pin compatible with the A10 Allwinner.

Furthermore, the average distances for the processors are shown in Table 5.5. As the table demonstrates, the class-distances for all four processors are larger than the self-distances, however, not as much as for the memory components.

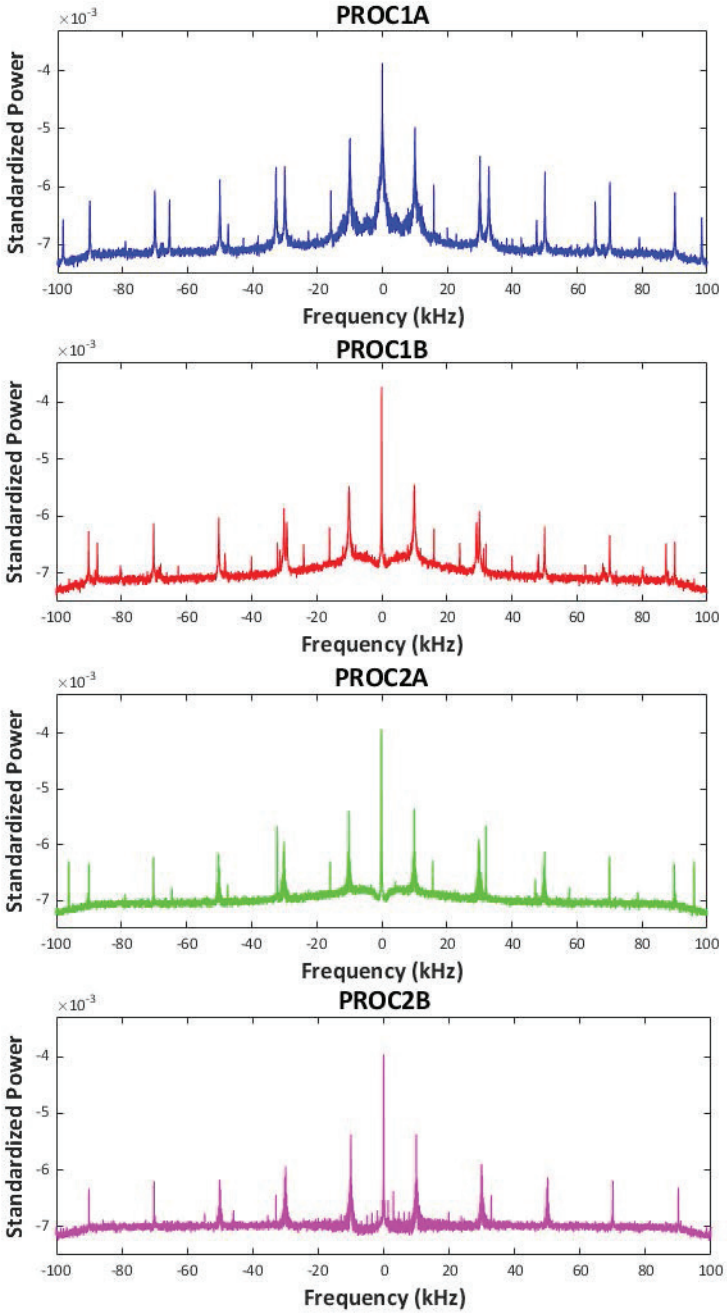

Fig. 5.9. Comparison of the EM signatures from PROC1A (top in blue), PROC1B (second from the top in red), PROC2A (third from the top in green), and PROC2B (bottom in magenta).

TABLE 5.4
Confusion Matrix for the Processors (in %).

	PROC1A	**PROC1B**	**PROC2A**	**PROC2B**
PROC1A: A10-LIME	**100.0**	0	0	0
PROC1B: A13-MAIN	0	**98.0**	2.0	0
PROC1B: A13-MICRO	0	**100.0**	0	0
PROC2A: A20-LIME1	0	0.2	**99.8**	0
PROC2A: A20-LIME2	0	0	**98.7**	1.3
PROC2A: A20-MICRO	0	0	**100.0**	0
PROC2B: A33-MAIN	0	0	0	**100.0**

TABLE 5.5
Average Distances Between Processor Components

	PROC1A	**PROC1B**	**PROC2A**	**PROC2B**
PROC1A	**2.7**	12.6	17.3	22.5
PROC1B	12.6	**2.9**	7.5	13.4
PROC2A	17.3	7.5	**3.0**	6.7
PROC2B	22.5	13.4	6.7	**2.4**

5.5.6 Recognition of Ethernet Transceiver Components

The Ethernet transceivers consist of three types: ETH1, ETH2, and ETH3. ETH1 is used on A10-LIME, ETH2 is used on A20-LIME1 and A20-MICRO, and ETH3 is used on A20-LIME2. The A13-MAIN, A13-MICRO and A33-MAIN do not have Ethernet transceivers.

Examples of the spectrums for ETH1, ETH2, and ETH3 are shown in Fig. 5.10. The spectrums are not as active as the memory and processor, indicating that the excitation program is not having as strong of an effect. Furthermore, the signatures for all three components share some similar features. For instance, all three spectrums have activity appearing every 8 kHz with varying magnitudes. However, despite these factors, the differences in the spectrums are still significant enough to distinguish the components. For example, the carrier for ETH1 has a larger spread than the others. At the same time, it has more instances of weak activity distributed throughout the spectrum. On the other hand, the spectrum from ETH2 has more activity within the first 10 kHz of the carrier. Finally, the spectrum for ETH3 has less interference than the other components.

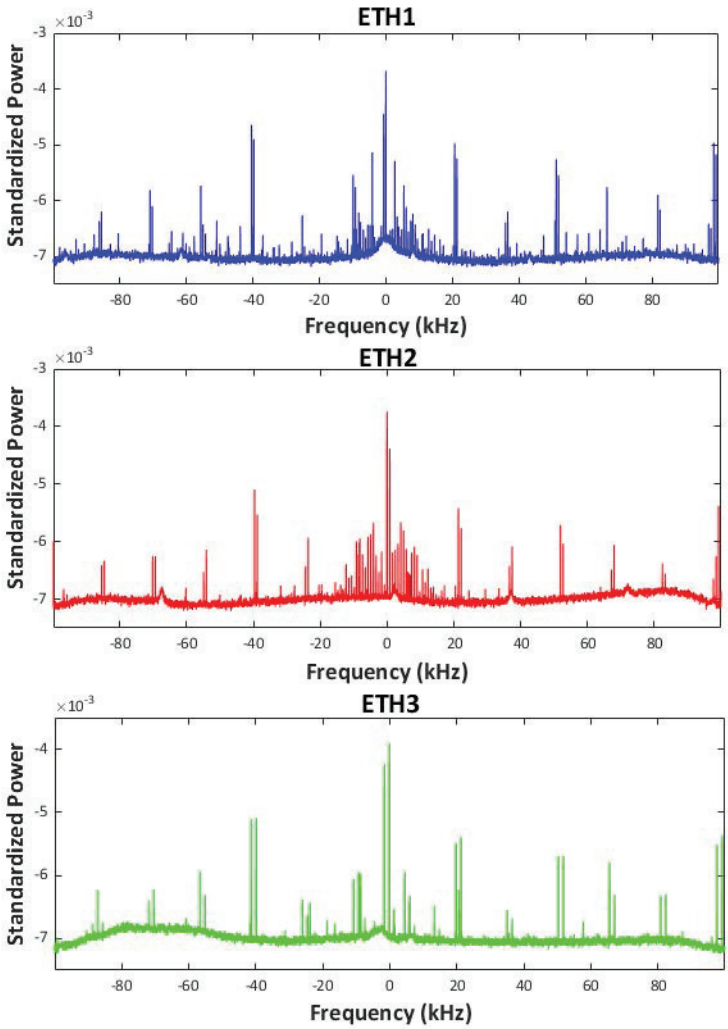

Fig. 5.10. Comparison of the EM signatures from ETH1 (top in blue), ETH2 (middle in red), and ETH3 (bottom in green).

Example measurements from A10-LIME, A20-MICRO, and A20-LIME2 are used for training. The classification accuracies for each individual Ethernet transceiver are shown in Table 5.6. The overall classification accuracy for the transceivers is 97.7%. As the table demonstrates, all the transceivers had some classification error, the worst being the A20-LIME1 with a total error of 3.8%. The errors are likely the result of several factors. First, the features of the signature are relatively weak, making them more vulnerable to noise. Second, there are variations in signatures from the same type of component, making it difficult for the algorithm to correctly group all the measurements from the same transceivers

TABLE 5.6
Confusion matrix for Ethernet Transceivers (in %)

	ETH1	**ETH2**	**ETH3**
ETH1: A10-LIME	**98.2**	1.8	0
ETH2: A20-LIME1	2.3	**96.2**	1.5
ETH3: A20-LIME2	0	3.2	**96.8**
ETH2: A20-MICRO	0	**99.7**	0.3

TABLE 5.7
Average Distances Between Ethernet Transceivers

	ETH1	**ETH2**	**ETH3**
ETH1	**3.5**	11.5	14.4
ETH2	11.5	**3.5**	5.4
ETH3	14.4	5.4	**2.5**

together. Third, features shared between signatures from different types of transceivers make it difficult for the algorithm to distinguish the different types of transceivers. The classification accuracy could be potentially improved by changing the measurement settings. Some possibilities include increasing the measurement bandwidth (to increase the number of features for classification), increasing the measurement time (to decrease the influence of noise), or using a different set of instructions for exciting the component.

The average distances for the Ethernet transceivers are shown in Table 5.7. As the table demonstrates, the class-distances are larger than the self-distances for each transceiver.

5.6 Conclusions

This chapter proposes leveraging EM side-channels to recognize/authenticate components integrated onto a motherboard. By focusing on components on a motherboard, the proposed method provides an opportunity to authenticate devices assembled by third parties. This method identifies components based on the modulated signals emanated during the component's operation. These signals are generated by exciting the component in a controlled and repeatable manner. When testing an unknown component, the spectrum is compared to

previously recorded signatures taken during training. To improve the efficiency, the size of the spectrum is first reduced by projecting it into a vector space generated from training signatures. The identity of the tested component is then determined using a k-NN algorithm. The method has successfully classified components such as external memories, processors, and Ethernet transceivers integrated on seven types of IoT devices. Nine to ten different instances of each device are used in the experiments and then cross-validated during classification. Cross-type testing of the motherboards is conducted as well. Since manufacturers commonly use the same components in multiple designs, being able to collect the training signatures on one motherboard and test components from different motherboards significantly speeds up the process and decreases the cost. Using the measurements taken while the components were excited for 1 s, our method achieved a classification accuracy greater than 96% for all the tested components. These results demonstrate that this method can recognize components based on their EM emanations even if they are integrated on a completely different motherboard.

CHAPTER 6
DETECTION OF RECYCLED ICS USING BACKSCATTERING SIDE-CHANNEL ANALYSIS

6.1 Overview

Recycled ICs are a major concern when manufacturing electronic devices. Not only do they cost designers money, they decrease the lifespan of the device they are integrated onto. The drawback of most detection methods is that they require additional circuitry added to the IC, the PCB, or some other direct interaction. These methods are costly and are limited to PCBs that were designed for them. One of the few methods that does not require additional circuitry, visual inspection, lacks the reliability of the others since it relies on the recyclers making mistakes and not adequately hiding signs of previous use.

SCA provides an attractive alternative for detecting recycled ICs. Since most side channels are a consequence of the device performing its normal operations, they do not require any modifications to the IC. SCA has commonly been used in types of IC authentication, but not for detecting recycled ICs. Unfortunately, most types of SCA are poorly suited for detecting recycled ICs since their causes are not related to aging. However, the newly defined BSCA opens up new opportunities.

To address the challenges of detecting recycled ICs, [17] proposed a new BSCA-based method for detecting recycled ICs. This method uses BSCA to detect changes in the IC caused by aging. Being impedance-based, the backscattering side channel is directly impacted by the increase in impedance caused by aging. It is intended to assist designers in checking questionable components already integrated into their designs by third-party assemblers. Unlike most other detection methods, BSCA can evaluate the ICs non-destructively and without directly interfacing or modifying the IC, making it low-cost and

convenient to use.

This chapter demonstrates the impact aging has on the backscattering side channel and then verifies this behavior in simulations. Both cases show that one of most reliable frequency ranges to monitor is based on the rise and fall times of the transistors (F_{fr}). Other useful frequencies, such as F_{duty}, will depend on the properties of the backscattered signal. Then, we introduce an identification algorithm based on SVD to distinguish unaged and aged circuits from the backscattered measurements. Afterwards the detection method is validated through experiments performed on a series of circuits implemented on FPGAs. The results show that recycled ICs can be detected after being aged for a small faction of the IC's lifetime (roughly 66 days). These experiments also illustrate how the size and complexity of the circuit impacts the accuracy of the detection method. Larger circuits will have a stronger backscattered signal, making them less sensitive to noise and manufacturing variation.

The rest of this chapter is organized as follows. Sections 6.2 and 6.3 discuss the relationship between BSCA and digital switching and demonstrate that relationship through simulations. Section 6.4 describes the approach used for identifying recycled ICs based on the backscattering measurements. Section 6.5 validates our method on experimental data. Section 6.6 concludes this work.

6.2 Understanding the Relationship Between Aging and BSCA

Being impedance-based, the backscattering side channel is directly impacted by transistor aging. As discussed in Section 2.5.1, one of the biggest changes in a transistor is that its on impedence increases. This increases in impedance causes the switching speed of the transistor to slowdown (i.e. the rise and fall times increases).

From the prospective of BSCA, aging affects how quickly the impedance seen by the carrier signal changes states and the overall magnitude of the impedance seen at any point in time. The decrease in the switching speed from aging will strongly impact the backscat-

tered signal's high frequency harmonics. Increasing the impedance of the IC's PMOS transistors increases the rise time, and increasing the impedance of the NMOS increases the fall time. Given how short the rise and fall times are compared to the period, the slowdown is most noticeable in the higher harmonics. The trade-off is that these higher frequencies can be somewhat difficult to measure since they are usually weak and are sensitive to timing variations.

On the other hand, the backscattered signal's lower frequencies will be impacted by the change in modulation factor and duty cycle. The modulation factor of the carrier signal is based on how great the difference in impedance is at the different circuit states [65]. Furthermore, if the changes in the rise and fall times are not equal, they will change the duty cycle of the signal, shifting the distribution of energy in the spectrum. Both of these properties will be affected by the how much the impedance of the NMOS and PMOS transistor differ. Since the effects of BTI and HCI can vary greatly based on the type of MOSFET, it is common for the change in NMOS and PMOS impedance to differ noticeably. If aging has a similar impact on the impedance of both types of transistors (such as in high-k dielectrics), the low frequencies will not change as much; however, the aging can still be detected based on the changes in rise and fall time.

Note that the impact of aging is distinct from the impact a hardware Trojan has on the backscattering side channel. Adding a hardware Trojan affects the modulation factor and, as a result, the lower frequencies of the backscattered signal. Assuming the hardware Trojan is made using the same transistors as the rest of the IC, its transistors will have the same switching speed as the rest of the IC. Therefore, a hardware Trojan will not affect the higher frequencies. Furthermore, the hardware Trojan's impedance will be isolated to one or more small sections of the IC. As a result, its detection can be sensitive to the location measurements are taken at on the IC. On the other hand, the impact of aging can be distributed throughout the IC, making detection less sensitive to measurement position.

6.2.1 Advantages Over Other Detection Methods

BSCA has several advantages over relying on ring oscillators. The most obvious advantage is that it does not require any modifications to the IC. Furthermore, a ring oscillator can only monitor the specific circuit path they are connected to. To test the whole circuit, a ring oscillator needs to be connected to each path. On the other hand, the resolution of backscattering SCA can be changed based on the carrier frequency, measurement height, and the directivity of the antennas. By changing these properties, backscattering can be used to evaluate a small segment of the IC or the entire IC. BSCA is also more resilient to other factors in the environment that ring oscillators are sensitive to.

On the other hand, different types of SCA might seem to be attractive alternatives for detecting aging. Since they are a result of the IC performing its normal operations, SCA usually does not require any modification to the device. In fact, measuring path delay using ring oscillators could be seen as an example of timing SCA, but is not usually defined this way.

While SCA has been used in a wide variety of IC authentication techniques; it is not commonly used for detecting recycled ICs. Two of the most well known types of physical side channels are power and EM [156], [157]. Both side channels are well suited for monitoring the behavior of the device [46]. Since the execution of instructions changes the current flow, both side channels are the direct result of the device's behavior.

Unfortunately, neither are well suited for monitoring gate-level properties of the IC, such as transistor aging. It is possible to measure the slowdown using current-based SCA since the shape of the current burst is related to the switching speed; however, this is not a trivial task. The relationship between switching speed and current flowing through a circuit is nonlinear. Furthermore, the high frequency components needed to determine switching speed are filtered by the power supply or are too weak to accurately measure due to the weak EM emanations.

6.2.2 Impact of Aging on the Backscattered Signal

Since most digital circuits are synchronous, the modulation of the backscattered signal will follow the IC's clock. Its presence throughout the IC makes the clock sensitive to any physical alterations. Furthermore, the clock network is one of the most used parts of the IC, making it likely to experience degradation due to aging.

Given that even new transistors have nonzero rise and fall times, clock signals are commonly approximated as a trapezoidal waveform. The expansion coefficients, c_n, for representing a trapezoidal waveform as a Fourier series can be found using:

$$c_n = -\,\mathrm{j}\frac{A}{2\pi n}\mathrm{e}^{\frac{\mathrm{j}n\pi(\tau-\tau_\mathrm{r})}{T}} \times \left(\mathrm{sinc}(\frac{\pi n\tau_\mathrm{r}}{T})\mathrm{e}^{\frac{\mathrm{j}n\pi\tau}{T}} - \mathrm{sinc}(\frac{\pi n\tau_\mathrm{f}}{T})\mathrm{e}^{\frac{-\mathrm{j}n\pi\tau}{T}}\right), \tag{6.1}$$

where A is the amplitude of the pulse, n is the harmonic number, T is the period of the waveform, τ is the pulse width, τ_r is the rise time, and τ_f is the fall time [158]. According to [158], (6.1) can be estimated as

$$c_n = A\frac{\tau}{T}\mathrm{sinc}(\frac{n\pi\tau}{T})\mathrm{sinc}(\frac{n\pi(\tau_\mathrm{r}+\tau_\mathrm{f})}{2T})\mathrm{e}^{-\mathrm{j}n\pi(\tau+\tau_\mathrm{r})/T}. \tag{6.2}$$

While the PMOS and NMOS transistors are usually designed to produce close rise and fall times, they will not be equal. Uneven changes in in the rise and fall times due to aging will change the duty cycle (τ/T).

As (6.2) demonstrates, the spectral content of the trapezoidal waveform is based on the duty cycle and the rise and fall time. A change in any of these factors shifts the distribution of the energy in the frequency spectrum. This shift is most noticeable near parts of the spectrum where (6.2) approaches 0. The most relevant for this work are frequencies near $2/(\tau_\mathrm{f}+\tau_\mathrm{r})$. This minimum is the result of the sinc function in (6.2) containing τ_r and τ_f. For simplicity, this frequency is referred to as F_fr.

To demonstrate the effect of aging, the coefficients of two trapezoidal waveforms are plotted in Fig. 6.1. For the plots, A is 1 and the frequency ($1/T$) is 50 MHz.

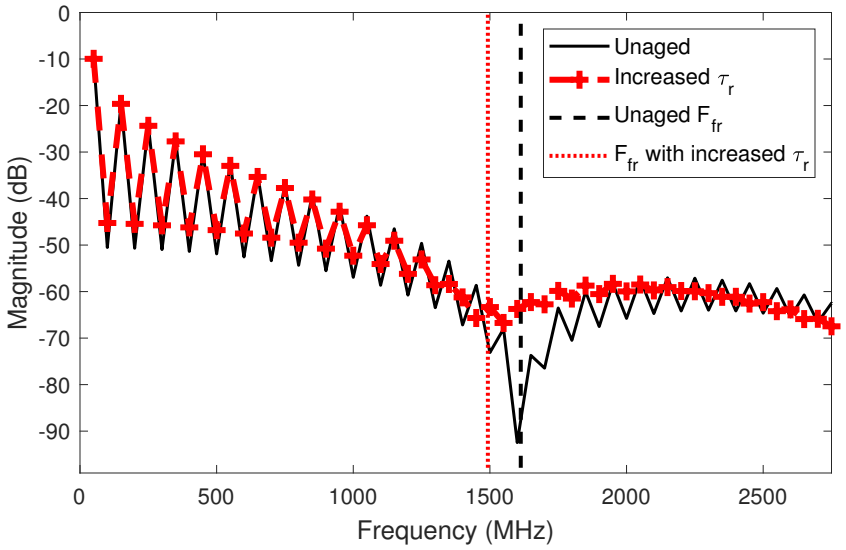

Fig. 6.1. Comparison of the trapezoidal spectrum before and after τ_{r} is increased.

In the figure, the solid black line represents the unaged scenario. It has a rise and fall time of 620 ps and a duty cycle of 49.7%. The unaged plot has a local minimum at F_{fr} (1.613 GHz) and is marked with a black dashed line. The red dashed line with + markers, represents the case where the IC has aged, but only the rise time is affected. This case represents the situation where the effect of NBTI is the dominant aging mechanism. This situation is common for larger ICs or those who are not high-k metal gate devices. In this case, the rise time is increased to 720 ps, causing F_{fr} to decrease to 1.493 GHz. The new F_{fr} is marked with a doted red line. The most noticeable difference between the unaged and aged plots occurs near the change in F_{fr}. Furthermore, the change in the rise time causes the duty cycle to decrease to 49.45%. Since the duty cycle has moved further from 50%, energy shifted from the odd to the even harmonics. This can be seen in the lower frequencies. If the fall time had been increased, the opposite would occur. If both the rise and fall times are increased by the same amount, the duty cycle would remain the same, and the only divergence would be caused by the change in F_{fr}.

Note that other properties of the original signal can impact how easy it is to detect transistor aging. For example, the signal's duty cycle can also produce another local minimum

at $|1/(T-2\tau)|$ or F_{duty}. This minimum is based on the how much the duty cycle diverges from 50%. However, F_{duty} is only useful if the duty cycle diverges from the clock within a small range. If the duty cycle deviates too much from 50%, F_{duty} will be lower than the fundamental frequency of the signal. As a result, aging is unlikely to have a strong enough impact on F_{duty} to be noticeable. On the other hand, if the duty cycle is almost 50%, harmonics near F_{duty} risk being too large to accurately measure.

Fig. 6.2 shows example of how F_{duty} is influenced by aging. All of the settings are the same as in example in Fig. 6.1 except that the duty cycle is decreased to 49%. The original plot again represents the unaged signal with a rise/fall time of 620 ps, while the red dashed line with + markers represents the case where the rise time is increased to 720 ps. The black dashed line represents the unaged F_{duty} while the red doted line represents the aged F_{duty}. The change in rise time causes the duty cycle of the aged signal to decrease to 48.25%, resulting in F_{duty} to shift from 2.5 GHz to 2 GHz. This shift results in a noticeable deviation between the unaged and aged plots in that frequency range. This deviation helps further differentiate the unaged and aged plots. However, given that F_{duty} has a similar effect as F_{fr} and has a small useful range, this work mostly focuses on F_{fr}.

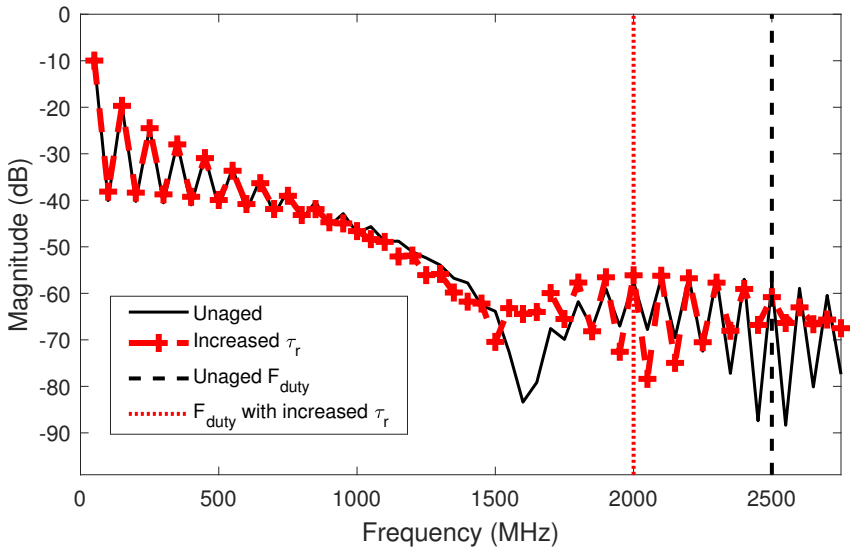

Fig. 6.2. Comparison of the trapezoidal spectrum before and after τ_{r} is increased with a duty cycle of 49%.

6.2.3 Distinguishing Differences Between Unaged and Aged Measurements

As discussed in [64], it can be difficult to directly compare the recorded amplitudes. Slight differences in the measurement position will impact the measured amplitude, potentially causing an IC to be incorrectly classified. Since the attenuation from position errors will be roughly the same for all harmonics, errors from misalignment can be removed by using the amplitude ratios.

To calculate the amplitude ratios, first the unaged measurements and the measurements from the ICs being tested are separated into two sets. Both sets of measurements are converted to decibels, and then each amplitude is subtracted by the amplitude of the following harmonic (harmonic 1 by harmonic 2, harmonic 2 by harmonic 3, etc.). Next the means of each of the unaged ratios is calculated. Both sets of ratios are then normalized by subtracting these means from the corresponding ratio number. As a result, all the unaged ratios will be centered around 0 dB, while the test ratios illustrate how much they deviate from the unaged.

To demonstrate how the measurements are processed, Fig. 6.3 shows a plot of the amplitude ratios for the example in Fig. 6.1. In the figure, the black circles represent the unaged case, and the red triangles represent the aged. This figure demonstrates how much the aged IC's ratios deviate from the unaged. Note that while the aged ratios are also normalized to the unaged, they diverge from 0 dB most noticeably at ratio numbers 25 to 35, near the unaged F_{fr}. This divergence indicates that something about the IC has changed between measurements (in this case that it has aged).

6.3 Simulating Aging

To further demonstrate that our intuition on the effect aging has on BSCA is correct, this section illustrates the effect of aging on a pair of inverters when BSCA is monitored. They are simulated in LTspice, and the schematic is shown in Fig. 6.4. The output of the final

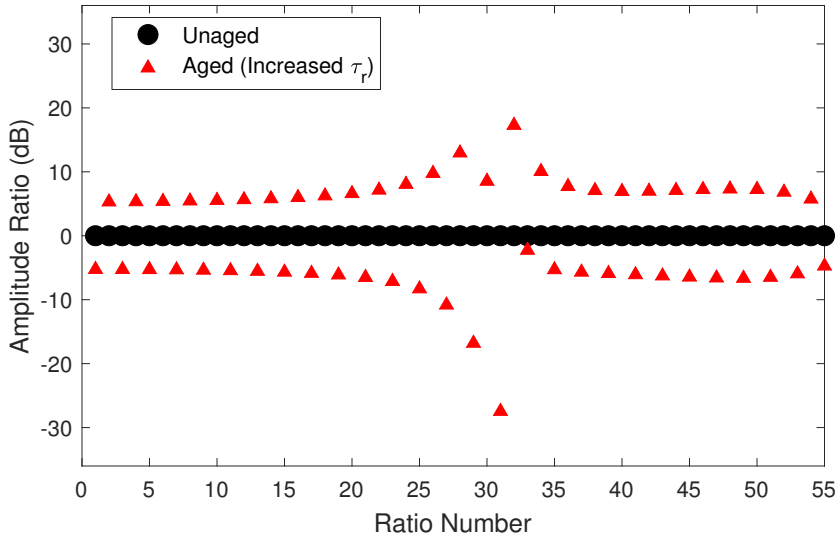

Fig. 6.3. Example of amplitude ratios from unaged and aged cases in Fig. 6.1 normalized to the unaged ratios.

inverter represents the backscattered signal. Each inverter is comprised of an NMOS and PMOS transistor. The properties for the transistors are based on the Predictive Technology Model (PTM) 22 nm low power, high-metal gate NMOS and PMOS models [159]. The sizes of the NMOS and PMOS transistors are selected to ensure that the fall time and the rise time of the inverters are roughly equal. Therefore, the NMOS transistors have a width-to-length ratio (W/L) of $2/1$ and the PMOS transistors have a W/L of $3.63/1$, where L is 22 nm. To represent the capacitance from the connections between inverters, a 10 fF capacitor is connected to the output from each inverter. The value of the capacitor is selected to give each inverter a rise/fall time of roughly 620 ps.

The input of the inverter chain is 1 V, 50 MHz square wave with a duty cycle of 49.7%. To represent the carrier used in the backscattering measurements, a 2.5 mV, 3.011 GHz sinusoid signal is injected into the power supply using a pair of coupled inductors. The properties of the antennas used to transmit the carrier and receive the backscattered signal are ignored.

To represent the impact aging has on the circuit, the magnitudes of the threshold voltages, $|V_{\text{th}}|$, for the NMOS and PMOS transistors are increased by ΔV_{th}. Note that, in

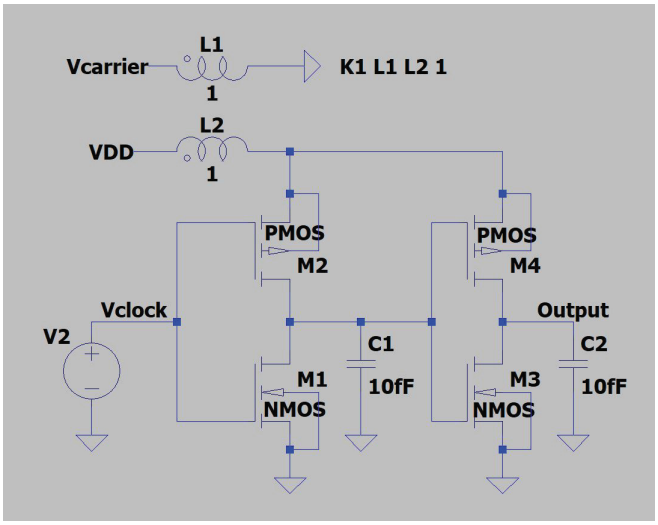

Fig. 6.4. Inverter chain example simulated in LTspice.

reality, ΔV_{th} for the NMOS and PMOS transistors is unlikely to be the same due to differences in the physical processes and materials. To account for this difference, the circuit is simulated in four different scenarios. First, to represent the unaged case, the circuit is simulated while ΔV_{th} is 0. In the second case, only the ΔV_{th} of the PMOS transistors is increased, representing the same situation as in Fig. 6.1. The third case is the opposite of the second, where only the NMOS transistors are aged. The fourth case is the most challenging scenario, where both types of MOSFETs are affected the same amount. To demonstrate the effect increasing ΔV_{th} has on the signal, it is swept from $10\ \mathrm{mV}$ to $30\ \mathrm{mV}$.

The simulated amplitude ratios are plotted in Figs. 6.5–6.7. The first 55 harmonics from each simulation are used in the plots, placing the harmonics between $50\ \mathrm{MHz}$ to $2.75\ \mathrm{GHz}$ away from f_{carrier}. In all three plots, F_{fr} of the original unaged circuit is near ratio number 32 (roughly $1.613\ \mathrm{GHz}$) and is marked with a dashed line.

In all three figures, increasing ΔV_{th} increases how much the ratios deviate from the unaged case. Furthermore, a strong deviation occurs near the unaged F_{fr}. This deviation is a result of the rise/fall time increasing with ΔV_{th}, causing the F_{fr} to decrease. Therefore, F_{fr} provides an estimate of the most useful parts of the spectrum to monitor for signs of aging. The divergence at F_{fr} is not necessary the strongest divergence; however, it is a

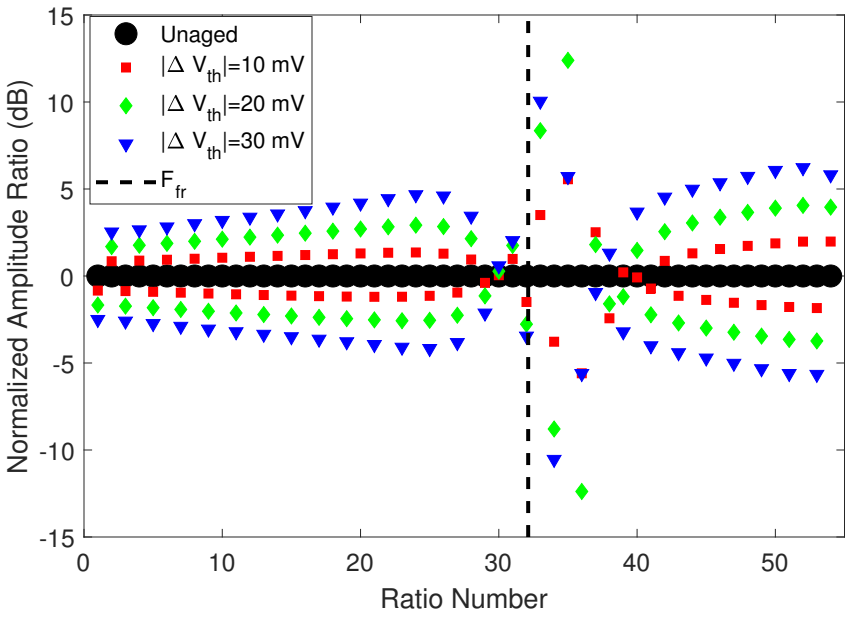

Fig. 6.5. Simulated amplitude ratios as only the PMOS transistors are aged.

reliable location to observe the impact of aging. Despite their similarities, the shapes of the plots differ greatly. These differences are caused by how the timing of the signal changes.

Fig. 6.5 represents the same situation as the example in Fig. 6.3. In both figures, the deviation increases with the ratio number before reaching the unaged F_{fr}. In Fig. 6.3, the deviation rapidly increases at F_{fr}. In Fig. 6.5, the deviation decreases abruptly before F_{fr}, before starting to increase rapidly again. This difference is the result of the clock signal in the simulation not being a perfect trapezoid waveform. The properties of the transistors distort the shape of the waveform. Despite this difference, the figures are fairly similar.

The behavior of the lower frequency ratios in Fig. 6.6 are the opposite of the ratios in Fig. 6.5. While in Fig. 6.5 the ratios' magnitude slowly increases as the number increases, in Fig. 6.6 the magnitude decreases. This difference is based on the impact the rise and fall times have on the signal. Increasing the rise time in Fig. 6.5 shifts the duty cycle further from 50% and causes energy to shift from odd to even harmonics in the spectrum. On the other hand, increasing the fall time in Fig. 6.6 causes the opposite. It shifts the duty cycle closer to 50% and causes energy to shift from even to odd harmonics in the spectrum. If

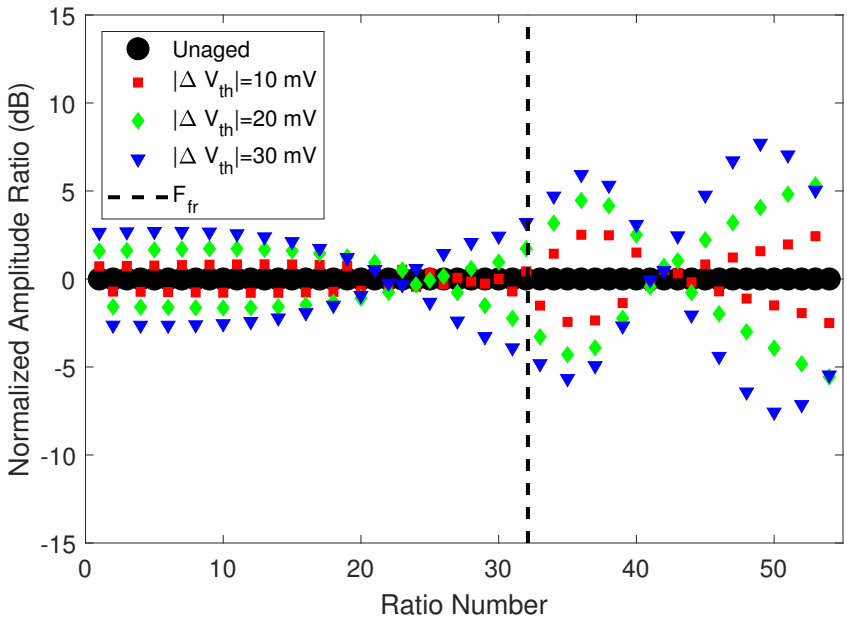

Fig. 6.6. Simulated amplitude ratios as only the NMOS transistors are aged.

the duty cycle had been 50.3%, Figs. 6.5 and 6.6 would be switched.

On the other hand, in Fig. 6.7, there is little deviation in the lower harmonics as ΔV_{th} increases. Since the V_{th} for both types of transistors increases by the same amount, the rise and fall times increase by the same amount. As a result, the duty cycle remains the same, ensuring that most of the deviation occurs near the original F_{fr}. Since the deviation is limited to a smaller, higher frequency range, this scenario can be considered the hardest situation for real measurements.

In all three cases, the deviations shown in the plots are strong enough to determine that the IC has been aged. If the only goal is to detect that the IC is aged, the overall strength of the deviation is the only important factor. However, the shape of the plots is useful in situations where more information about the IC is wanted. The behavior of the lower frequency harmonics is related to the types of NMOS and PMOS transistors used in the IC and how rapidly they age. This information is not commonly provided by the IC manufacturers and could not be determined from ring oscillator measurements.

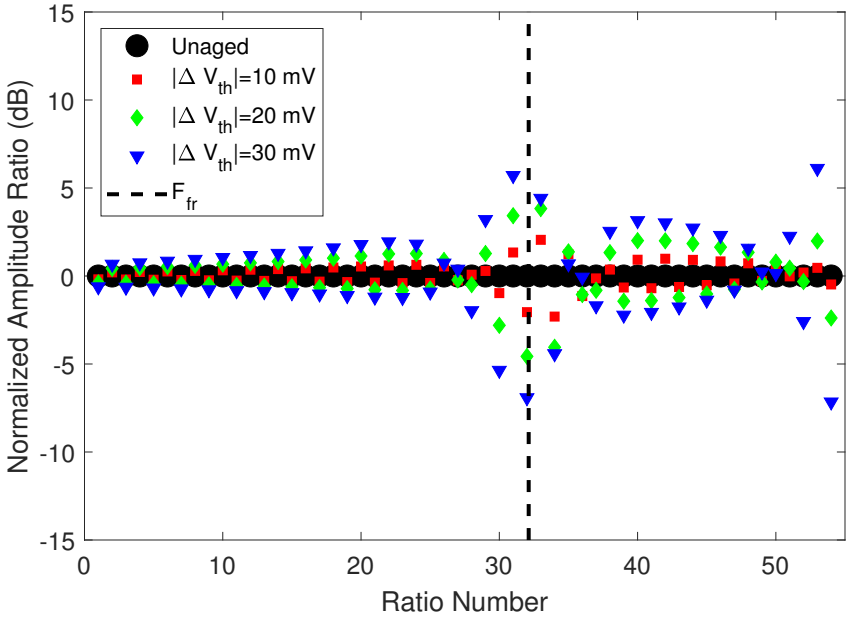

Fig. 6.7. Simulated amplitude ratios as both types of transistors are aged.

6.4 Detecting Aging Using BSCA

In this section we describe the new detection algorithm based on singular value decomposition. This approach is based on work in [16], but to limit the influence of noise and to improve the efficiency of detection, additional improvements were needed. The approach is broken into training and testing phases.

In the training phase, measurements are first taken on n ICs that are known to be new. On each IC, the amplitudes of some number, m, of the clock's harmonics are recorded five times. If the clock has a frequency of f_{C} and the carrier has a frequency of f_{carrier}, measurements are taken at $f_{\mathrm{carrier}} \pm f_{\mathrm{C}}$, $f_{\mathrm{carrier}} \pm 2f_{\mathrm{C}}$, $f_{\mathrm{carrier}} \pm 3f_{\mathrm{C}}$, ..., and $f_{\mathrm{carrier}} \pm mf_{\mathrm{C}}$. Since the upper and lower sidebands contain the same information, only the upper sideband is used in this work. Afterwards, the measurements are combined in the $5n \times m$ matrix, $\mathbf{Y}$. To improve accuracy, harmonics with large variations between ICs or with outliers are removed from the matrix. Next, the amplitude ratios of the measurements are calculated using the approach described in Subsection 6.2.3 and results are stored in a $5n \times m - 1$

matrix. The means used to normalized the training ratios are saved for later use in the testing phase.

To reduce the size of the data and simplify comparison, singular value decomposition is applied to the ratio matrix. The data matrix is decomposed into

$$\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \tag{6.3}$$

where $\mathbf{U}$ is the left singular vectors matrix, $\mathbf{\Sigma}$ is the singular values matrix, and $\mathbf{V}^T$ is the transpose of the right singular vectors matrix. The matrix $\mathbf{V}_K$ can then be used to project the training (and later the testing) data into a new vector space, $\boldsymbol{Z} \in \Re^{H \times K}$, by

$$\boldsymbol{Z} = \mathbf{Y}\mathbf{V}_K, \tag{6.4}$$

where $\mathbf{V}_K$ has been reduced to $K \times K$ matrix. The size of K is smaller than the length of the original measurements and corresponds to the K largest singular values in $\mathbf{\Sigma}$. These K vectors represent the components of the data matrix that correspond to the largest K singular values. The result is a model ($\mathcal{Y}$, $\mathbf{V}$) where the original $5n \times m - 1$ dimensional feature space has been reduced to $5n \times K$ dimensions, without much loss in information. This model is later used in the testing phase.

The testing phase is similar to the training; however, each test IC is processed separately. First, the same harmonics as in the training phase are recorded five times on each of the test ICs. These five measurements are then averaged to minimize noise. Next, the amplitude ratios are calculated using the same process as before, except that the means of the training ratios are used for normalization.

While visually distinguishing between the unaged and aged ratios in the examples from Sections 6.2 and 6.3 is easy, these examples represent idealized cases. In real measurements, process variations between different ICs and noise will make distinguishing unaged and aged measurements more difficult. Therefore, when evaluating actual measure-

ments, the training and testing amplitude ratios are projected into the new vector space using $(\mathcal{Y}, \mathbf{V})$. In the new vector space, all the training points are assumed to belong to the same cluster. The distance between the centroid of the training cluster and the test points is calculated using the city block distance metric. This distance is compared to the distance of every training point from the training cluster's centroid. If the distance for the test point is less than the distance of any of the training points, the IC is assumed to belong to the training cluster (i.e., it is unaged). If the distance of the test point is greater, the IC is classified as not belonging to the training cluster (i.e., it is aged).

An example of unaged and aged measurements projected into the new vector space is shown in Fig. 6.8. The black circles represent the unaged training ratios and each of the red triangles represents a different aged IC. As the figure demonstrates, all the unaged points are clustered close together. On the other hand, the aged points are noticeably far from the unaged points, making it easy to distinguish between the two groups. While only the first three dimensions are represented in the plot, more can be used during analysis.

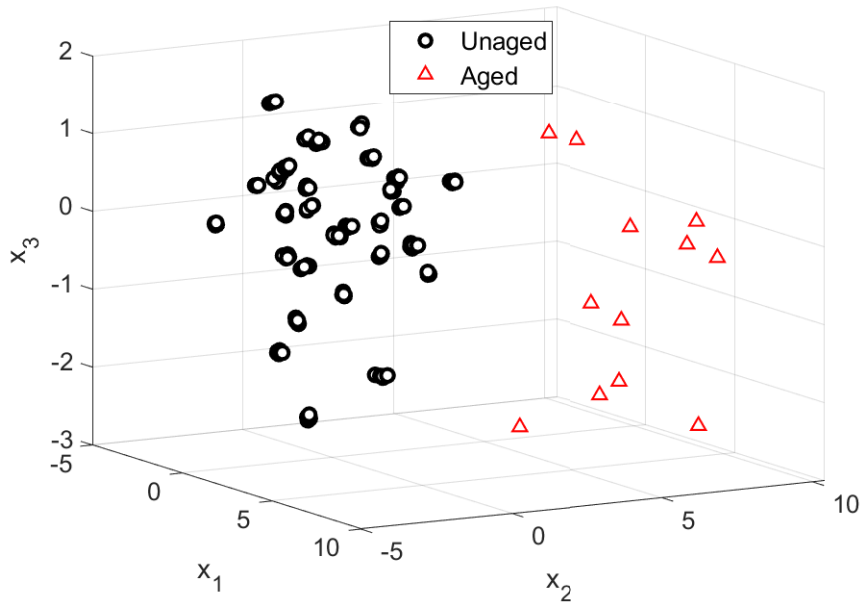

Fig. 6.8. Example of unaged and aged amplitude ratios after being projected into the new vector space.

6.5 Experimental Validation

6.5.1 Setup

To demonstrate the effectiveness of our method, a series of experiments are performed on several Cyclone V FPGAs from Altera and are integrated into DE0-CV FPGA development boards [138]. FPGAs are attractive targets for counterfeiting given their wide use and high price. Furthermore, the FPGAs provide a convenient way of demonstrating the effectiveness over a variety of circuit designs and sizes. However, this approach should also work on other types of ICs.

The Cyclone V is manufactured for Altera by the Taiwan Semiconductor Manufacturing Company (TSMC) using their $28\ \mathrm{nm}$ CMOS low power process (28LP) [160]. This process uses high-k metal gate transistors. Based on the transistor size and manufacturing process, it can be assumed that both NBTI and PBTI will have a notable impact on the FPGA's degradation.

An EM Probe Station Motorized XYZ Table from Riscure is used to position the measurement probes [128]. An X-LSQ150B motorized linear stage positions the FPGA boards. It allows for two FPGAs to be tested before new boards need to be added [161]. An Aaronia E1 electric field probe transmits the carrier at the FPGA, while an Aaronia H2 magnetic field probe receives the backscattered signal [162]. A Keysight M9391 A PXIe vector signal analyzer supplies the $3.011\ \mathrm{GHz}, 20\ \mathrm{dBm}$ carrier and records the backscattered signal.

To accelerate the aging process, the voltage supplied to the Cyclone V was increased from $1.1\ \mathrm{V}$ to $1.96\ \mathrm{V}$. The voltage is returned to $1.1\ \mathrm{V}$ when taking the measurements. During the accelerated aging, the FPGA implements a circular shift register circuit that alternates at $3\ \mathrm{MHz}$. The circuit is comprised of a chain of flip-flops whose inverting output is connected to the input of the next flip-flop in the chain. The inputs are all initialized to the same value and have the same clock [64]. As the cyclical shift register is clocked, the outputs of each flip-flop toggle back and forth. The size of the circuit can be increased

by adding more flip-flops in the chain. To age as much of the FPGA as possible, the shift register is comprised of as many stages as can fit in the FPGA, $37\,000$.

It is difficult to accurately calculate how much this setup accelerates the FPGA's aging. Properties such as the location on the FPGA, its usage in the circuit, temperature, and other manufacturing variations will ensure that transistors at different parts of the FPGA will experience slightly different rates of degradation. Furthermore, the rate of degradation due to BTI and HCI depends on the switching frequency of the transistors. Unfortunately, more accurate aging models require information about the physical properties of the IC that the manufacturers do not supply to their customers. Given the size and composition of the transistors in the Cyclone V, the acceleration factor can be estimated based on the NBTI since it will be one of the dominant processes. The acceleration factor of NBTI is commonly estimated using the voltage power law relation

$$t_{\mathrm{F}} \propto (V_{\mathrm{gs}})^{-\gamma}, \tag{6.5}$$

where t_{F} is the time for the threshold voltage to change by a fixed value, V_{gs} is the voltage applied across the transistor's gate and source, and γ is the voltage-acceleration factor and is between 6 to 8 [93], [163]. Based on this equation, increasing the voltage from $1.1\ \mathrm{V}$ to $1.96\ \mathrm{V}$ accelerates the degradation (i.e., decreases t_{f}) by roughly $66\times$.

For the experiments, 12 Cyclone V FPGAs are aged using this process for three days (which is equivalent to roughly 198 days). Since ΔV_{th} increases logarithmically with time, the most abrupt changes are going to occur early in the FPGA's lifespan, near these first three days [79]. Backscattering measurements are taken after each day of aging. Prior to the measurements, the FPGAs are powered off for a period of time to allow for the reversible component of the BTI to recover [94]. Once the aging is complete, the results are compared to the measurements taken on 30 unaged FPGAs.

Measurements are taken while the FPGA implements several circuits. First, to deter-

mine the impact size and complexity of the circuit-under-test, six different-sized cyclical shift register circuits are tested. For the experiments, circuits that utilize 100%, 50%, 25%, 10%, 5%, and 1% of the Cyclone V's $49\,000$ logic elements are used. As mentioned previously, smaller circuits are more likely to produce a weaker backscattered signal. Since detecting aged ICs already relies on the weaker high harmonics, noise and manufacturing variations risk obscuring the effects of aging. Testing multiple circuit sizes highlights how resilient our method is to this concern.

Then to demonstrate that the backscattering approach works on more practical circuits as well, measurements are taken while the FPGA implements an AES-128 cryptographic processor supplied by TrustHub [164], [165]. This circuit performs 10 stages of AES encryption on a 128-bit block and utilizes 18% of the FPGA's resources.

The clock frequency, f_{C}, is $22\ \mathrm{MHz}$ for the shift register circuit and $21\ \mathrm{MHz}$ for the AES circuit. Based on previous measurements, it is estimated that the FPGA's transistors have an average rise/fall time of roughly $548\ \mathrm{ps}$ ($F_{\mathrm{fr}} = 1.916\ \mathrm{GHz}$). Given that the NMOS and PMOS transistors are both going to have noticeable degradation, it can be assumed that results will be closer to the last scenario in the simulations, where most of the divergence between the unaged and the aged measurements occurred at harmonics near F_{fr}. To save time, only harmonics 51 to 90 (between 4.110 and $4.991\ \mathrm{GHz}$) for the shift register circuits are measured. For the AES circuit, harmonics 53 to 93 (between $4.124\ \mathrm{GHz}$ and $4.964\ \mathrm{GHz}$) are used for evaluating the age.

The approach described in Section 6.4 is used for classifying all of the measurements. Only the first five dimensions of the new vector space are used for evaluating the measurements, i.e., $K = 5$. To account for the limited number of ICs, a six-fold k-fold cross-validation is run for 200 iterations. Such schemes are commonly used since the results have relatively low bias and variance compared to other cross-validation schemes [129]. To evaluate how accurate our approach classified the ICs, the results for each iteration are used to calculate receiver operation characteristics (ROC) and are then averaged.

6.5.2 Impact of Circuit Size

Figure 6.9 shows the amplitude ratios of one FPGA while it is running the shift register with 100% utilization after being aged several days. The black x's are the unaged ratios used as training. Variations in the training ratios are caused by manufacturing differences in the FPGAs. However, as the figure demonstrates, the training varies more as the the ratio number increases. This increase is the result of the amplitudes of the harmonics decreasing, becoming more sensitive to noise. This deviation can become a concern in small circuits, where the backscattered signal is already weak.

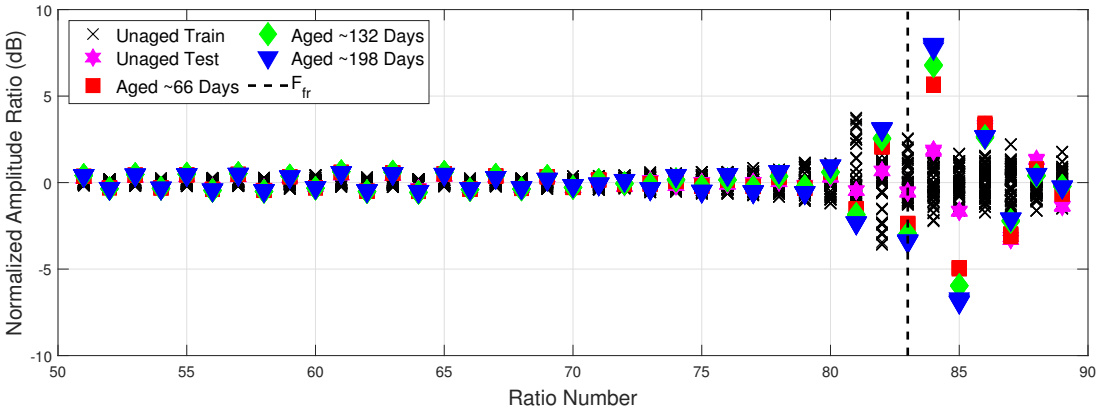

Fig. 6.9. The amplitude ratios of one FPGA implementing the 100% utilization shift register circuit after the equivalent of 198 days of accelerated aging.

The rest of the points represent the amplitude ratios recorded from one FPGA while it is unaged (the purple hexagrams), after being aged the equivalent of 66 days (the red squares), aged 132 days (the green diamonds), and aged 198 days (the blue triangles). The unaged F_{fr} is marked with a dashed line. The ratios' behavior is similar to the simulation where the NMOS and PMOS transistors are both aged. As in Fig. 6.7, the lower aged harmonics only slightly deviate from the unaged ratios. The lower frequency harmonics that are not recorded would have followed the same trend. Therefore, they were not included. As the ratio number increases, the deviation slightly increases until reaching ratio 83, where there is a noticeable difference between the training and testing. This difference is caused by

the change in rise and fall time. Furthermore, as in the simulations, the deviation from the unaged ratios increases the longer the FPGA is aged. The ratios for the rest of the FPGAs and for the other shift register utilizations follow a similar trend as the Fig. 6.9.

Figs. 6.10 (a)-(f) summarize the classification results for the shift register circuits with the ROC curves. Each curve shows the ROC plot for the circuit found using the ratios taken after aging the FPGAs 66, 132, and 198 days. The area under the curve (AUC) is also provided for each case. All six circuits follow the same trend. After being aged for 66 days, all the circuits except the 1% have an AUC of greater than 0.93. After 132 days of aging, the AUC of the 100% and 50% are greater than 0.99, while the 25%, 10% and 5% are greater than 0.96. After 198 days of aging, the 100% through 10% utilizations have an AUC of greater than 0.99, while the 5% is greater than 0.98 and the 1% is almost 0.98. As the plots demonstrate, our BSCA-based method can easily identify that large ICs have been aged after only short period of time. For the 100% and 50% utilizations, our method had almost 100% accuracy after 198 days of aging. Unsurprisingly, the accuracy deceases as the size of the IC decreases; however, the AUCs of the 5% and 1% are still over 0.97. In these circuits, the backscattered signal is weak enough that the effects of aging can become obscured by noise and manufacturing variation. These results can be improved by increasing the power of the carrier or aging the IC's longer.

6.5.3 AES Circuit Results

The shift register circuits used in the previous example are fairly simple. During operation, these circuits are switching between only two states, causing the backscattering side channel to only see one of two impedances at any point in time. Similar to the example in Section 6.3, the impedance change will have a trapezoidal pattern. Most circuits are going to have more than two states, making them more complicated. As a result, the impedances seen by the backscattered signal will vary every clock cycle as different transistors are toggled off and on to perform specific functions.

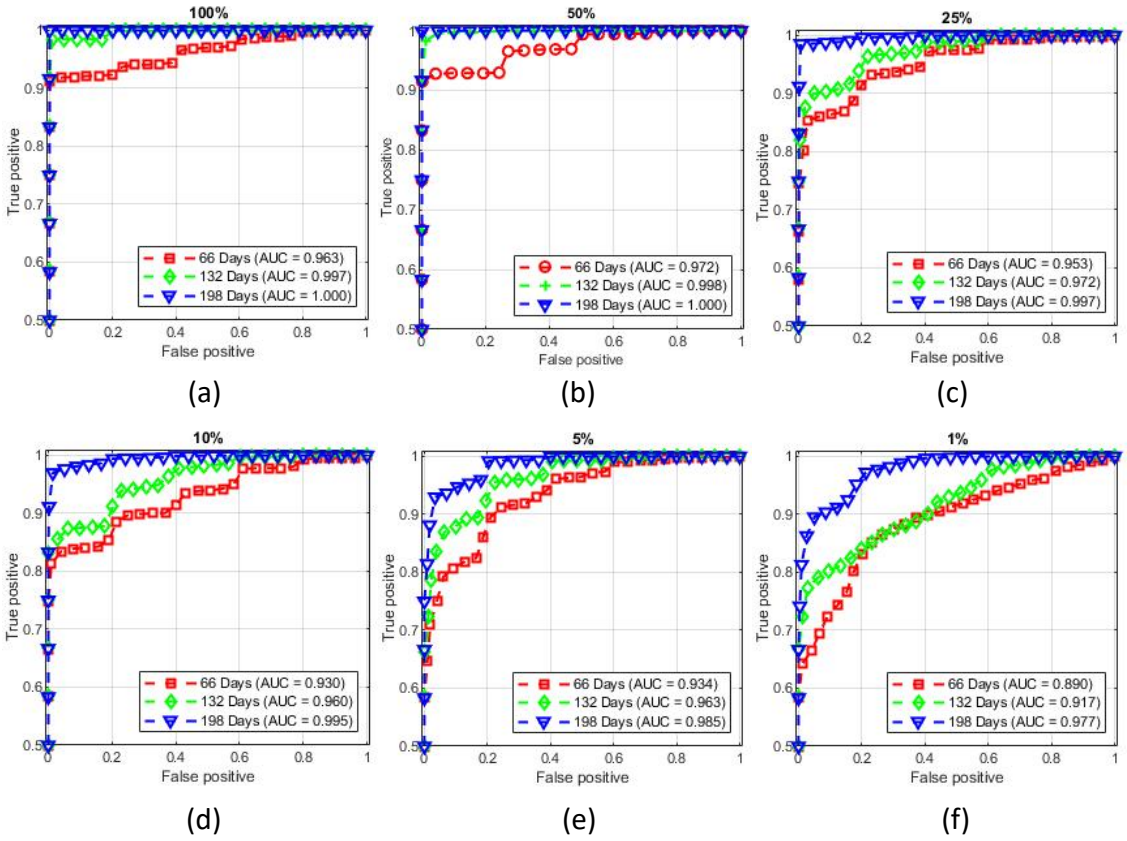

Fig. 6.10. ROC curves for the six shift register circuits after each round of aging. (a) 100%, (b) 50%, (c) 25%, (d) 10%, (e) 5%, and (f) 1% utilization.

To demonstrate that our method works on realistic circuits, the AES circuit is tested in the same manner as the shift register circuits. Fig. 6.11 shows how aging impacts the amplitude ratios from one of the FPGAs as it implements the AES circuit. The black x's are the unaged ratios used as training, the purple hexagrams are the ratios from the FPGA prior to aging, the red squares are the ratios after 66 days, the green diamonds are the ratios after 132 days, and the blue triangles are the ratios after 198 days. The dashed black line again represents the original F_{fr}, which, since the AES has a clock of 21 MHz, is near harmonic 86.

The overall shape of the amplitude ratios is distinct from Fig. 6.9. While both have large deviations occurring near F_{fr}, the AES ratios also have a noticeable deviation at ratios 53 to 57. This second deviation is the result of another local minimum located at roughly

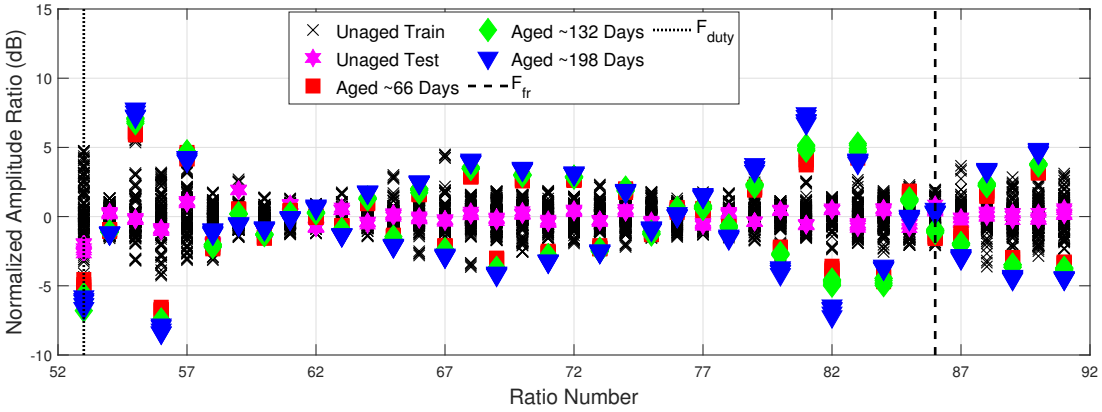

Fig. 6.11. The amplitude ratios of one FPGA implementing the AES circuit after the equivalent of 198 days aging.

1.113 GHz from the carrier frequency. This second minimum is likely F_{duty} (which is marked with a dotted line near ratio 53), indicating that the impedance change has a duty cycle of roughly 49.06% or 50.94%. While the clock's duty cycle is set to 50% in the FPGA's programming files, the difference in the duty cycle is likely the result of limitations of the FPGA's routing program and how the impedances change every clock cycle. F_{duty} was not present in the shift register circuits, indicating that the duty cycle is closer to 50% and out of the measurement range.

The ROC curves for the AES circuit are shown in Fig. 6.12. Despite utilizing only 18% of the FPGA's resources, our method had an accuracy of 100% (with an AUC of 1) after only one day of aging. This is better than the performance on the 100% shift register circuit. The extra deviations due to F_{duty} help the detection algorithm differentiate unaged and aged measurements.

6.6 Conclusions

In this chapter, a new method for detecting recycled ICs is presented. This method uses BSCA to detect changes in the IC's transistors' impedances caused by aging. Unlike other types of side channels, backscattering is directly impacted by aging, making it well suited for this application. Furthermore, we introduce an identification algorithm to correctly

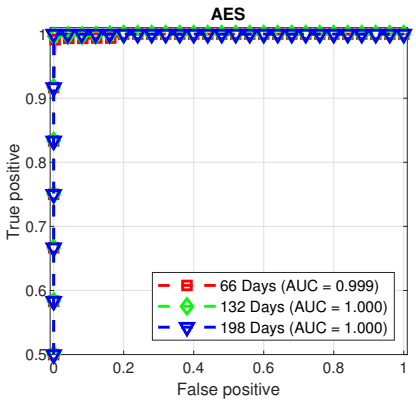

Fig. 6.12. ROC curves for the AES circuit after each day of aging.

distinguish unaged and aged circuits from the backscattered measurements. This method is intended to assist designers in checking questionable components already integrated into their designs by third-party assemblers. Unlike most other detection methods, our method can evaluate the ICs non-destructively and without directly interfacing or modifying the IC, making it low-cost and convenient to use.

First the impact aging has on the backscattering side channel is demonstrated and then verified in simulations. Both cases show that one of most reliable frequency ranges to monitor is based on the rise and fall times of the transistors (F_{fr}). Other useful frequencies, such as F_{duty} will depend on the properties of the backscattered signal. We then validated our detection method through experiments performed on a series of circuits implemented on FPGAs. These experiments demonstrate that the impact of aging outweighs the manufacturing variations between individual ICs. After the equivalent of only 66 days, a fraction of the IC's lifespan, a majority of the aged circuits could be correctly classified. Furthermore, these experiments demonstrate that the size and complexity of the circuit impacts the accuracy of the detection method. Larger circuits will have a stronger backscattered signal, making them less sensitive to noise and manufacturing variation.

CHAPTER 7
SUMMARY AND FUTURE WORK

7.1 Summary

As discussed in Chapter 1, this dissertation is focused on providing hardware designers with methods of accessing the security of their devices.

The first half of this work was focused on the sources of EM side-channel leakage on PCBs. To help address this concern, two methods for locating the sources of EM side channels were introduced. Not only does knowing the sources of the EM side channel help designers address leakage from their devices, these sources can also help detect other security concerns such as hardware Trojans and malware. Instead of trying to identify the leakage sources for entire programs, like other techniques, both of our methods focus on relating the leakage sources to the basic instructions executed on the IC. This allows designers to address the root cause of the leakage. The first method focused on low-frequency (less than 100 MHz) magnetic field sources, while the second method focused on higher frequency electric and magnetic field sources. To accurately locate the higher frequency signals, we also developed a measurement setup that limits errors in the measurements due to reflections from nearby objects. Both methods require only measurements taken around the edges of the device, drastically speeding up the measurement time.

The effectiveness of both methods were demonstrated through a series of experiments performed on two devices that represent a wide variety of embedded and IoT systems. These experiments demonstrated that at low frequencies, the sources of the side channel are the larger current carrying structures on the device, such as the decoupling capacitors. However, the location of the sources change as frequency increases since not only does the radiating efficiency of smaller parts of the device improve, the causes of the sig-

nal change. Furthermore, the number and locations of the leakage sources depend on the specific program activity. As a result, when executing more complex programs, the side-channel sources can rapidly shift across the device as different resources are used to execute different instruction. This shift emphasizes the importance of focusing on leakage from basic processor instructions instead on specific programs.

The second half of this work introduced two methods for authenticating components integrated onto a PCB. Both methods are intended to provide reliable ways for designers to authenticate devices assembled by third parties. The first leveraged EM SCA to differentiate between similar types of components. This method used an SVD-based algorithm to efficiently identify different components. The experiments demonstrated that the method could efficiently distinguish between several commonly used types of components integrated onto multiple PCBs. We also performed cross-type testing to demonstrate that the method can identify a component even when integrated onto several different PCBs. Since designers commonly use the same components in multiple designs, this significantly speeds up the test process.

The second method focused on detecting recycled ICs. This method leverages BSCA for detection. Unlike most detection approaches, BSCA makes it possible to detect aging without integrating additional circuitry onto the IC. This cuts the cost and allows for testing a variety of designs without needing to modify them. We described the impact aging has on the backscattering side channel and demonstrate it through simulation. At the same time, we identified the most reliable frequency ranges to monitor. Furthermore, we outlined a new SVD-based algorithm for efficient distinguishing between unaged and aged IC from the backscattered measurements. Finally the detection method was validated through a series of experiments. These experiments also investigated the impact the size and complexity had on the accuracy.

7.2 Future Work

This work has several avenues for further research. One practical improvement would be unifying the component authentication methods described in Chapters 5 and 6 and the hardware Trojan detection approach described in [64]. Having a single approach for detecting several of the biggest types of counterfeit ICs would greatly simplify the testing. While the authentication approach in Chapter 5 relies on EM SCA, it could potentially be adapted to work with BSCA. The main concern is distinguishing what caused the change in the tested component (whether its a different IC, has a hardware Trojan, or recycled). As discussed in Chapter 6, hardware Trojans and recycling are likely to effect different parts of the backscattered spectrum. However, these differences need to be quantified and compared with how using a completely different IC will impact the spectrum.

By itself, the recycled IC detection method has several areas that can be further investigated. Like most detection techniques, the current need for golden (known to be unaged) ICs can be inconvenient for designers. However, this requirement can be eliminated with enough information about the physical properties of the ICs' transistors [166]. In theory, this information can be obtained from the ICs' manufacturers and their foundries; however, manufacturers are usually resistant to supply this information. The other option is to obtain this information from other parametric tests on the IC. While this can be time consuming, it needs to be performed only once.

Furthermore, the recycled IC detection method can also be used for evaluating the age of ICs currently in use. This would provide a reliable way of detecting the potential failure of important equipment, giving an opportunity to replace it before it fails at a critical moment. To accomplish this, the effect on the IC as it ages needs to be better quantified. While it is straightforward to predict the impact aging will have on the backscattered signal, the state of the side channel prior to failure needs to be known. In other words, how much the backscattering side channel will deviate before the IC fails needs to be known.

To accurately model the state prior to failure requires several measurements or information from the manufacturer.

REFERENCES

[1] A. Nazari, N. Sehatbakhsh, M. Alam, A. Zajic, and M. Prvulovic, "EDDIE: EM-based detection of deviations in program execution," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, 2017, pp. 333–346.

[2] N. Sehatbakhsh *et al.*, "REMOTE: Robust external malware detection framework by using electromagnetic signals," *IEEE Transactions on Computers*, 2019.

[3] H. A. Khan, N. Sehatbakhsh, L. N. Nguyen, M. Prvulovic, and A. Zajić, "Malware detection in embedded systems using neural network model for electromagnetic side-channel signals," *Journal of Hardware and Systems Security*, vol. 3, no. 4, pp. 305–318, 2019.

[4] P. Maheshwari, V. Khilkevich, D. Pommerenke, H. Kajbaf, and J. Min, "Application of emission source microscopy technique to EMI source localization above 5 GHz," in *Electromagnetic Compatibility (EMC), 2014 IEEE International Symposium on*, IEEE, 2014, pp. 7–11.

[5] P. Maheshwari, H. Kajbaf, V. V. Khilkevich, and D. Pommerenke, "Emission source microscopy technique for EMI source localization," *IEEE Transactions on Electromagnetic Compatibility*, vol. 58, no. 3, pp. 729–737, 2016.

[6] J. Balasch, B. Gierlichs, and I. Verbauwhede, "Electromagnetic circuit fingerprints for hardware trojan detection," in *2015 IEEE International Symposium on Electromagnetic Compatibility (EMC)*, Aug. 2015, pp. 246–251.

[7] A. Gorbunova, A. Baev, M. Konovalyuk, and Y. Kuznetsov, "Localization of cyclostationary EMI sources based on near-field measurements," in *Electromagnetic Compatibility (EMC), 2015 IEEE International Symposium on*, IEEE, 2015, pp. 450–455.

[8] HINT-PROJECT. "Holistic Approaches for Integrity of ICT-Systems." (2012), [Online]. Available: http://www.hint-project.eu.

[9] U. Guin, K. Huang, D. DiMase, J. M. Carulli, M. Tehranipoor, and Y. Makris, "Counterfeit integrated circuits: A rising threat in the global semiconductor supply chain," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1207–1228, Aug. 2014.

[10] N. Kae-Nune and S. Pesseguier, "Qualification and testing process to implement anti-counterfeiting technologies into IC packages," in *2013 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, IEEE, 2013, pp. 1131–1136.

[11] M. Pecht and S. Tiku, "Bogus: Electronic manufacturing and consumers confront a rising tide of counterfeit electronics," *IEEE spectrum*, vol. 43, no. 5, pp. 37–46, 2006.

[12] F. Werner, D. A. Chu, A. R. Djordjević, D. I. Olćan, M. Prvulovic, and A. Zajić, "A method for efficient localization of magnetic field sources excited by execution of instructions in a processor," *IEEE Transactions on Electromagnetic Compatibility*, vol. 60, no. 3, pp. 613–622, 2017.

[13] R. Callan, A. Zajić, and M. Prvulovic, "FASE: Finding amplitude-modulated side-channel emanations," in *2015 ACM IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, IEEE, 2015, pp. 592–603.

[14] M. Prvulovic, A. Zajić, R. L. Callan, and C. J. Wang, "A method for finding frequency-modulated and amplitude-modulated electromagnetic emanations in computer systems," *IEEE Transactions on Electromagnetic Compatibility*, vol. 59, no. 1, pp. 34–42, 2016.

[15] F. T. Werner, J. Dinkić, D. Olćan, A. Djordjević, M. Prvulovic, and A. Zajić, "An efficient method for localization of magnetic field sources that produce high-frequency side-channel emanations," *IEEE Transactions on Electromagnetic Compatibility*, vol. 63, no. 6, pp. 1799–1811, 2021.

[16] F. T. Werner, B. B. Yilmaz, M. Prvulovic, and A. Zajić, "Leveraging EM side-channels for recognizing components on a motherboard," *IEEE Transactions on Electromagnetic Compatibility*, vol. 63, no. 2, pp. 502–515, 2020.

[17] F. T. Werner, M. Prvulovic, and A. Zajić, "Detection of recycled ICs using backscattering side-channel analysis," *Submitted to Transactions on Very Large Scale Integration (VLSI) Systems*,

[18] M. G. Kuhn, "Compromising emanations of LCD TV sets," *IEEE Transactions on Electromagnetic Compatibility*, vol. 55, no. 3, pp. 564–570, Jun. 2013.

[19] Y. I. Hayashi *et al.*, "Analysis of electromagnetic information leakage from cryptographic devices with different physical structures," *IEEE Transactions on Electromagnetic Compatibility*, vol. 55, no. 3, pp. 571–580, Jun. 2013.

[20] H. Sekiguchi and S. Seto, "Study on maximum receivable distance for radiated emission of information technology equipment causing information leakage," *IEEE Transactions on Electromagnetic Compatibility*, vol. 55, no. 3, pp. 547–554, Jun. 2013.

[21] K. Gandolfi, C. Mourtel, and F. Olivier, "Electromagnetic analysis: Concrete results," in *International workshop on cryptographic hardware and embedded systems*, Springer, 2001, pp. 251–261.

[22] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, "The EM side—channel (s)," in *International workshop on cryptographic hardware and embedded systems*, Springer, 2002, pp. 29–45.

[23] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Annual International Cryptology Conference*, Springer, 1999, pp. 388–397.

[24] M. Backes, M. Dürmuth, S. Gerling, M. Pinkal, and C. Sporleder, "Acoustic side-channel attacks on printers.," in *USENIX Security symposium*, 2010, pp. 307–322.

[25] M. Hutter and J.-M. Schmidt, "The temperature side channel and heating fault attacks," in *International Conference on Smart Card Research and Advanced Applications*, Springer, 2013, pp. 219–235.

[26] L. N. Nguyen, C.-L. Cheng, M. Prvulovic, and A. Zajić, "Creating a backscattering side channel to enable detection of dormant hardware trojans," *IEEE transactions on very large scale integration (VLSI) systems*, vol. 27, no. 7, pp. 1561–1574, 2019.

[27] P. C. Kocher, "Timing attacks on implementations of diffie-hellman, RSA, DSS, and other systems," in *Annual International Cryptology Conference*, Springer, 1996, pp. 104–113.

[28] Y. Yarom and K. Falkner, "FLUSH+ RELOAD: A high resolution, low noise, l3 cache side-channel attack," in *23rd USENIX Security Symposium, (USENIX Security 14)*, 2014, pp. 719–732.

[29] C. Percival, *Cache missing for fun and profit*, 2005.

[30] P. Kocher *et al.*, "Spectre attacks: Exploiting speculative execution," in *2019 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2019, pp. 1–19.

[31] M. Lipp *et al.*, "Meltdown: Reading kernel memory from user space," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 973–990.

[32] S. Sangodoyin *et al.*, "Side-channel propagation measurements and modeling for hardware security in IoT devices," *IEEE Transactions on Antennas and Propagation*, vol. 69, no. 6, pp. 3470–3484, 2020.

[33] M. Alam *et al.*, "One&done: A single-decryption EM-based attack on OpenSSL's constant-time blinded RSA," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 585–602.

[34] D. Easter, "The impact of 'Tempest' on anglo-american communications security and intelligence, 1943–1970," *Intelligence and National Security*, pp. 1–16, 2020.

[35] H. J. Highland, "Electromagnetic radiation revisited," *Computers and Security*, pp. 85–93, Dec. 1986.

[36] W. Van Eck, "Electromagnetic radiation from video display units: An eavesdropping risk?" *Computers & Security*, vol. 4, no. 4, pp. 269–286, 1985.

[37] M. A. Ahsan, S. R. Islam, and M. A. Islam, "A countermeasure for compromising electromagnetic emanations of wired keyboards," in *Computer and Information Technology (ICCIT), 2014 17th International Conference on*, IEEE, 2014, pp. 241–244.

[38] M. Vuagnoux and S. Pasini, "An improved technique to discover compromising electromagnetic emanations," in *2010 IEEE International Symposium on Electromagnetic Compatibility*, Jul. 2010, pp. 121–126.

[39] H. S. Lee, J.-G. Yook, and K. Sim, "An information recovery technique from radiated electromagnetic fields from display devices," in *Electromagnetic Compatibility (APEMC), 2016 Asia-Pacific International Symposium on*, IEEE, vol. 1, 2016, pp. 473–475.

[40] A. V. Ivanov, I. L. Reva, and A. E. Ushakov, "Features of identification and the analysis of collateral electromagnetic radiations from USB flash drives," in *Actual Problems of Electronics Instrument Engineering (APEIE), 2016 13th International Scientific-Technical Conference on*, IEEE, vol. 2, 2016, pp. 156–158.

[41] V. Carlier, H. Chabanne, E. Dottax, and H. Pelletier, "Generalizing square attack using side-channels of an AES implementation on an FPGA," in *Field Programmable Logic and Applications, 2005. International Conference on*, IEEE, 2005, pp. 433–437.

[42] K. Gandolfi, C. Mourtel, and F. Olivier, "Electromagnetic analysis: Concrete results," in *Proceedings of Cryptographic Hardware and Embedded Systems - CHES 2001*, 2001, pp. 251–261.

[43] A. Zajic and M. Prvulovic, "Experimental demonstration of electromagnetic information leakage from modern processor-memory systems," *IEEE Transactions on Electromagnetic Compatibility*, vol. 56, no. 4, pp. 885–893, Aug. 2014.

[44] M. G. Kuhn, "Compromising emanations: Eavesdropping risks of computer displays," Ph.D. dissertation, University of Cambridge, 2002.

[45] D. Genkin, I. Pipman, and E. Tromer, "Get your hands off my laptop: Physical side-channel key-extraction attacks on PCs," *Journal of Cryptographic Engineering*, vol. 5, no. 2, pp. 95–112, 2015.

[46] L. N. Nguyen, C.-L. Cheng, F. T. Werner, M. Prvulovic, and A. Zajic, "A comparison of backscattering, EM, and power side-channels and their performance in detecting software and hardware intrusions," *Journal of Hardware and Systems Security*, pp. 1–16, 2020.

[47] D. Agrawal, S. Baktir, D. Karakoyunlu, P. Rohatgi, and B. Sunar, "Trojan detection using IC fingerprinting," in *2007 IEEE Symposium on Security and Privacy (SP '07)*, 2007, pp. 296–310.

[48] M. M. Ahmed *et al.*, "Towards a robust and efficient EM based authentication of FPGA against counterfeiting and recycling," in *2017 19th International Symposium on Computer Architecture and Digital Systems (CADS)*, Dec. 2017, pp. 1–6.

[49] L. N. Nguyen, B. B. Yilmaz, C.-l. Cheng, M. Prvulovic, and A. Zajić, "A novel clustering technique using backscattering side channel for counterfeit IC detection," in *2020 SPIE Defense + Commercial Sensing Digital Forum*, 2020.

[50] L. Sauvage, S. Guilley, and Y. Mathieu, "Electromagnetic radiations of FPGAs: High spatial resolution cartography and attack on a cryptographic module," *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, vol. 2, no. 1, pp. 1–24, 2009.

[51] L. Wang, C. Zhou, and B. Yu, "Laboratory test and mechanism analysis on electromagnetic compromising emanations of PS/2 keyboard," in *Computational Intelligence and Security (CIS), 2012 Eighth International Conference on*, IEEE, 2012, pp. 657–660.

[52] L. Nowosielski, R. Przesmycki, and M. Nowosielski, "Compromising emanations from VGA and DVI interface," in *Progress in Electromagnetic Research Symposium (PIERS)*, IEEE, 2016, pp. 1024–1028.

[53] T.-L. Song, Y.-R. Jeong, and J.-G. Yook, "Modeling of leaked digital video signal and information recovery rate as a function of SNR," *IEEE Transactions on Electromagnetic Compatibility*, vol. 57, no. 2, pp. 164–172, 2015.

[54] D. Sun, J. Shi, X. Ding, M. Zhang, and W. Huang, "Method for detecting information leakage from computer display in electromagnetic radiation," in *2016 IEEE Trustcom/BigDataSE/ISPA*, IEEE, 2016, pp. 2041–2046.

[55] S. B. Dhia, M. Ramdani, and E. Sicard, *Electromagnetic Compatibility of Integrated Circuits: Techniques for low emission and susceptibility*. Springer Science & Business Media, 2006, ch. 3, pp. 3, 9, 120, 201.

[56] M. I. Montrose, "How decoupling capacitors may cause radiated EMI," in *Electromagnetic Compatibility Symposium-Perth (EMCSA), 2011*, IEEE, 2011, pp. 1–4.

[57] R. Callan, A. Zajić, and M. Prvulovic, "A practical methodology for measuring the side-channel signal available to the attacker for instruction-level events," in *2014 47th Annual IEEE ACM International Symposium on Microarchitecture*, IEEE, 2014, pp. 242–254.

[58] G. Perin, L. Torres, P. Benoit, and P. Maurine, "Amplitude demodulation-based EM analysis of different RSA implementations," in *2012 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, IEEE, 2012, pp. 1167–1172.

[59] D. Genkin, L. Pachmanov, I. Pipman, and E. Tromer, "Stealing keys from PCs using a radio: Cheap electromagnetic attacks on windowed exponentiation," in *International workshop on cryptographic hardware and embedded systems*, Springer, 2015, pp. 207–228.

[60] O. Meynard, D. Réal, F. Flament, S. Guilley, N. Homma, and J.-L. Danger, "Enhancement of simple electro-magnetic attacks by pre-characterization in frequency domain and demodulation techniques," in *2011 Design, Automation & Test in Europe*, IEEE, 2011, pp. 1–6.

[61] A. Yaghjian, "An overview of near-field antenna measurements," *IEEE Transactions on antennas and propagation*, vol. 34, no. 1, pp. 30–45, 1986.

[62] "IEEE recommended practice for near-field antenna measurements," *IEEE Std 1720-2012*, pp. 1–102, 2012.

[63] T. Stadtler, L. Eifler, and J. Ter Haseborg, "Double probe near field scanner, a new device for measurements in time domain," in *2003 IEEE Symposium on Electromagnetic Compatibility. Symposium Record (Cat. No. 03CH37446)*, IEEE, vol. 1, 2003, pp. 86–90.

[64] L. N. Nguyen, C.-L. Cheng, M. Prvulovic, and A. Zajić, "Creating a backscattering side channel to enable detection of dormant hardware trojans," *IEEE Transactions on Very Large Scale Integration (VLSI) systems*, vol. 27, no. 7, pp. 1561–1574, 2019.

[65] C.-L. Cheng, L. N. Nguyen, M. Prvulovic, and A. Zajić, "Exploiting switching of transistors in digital electronics for RFID tag design," *IEEE Journal of Radio Frequency Identification*, vol. 3, no. 2, pp. 67–76, 2019.

[66] M. Tehranipoor and F. Koushanfar, "A survey of hardware trojan taxonomy and detection," *IEEE Design Test of Computers*, vol. 27, no. 1, pp. 10–25, 2010.

[67] H. Dogan, D. Forte, and M. M. Tehranipoor, "Aging analysis for recycled FPGA detection," in *2014 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, Oct. 2014, pp. 171–176.

[68] X. Zhang and M. Tehranipoor, "Design of on-chip lightweight sensors for effective detection of recycled ICs," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 22, no. 5, pp. 1016–1029, 2014.

[69] C. Jin and M. van Dijk, "Secure and efficient initialization and authentication protocols for SHIELD," *IEEE Transactions on Dependable and Secure Computing*, vol. 16, no. 1, pp. 156–173, 2019.

[70] X. Dong *et al.*, "Detection and identification of vehicles based on their unintended electromagnetic emissions," *IEEE Transactions on Electromagnetic Compatibility*, vol. 48, no. 4, pp. 752–759, 2006.

[71] H. Göksu, D. C. Wunsch, X. Dong, A. Kökce, and D. G. Beetner, "Detection and identification of vehicles based on their spark-free unintended electromagnetic emissions," *IEEE Transactions on Electromagnetic Compatibility*, vol. 60, no. 5, pp. 1594–1597, 2018.

[72] C. Herder, M. Yu, F. Koushanfar, and S. Devadas, "Physical unclonable functions and applications: A tutorial," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1126–1141, Aug. 2014.

[73] J. M. Vann, T. P. Karnowski, R. Kerekes, C. D. Cooke, and A. L. Anderson, "A dimensionally aligned signal projection for classification of unintended radiated emissions," *IEEE Transactions on Electromagnetic Compatibility*, vol. 60, no. 1, pp. 122–131, 2017.

[74] J. M. Vann, T. Karnowski, and A. L. Anderson, "Classification of unintended radiated emissions in a multi-device environment," *IEEE Transactions on Smart Grid*, vol. 10, no. 5, pp. 5506–5513, 2018.

[75] C. Yang and A. P. Sample, "EM-ID: Tag-less identification of electrical devices via electromagnetic emissions," in *2016 IEEE International Conference on RFID (RFID)*, IEEE, 2016, pp. 1–8.

[76] G. Laput, C. Yang, R. Xiao, A. Sample, and C. Harrison, "EM-Sense: Touch recognition of uninstrumented, electrical and electromechanical objects," in *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, 2015, pp. 157–166.

[77] M. M. Ahmed *et al.*, "Authentication of microcontroller board using non-invasive EM emission technique," in *2018 IEEE 3rd International Verification and Security Workshop (IVSW)*, IEEE, 2018, pp. 25–30.

[78] M. M. Ahmed *et al.*, "Radiated electromagnetic emission for integrated circuit authentication," *IEEE Microwave and Wireless Components Letters*, vol. 27, no. 11, pp. 1028–1030, Nov. 2017.

[79] X. Zhang, K. Xiao, and M. Tehranipoor, "Path-delay fingerprinting for identification of recovered ICs," in *2012 IEEE International symposium on defect and fault tolerance in VLSI and nanotechnology systems (DFT)*, IEEE, 2012, pp. 13–18.

[80] K. Huang, Y. Liu, N. Korolija, J. M. Carulli, and Y. Makris, "Recycled IC detection based on statistical methods," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 6, pp. 947–960, 2015.

[81] A. Vakil, F. Niknia, A. Mirzaeian, A. Sasan, and N. Karimi, "Learning assisted side channel delay test for detection of recycled ICs," in *Proceedings of the 26th Asia and South Pacific Design Automation Conference*, 2021, pp. 455–462.

[82] F. Ahmed, M. Shintani, and M. Inoue, "Accurate recycled FPGA detection using an exhaustive-fingerprinting technique assisted by WID process variation modeling," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020.

[83] T. Alnuayri, S. Khursheed, A. L. H. Martinez, and D. Rossi, "Differential aging sensor using subthreshold leakage current to detect recycled ICs," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2021.

[84] U. Guin, X. Zhang, D. Forte, and M. Tehranipoor, "Low-cost on-chip structures for combating die and IC recycling," in *2014 51st ACM/EDAC/IEEE Design Automation Conference (DAC)*, IEEE, 2014, pp. 1–6.

[85] A. Amouri and M. Tahoori, "A low-cost sensor for aging and late transitions detection in modern FPGAs," in *2011 21st International Conference on Field Programmable Logic and Applications*, IEEE, 2011, pp. 329–335.

[86] C. Leong *et al.*, "Aging monitoring with local sensors in FPGA-based designs," in *2013 23rd International Conference on Field programmable Logic and Applications*, IEEE, 2013, pp. 1–4.

[87] M. Sadi, L. Winemberg, and M. Tehranipoor, "A robust digital sensor IP and sensor insertion flow for in-situ path timing slack monitoring in SoCs," in *2015 IEEE 33rd VLSI Test Symposium (VTS)*, IEEE, 2015, pp. 1–6.

[88] M. Alam, S. Chowdhury, M. M. Tehranipoor, and U. Guin, "Robust, low-cost, and accurate detection of recycled ICs using digital signatures," in *2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, IEEE, 2018, pp. 209–214.

[89] H. Dogan, D. Forte, and M. M. Tehranipoor, "Aging analysis for recycled FPGA detection," in *2014 IEEE international symposium on defect and fault tolerance in VLSI and nanotechnology systems (DFT)*, IEEE, 2014, pp. 171–176.

[90] M. M. Alam, M. Tehranipoor, and D. Forte, "Recycled FPGA detection using exhaustive LUT path delay characterization," in *2016 IEEE International test conference (ITC)*, IEEE, 2016, pp. 1–10.

[91] M. M. Alam, M. Tehranipoor, and D. Forte, "Recycled FPGA detection using exhaustive LUT path delay characterization and voltage scaling," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 12, pp. 2897–2910, 2019.

[92] E. A. Stott, J. S. Wong, P. Sedcole, and P. Y. Cheung, "Degradation in FPGAs: Measurement and modelling," in *Proceedings of the 18th annual ACM/SIGDA international symposium on Field programmable gate arrays*, 2010, pp. 229–238.

[93] X. Li, J. Qin, and J. B. Bernstein, "Compact modeling of MOSFET wearout mechanisms for circuit-reliability simulation," *IEEE Transactions on Device and Materials Reliability*, vol. 8, no. 1, pp. 98–121, 2008.

[94] B. Tudor, J. Wang, Z. Chen, R. Tan, W. Liu, and F. Lee, "An accurate MOSFET aging model for 28 nm integrated circuit simulation," *Microelectronics Reliability*, vol. 52, no. 8, pp. 1565–1570, 2012.

[95] A. Amouri and M. Tahoori, "High-level aging estimation for FPGA-mapped designs," in *22nd International Conference on Field Programmable Logic and Applications (FPL)*, IEEE, 2012, pp. 284–291.

[96] R. Vattikonda, W. Wang, and Y. Cao, "Modeling and minimization of PMOS NBTI effect for robust nanometer design," in *2006 43rd ACM/IEEE Design Automation Conference*, IEEE, 2006, pp. 1047–1052.

[97] S. Kiamehr, A. Amouri, and M. B. Tahoori, "Investigation of NBTI and PBTI induced aging in different LUT implementations," in *2011 International Conference on Field-Programmable Technology*, IEEE, 2011, pp. 1–8.

[98] S. Deora *et al.*, "Positive bias instability and recovery in InGaAs channel nMOSFETs," *IEEE Transactions on Device and Materials Reliability*, vol. 13, no. 4, pp. 507–514, 2013.

[99] S. Mukhopadhyay, N. Goel, and S. Mahapatra, "A comparative study of NBTI and PBTI using different experimental techniques," *IEEE Transactions on Electron Devices*, vol. 63, no. 10, pp. 4038–4045, 2016.

[100] S. Zafar *et al.*, "A comparative study of NBTI and PBTI (charge trapping) in SiO2/HfO2 stacks with FUSI, TiN, Re gates," in *2006 Symposium on VLSI Technology, 2006. Digest of Technical Papers.*, 2006, pp. 23–25.

[101] A. Kerber and E. A. Cartier, "Reliability challenges for CMOS technology qualifications with hafnium oxide/titanium nitride gate stacks," *IEEE Transactions on Device and Materials Reliability*, vol. 9, no. 2, pp. 147–162, 2009.

[102] M. Prvulovic, A. Zajic, R. L. Callan, and C. J. Wang, "A method for finding frequency-modulated and amplitude-modulated electromagnetic emanations in computer systems," *IEEE Transactions on Electromagnetic Compatibility*, vol. 59, no. 1, pp. 34–42, Feb. 2017.

[103] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The computer journal*, vol. 7, no. 4, pp. 308–313, 1965.

[104] J. Robinson and Y. Rahmat-Samii, "Particle swarm optimization in electromagnetics," *IEEE Transactions on Antennas and Propagation*, vol. 52, no. 2, pp. 397–407, Feb. 2004.

[105] A. R. Djordjević, M. B. Bazdar, V. V. Petrović, D. I. Olćan, T. K. Sarkar, and R. F. Harrington, *AWAS for Windows: Analysis of Wire Antennas and Scatterers, Software and User's Manual*. Boston:Artech House, 2002.

[106] H. Shimada *et al.*, "Efficient mapping of EM radiation associated with information leakage for cryptographic devices," in *2012 IEEE International Symposium on Electromagnetic Compatibility*, IEEE, 2012, pp. 794–799.

[107] V. Kresalek, M. Smola, and T. Kosina, "Scanning of electromagnetic radiation for EMC and data security purposes," in *2008 42nd Annual IEEE International Carnahan Conference on Security Technology*, IEEE, 2008, pp. 117–120.

[108] L. Sauvage, S. Guilley, J.-L. Danger, N. Homma, and Y.-i. Hayashi, "Practical results of EM cartography on a FPGA-based RSA hardware implementation," in *2011 IEEE International Symposium on Electromagnetic Compatibility*, IEEE, 2011, pp. 768–772.

[109] "C-RAM SFC 4," Cuming Microwave. (2015), [Online]. Available: http://stores.cumingmicrowave-online-store.com/c-ram-sfc-4-4-3-x-24-x-24/.

[110] J. Fan, "Near-field scanning for EM emission characterization," *IEEE Electromagnetic Compatibility Magazine*, vol. 4, no. 3, pp. 67–73, 2015.

[111] "USRP B210," Ettus Research. (2019), [Online]. Available: https://www.ettus.com/all-products/ub210-kit/.

[112] F. T. Werner, A. R. Djordjević, and A. G. Zajić, "A compact probe for EM side-channel attacks on cryptographic systems," in *2019 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting*, IEEE, 2019, pp. 613–614.

[113] Y. Chou and H. Lu, "Magnetic near-field probes with high-pass and notch filters for electric field suppression," *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, no. 6, pp. 2460–2470, Jun. 2013.

[114] K. Solbach, "Compensation of electric cross-field response in shielded loop probe," *Electronics letters*, vol. 47, no. 2, pp. 95–97, 2011.

[115] H. Whiteside and R. King, "The loop antenna as a probe," *IEEE Transactions on Antennas and Propagation*, vol. 12, no. 3, pp. 291–297, 1964.

[116] X. Tong, D. W. P. Thomas, A. Nothofer, P. Sewell, and C. Christopoulos, "Modeling electromagnetic emissions from printed circuit boards in closed environments using equivalent dipoles," *IEEE Transactions on Electromagnetic Compatibility*, vol. 52, no. 2, pp. 462–470, 2010.

[117] Y. Vives-Gilabert, C. Arcambal, A. Louis, F. de Daran, P. Eudeline, and B. Mazari, "Modeling magnetic radiations of electronic circuits using near-field scanning method," *IEEE Transactions on Electromagnetic Compatibility*, vol. 49, no. 2, pp. 391–400, 2007.

[118] Q. Huang, F. Zhang, T. Enomoto, J. Maeshima, K. Araki, and C. Hwang, "Physics-based dipole moment source reconstruction for RFI on a practical cellphone," *IEEE Transactions on Electromagnetic Compatibility*, vol. 59, no. 6, pp. 1693–1700, 2017.

[119] "SMA four (4) hole panel receptacle," Amphenol Connex. (2016), [Online]. Available: https://www.amphenolrf.com/library/download/link/link_id/307286/parent/132146/.

[120] "HFSS," Ansys. (2016), [Online]. Available: https://www.ansys.com/products/electronics/ansys-hfss.

[121] "DE1 development and education board," Altera Coporation. (2013), [Online]. Available: https://www.intel.com/content/dam/altera-www/global/en_US/portal/dsn/42/doc-us-dsnbk-42-4904342209-de1-usermanual.pdf.

[122] "A13-olinuxino," Olimex. (2013), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A13/A13-OLinuXino/open-source-hardware.

[123] N. Neshenko, E. Bou-Harb, J. Crichigno, G. Kaddoum, and N. Ghani, "Demystifying IoT security: An exhaustive survey on IoT vulnerabilities and a first empirical look on internet-scale IoT exploitations," *IEEE Communications Surveys Tutorials*, vol. 21, no. 3, pp. 2702–2733, 2019.

[124] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.

[125] C. Wang, R. Callan, A. Zajic, and M. Prvulovic, "An algorithm for finding carriers of amplitude-modulated electromagnetic emanations in computer systems," in *2016 10th European Conference on Antennas and Propagation (EuCAP)*, IEEE, 2016, pp. 1–5.

[126] R. Callan, N. Popovic, A. Daruna, E. Pollmann, A. Zajic, and M. Prvulovic, "Comparison of electromagnetic side-channel energy available to the attacker from different computer systems," in *2015 IEEE International Symposium on Electromagnetic Compatibility (EMC)*, IEEE, 2015, pp. 219–223.

[127] B. B. Yilmaz, E. M. Ugurlu, M. Prvulovic, and A. Zajic, "Detecting cellphone camera status at distance by exploiting electromagnetic emanations," in *MILCOM 2019-2019 IEEE Military Communications Conference (MILCOM)*, IEEE, 2019, pp. 1–6.

[128] "EM probe station," Riscure. (2018), [Online]. Available: https://getquote.riscure.com/en/quote/2101064/em-probe-station.htm.

[129] C. Beleites *et al.*, "Variance reduction in estimating classification error using sparse datasets," *Chemometrics and Intelligent Laboratory Systems*, vol. 79, no. 1-2, pp. 91–100, 2005.

[130] M. Dey, A. Nazari, A. Zajić, and M. Prvulovic, "EMPROF: Memory profiling via EM-emanation in IoT and hand-held devices," in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, IEEE, 2018, pp. 881–893.

[131] "A10-olinuxino-lime," Olimex. (2019), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A10/A10-OLinuXino-LIME-n4GB/open-source-hardware.

[132] "A13-olinuxino-micro," Olimex. (2013), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A13/A13-OLinuXino-MICRO/open-source-hardware.

[133] "A20-olinuxino-lime," Olimex. (2015), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A20/A20-OLinuXino-LIME/open-source-hardware.

[134] "A20-olinuxino-lime2," Olimex. (2015), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A20/A20-OLinuXino-LIME2/open-source-hardware.

[135] "A20-olinuxino-micro," Olimex. (2013), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A20/A20-OLinuXino-MICRO/open-source-hardware.

[136] "A33-olinuxino," Olimex. (2017), [Online]. Available: https://www.olimex.com/Products/OLinuXino/A33/A33-OLinuXino/open-source-hardware.

[137] "Olinuxino is open source," Olimex. (2012), [Online]. Available: https://github.com/OLIMEX/OLINUXINO.

[138] "DE0-CV cyclone V board," Zentel Electronics Corp. (2017), [Online]. Available: https://www.intel.com/content/www/us/en/programmable/solutions/partners/partner-profile/terasic-inc-/board/de0-cv-cyclone-v-board.html.

[139] "Altera DE1 board," Intel Corporation. (2017), [Online]. Available: https://www.intel.com/content/www/us/en/programmable/solutions/partners/partner-profile/terasic-inc-/board/altera-de1-board.html.

[140] "K4B4G1646E-BYK0," Samsung. (2018), [Online]. Available: https://www.samsung.com/semiconductor/dram/ddr3/K4B4G1646E-BYK0/.

[141] "H5TQ2G83FFR," Hynix Semiconductor. (2014), [Online]. Available: https://www.skhynix.com/eolproducts.view.do?pronm=DDR3+SDRAM&srnm=H5TQ2G83FFR&rk=19&rc=consumer.

[142] "H5TQ2G63BFR," Hynix Semiconductor. (2010), [Online]. Available: https://www.skhynix.com/eolproducts.view.do?pronm=DDR3+SDRAM&srnm=H5TQ2G63BFR&rk=19&rc=computing.

[143] "MT41K256M16HA-125 XIT," Micron. (2019), [Online]. Available: https://www.micron.com/products/dram/ddr3-sdram/part-catalog/mt41k256m16ha-125-xit.

[144] "IS42/45R86400D/16320D/32160D IS42/45S86400D/16320D/32160D," Integrated Silicon Solution Inc. (2015), [Online]. Available: http://www.issi.com/WW/pdf/42-45R-S_86400D-16320D-32160D.pdf.

[145] "A3V64S40GTP/GBF," Zentel Electronics Corporation. (2017), [Online]. Available: https://zentel-europe.com/datasheets/A3V64S40GTP_v1.3_Zentel.pdf.

[146] "Allwinner A10," Allwinner Technology. (2014), [Online]. Available: https://web.archive.org/web/20160729114202/http://www.allwinnertech.com/en/clq/processora/A10.html.

[147] "Allwinner A13," Allwinner Technology. (2014), [Online]. Available: https://web.archive.org/web/20160811122523/http://www.allwinnertech.com/en/clq/processora/A13.html.

[148] "Allwinner A20," Allwinner Technology. (2016), [Online]. Available: http://www.allwinnertech.com/index.php?c=product&a=index&id=45.

[149] "Allwinner A33," Allwinner Technology. (2016), [Online]. Available: http://www.allwinnertech.com/index.php?c=product&a=index&id=23.

[150] "RTL8201CP," Realtek. (2003), [Online]. Available: http://realtek.info/pdf/rtl8201cp.pdf.

[151] "LAN8710A," Microchip Technology. (2016), [Online]. Available: https://www.microchip.com/wwwproducts/en/LAN8710A.

[152] "KSZ9031," Microchip Technology. (2016), [Online]. Available: https://www.microchip.com/wwwproducts/en/KSZ9031.

[153] J. H. Anderson, "A PUF design for secure FPGA-based embedded systems," in *2010 15th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2010, pp. 1–6.

[154] "Cortex-a8," ARM Developer. (2019), [Online]. Available: https://developer.arm.com/ip-products/processors/cortex-a/cortex-a8.

[155] "Cortex-a7," ARM Developer. (2019), [Online]. Available: https://developer.arm.com/ip-products/processors/cortex-a/cortex-a7.

[156] Y. Zhou and D. Feng, "Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing," *Cryptology ePrint Archive*, 2005.

[157] C. Lavaud, R. Gerzaguet, M. Gautier, O. Berder, E. Nogues, and S. Molton, "Whispering devices: A survey on how side-channels lead to compromised information," *Journal of Hardware and Systems Security*, vol. 5, no. 2, pp. 143–168, 2021.

[158] C. R. Paul, *Introduction to Electromagnetic Compatibility*. John Wiley & Sons, Inc, 2006.

[159] "Predictive technology model (PTM)," Nanoscale Integration and Modeling (NIMO) Group, ASU. (2011), [Online]. Available: http://ptm.asu.edu/.

[160] S. Dixon-Warren. "A review of TSMC 28 nm process technology," TechInsights. (2011), [Online]. Available: https://www.techinsights.com/blog/review-tsmc-28-nm-process-technology.

[161] "X-LSQ series: High-speed motorized linear stages with built-in controllers," Zaber Technologies. (2021), [Online]. Available: https://www.zaber.com/products/linear-stages/X-LSQ/specs?part=X-LSQ150B.

[162] "RF & near-field probes," Aaronia AG. (2016), [Online]. Available: https://aaronia.com/antennas/rf-and-near-field-probes/.

[163] Y. Chen *et al.*, "Negative bias temperature instability (NBTI) in deep sub-micron p/sup+/-gate pMOSFETs," in *2000 IEEE International Integrated Reliability Workshop Final Report (Cat. No. 00TH8515)*, IEEE, 2000, pp. 98–101.

[164] H. Salmani, M. Tehranipoor, and R. Karri, "On design vulnerability analysis and trust benchmarks development," in *2013 IEEE 31st international conference on computer design (ICCD)*, IEEE, 2013, pp. 471–474.

[165] B. Shakya, T. He, H. Salmani, D. Forte, S. Bhunia, and M. Tehranipoor, "Benchmarking of hardware trojans and maliciously affected circuits," *Journal of Hardware and Systems Security*, vol. 1, no. 1, pp. 85–102, 2017.

[166] L. N. Nguyen, B. B. Yilmaz, M. Prvulovic, and A. Zajic, "A novel golden-chip-free clustering technique using backscattering side channel for hardware trojan detection," in *2020 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, IEEE, 2020, pp. 1–12.

VITA

Frank Thompson Werner received his B.S. degree (2013) and his M.S. (2016) in electrical engineering from Auburn University, Alabama. During his time at Auburn, he worked on electromagnetic modeling, power systems, and chaos electronics.

Currently, he is completing his Ph.D. in electrical engineering at the Georgia Institute of Technology. At Georgia Tech, he has worked on studying and leveraging side channels on various devices. His research interests include electromagnetic compatibility, hardware security, component authentication, signal processing, and applied electromagnetics.