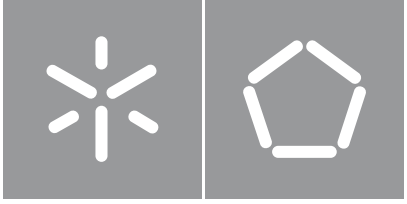
Universidade do Minho
Escola de Engenharia

Luís Miguel Marinho Novais

Electromagnetic Vulnerability Discovery and Mitigation in Safety-Critical CPS

julho de 2026

Electromagnetic Vulnerability Discovery and Mitigation in Safety-Critical CPS
Luís Miguel Marinho Novais
UMinho | 2026



Universidade do Minho
Escola de Engenharia

Luís Miguel Marinho Novais

Electromagnetic Vulnerability Discovery and Mitigation in Safety-Critical CPS

Tese de Doutoramento
Programa Doutoral Engenharia
Eletrónica e de Computadores

Trabalho efetuado sob a orientação do
Professor Doutor Jorge Miguel Nunes dos Santos Cabral
Doutor Sandro Emanuel Salgado Pinto

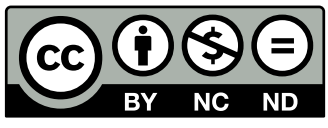
julho de 2026

DIREITOS DE AUTOR E CONDIÇÕES DE UTILIZAÇÃO DO TRABALHO POR TERCEIROS

Este é um trabalho académico que pode ser utilizado por terceiros desde que respeitadas as regras e boas práticas internacionalmente aceites, no que concerne aos direitos de autor e direitos conexos. Assim, o presente trabalho pode ser utilizado nos termos previstos na licença abaixo indicada.

Caso o utilizador necessite de permissão para poder fazer um uso do trabalho em condições não previstas no licenciamento indicado, deverá contactar o autor, através do RepositóriUM da Universidade do Minho.

Licença concedida aos utilizadores deste trabalho



Atribuição-NãoComercial-SemDerivações
CC BY-NC-ND

https://creativecommons.org/licenses/by-nc-nd/4.0/

Acknowledgements

During the past years, I had the opportunity to grow both personally and professionally while developing the work presented in this thesis. This journey was far from easy, and reaching this point would not have been possible without the support of several remarkable people, to whom I am deeply grateful.

I would like to begin by expressing my sincere gratitude to my supervisor, Dr. Jorge Cabral, for his guidance, support, and trust throughout this journey. His experience and perspective were fundamental to the completion of this work. I also thank my co-supervisor, Dr. Sandro Pinto, for his availability, insight, and valuable contributions.

A very special acknowledgment goes to Nelson Naia, who accompanied me throughout this long journey and was truly the most important person in helping me reach this milestone and publish this thesis. Your support, presence, and encouragement made a decisive difference. I would also like to thank Pedro Lages, who joined the laboratory as a student and steadily earned his place through dedication and hard work. It is extremely rewarding to see him now with strong prospects of pursuing a PhD, building upon the research lines initiated in this work. Last but not least, my sincere thanks go to Paulo Santos, a companion of more than a decade, a true friend, and one of the most outstanding professionals I have had the privilege to work with.

To my parents and brothers, thank you for your unconditional support, patience, and for always standing by me.

To Carina Faria, thank you for your patience, your strength, and for standing by me throughout these demanding years. I hope this achievement reflects, in part, everything you have given along the way.

And finally, to my son, Matias Novais, who arrived before his father could submit this thesis — you are my greatest achievement, by far.

Luís Miguel Marinho Novais

April 17th, 2026

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the University of Minho.

Resumo

Os *cyber-physical systems* (CPS) utilizados em aplicações críticas dependem de *front-ends* analógicos de sensores para assegurar a interface com o mundo físico. Ainda assim, estes *front-ends* analógicos permanecem entre os elementos mais expostos e vulneráveis da arquitetura, e, embora as práticas existentes de *safety*, *security* e de garantia eletromagnética sejam essencias, estas continuam a oferecer uma visão fragmentada de como perturbações ao nível do *hardware* se propagam e originam efeitos funcionais indesejados. Embora o *Risk-Based Electromagnetic Compatibility* (RB-EMC) tenha reforçado o tratamento da *Electromagnetic Interference* (EMI) como risco funcional, a descoberta de *Electromagnetic Vulnerabilities* (EMV) em *front-ends* analógicos continua a ocorrer de forma essencialmente *ad hoc*.

A tese introduz a *Architecture-Aware Electromagnetic Methodology* (A^2EM), que estende o RB-EMC ao relacionar a exposição eletromagnética com a arquitetura e os mecanismos de propagação, permitindo a descoberta sistemática de vulnerabilidades. Para tal, é desenvolvida uma estratégia de modelação que preserva mecanismos de acoplamento fisicamente significativos, recorrendo a um modelo eletromagnético ao nível das pistas baseado no *Method of Moments* (MoM) para fios finos, tratando as interligações de PCB como pontos de entrada explícitos. As respostas induzidas são reformuladas como fontes equivalentes e injetadas em modelos arquiteturais, permitindo analisar a propagação de perturbações ao longo da cadeia de sinal. A metodologia é demonstrada num caso de estudo de um sensor automóvel, evidenciando a sua capacidade para identificar vulnerabilidades e relacioná-las com desvios funcionais à saída do sistema.

Globalmente, a tese propõe um processo sistemático e fisicamente fundamentado para a descoberta de EMV em CPS, suportando uma avaliação mais precoce e informada do risco eletromagnético em *front-ends* analógicos de sensores em aplicações críticas.

Palavras-chave: Architecture-Aware Electromagnetic Methodology (A^2EM), Cyber-physical Systems (CPS), Electromagnetic Interference (EMI), Electromagnetic Vulnerabilities(EMV), Front-ends Analógicos de Sensores

Abstract

Safety-Critical Cyber-Physical Systems (CPS) increasingly depend on tightly constrained analog sensor front-ends to interface computation with the physical world. Yet these front-ends remain among the most expose and fragile parts of the system architecture, while existing safety, security, and electromagnetic assurance practices are essential, they still offer a fragmented account of how hardware-level disturbances couple into sensor architectures, propagate through signal chains, and ultimately produce unsafe functional effects. In particular, although Risk-Based Electromagnetic Compatibility (RB-EMC) has strengthened the treatment of Electromagnetic Interference (EMI) as a functional hazard, the discovery of Electromagnetic Vulnerabilities (EMV) in analog front-ends remains largely ad hoc.

This thesis introduces the Architecture-Aware Electromagnetic Methodology (A^2EM), which extends RB-EMC by linking electromagnetic exposure to the system architecture and propagation mechanisms, enabling systematic vulnerability discovery at the system level. A modelling strategy is developed to preserve physically meaningful coupling mechanisms while remaining compatible with system-level workflows. A trace-level electromagnetic model based on a thin-wire Method of Moments (MoM) formulation is proposed, allowing PCB interconnects to be treated as explicit entry points. Induced responses are reformulated as equivalent sources and injected into architectural models, enabling the analysis of disturbance propagation across the signal chain. The methodology is demonstrated through a representative automotive sensor case study, showing its ability to identify vulnerable coupling paths and relate them to functional deviations at the system output.

Overall, this work provides a systematic and physically grounded framework for EMV discovery in CPS, supporting earlier and more informative assessment of electromagnetic risk in analog safety-critical sensor front-ends.

Keywords: Analog Sensor Front-Ends, Architecture-Aware Electromagnetic Methodology (A^2EM), Cyber-Physical Systems (CPS), Electromagnetic Interference (EMI), Electromagnetic Vulnerability (EMV)

Table of Contents

Acknowledgements iii

Resumo v

Abstract vi

List of Figures xi

List of Tables xiii

List of Algorithms xiv

List of Acronyms xv

1 Introduction 1

1.1 Motivation . . . 2

1.2 Research Objectives . . . 3

1.3 Associated Projects and Institutional Context . . . 4

1.4 Thesis Structure . . . 6

2 Background 8

2.1 Embedded Systems . . . 8

2.2 Cyber-Physical Systems . . . 10

2.3 The Internet of Things . . . 11

2.4 CPS vs. IoT . . . 12

2.5 CPS Layers . . . 13

2.6 Safety-Critical CPS Landscape . . . 14

2.6.1 Analog Fault Coverage . . . 15

2.6.2 P2427 Draft Standard for Analog Defect Modeling and Coverage 16
2.6.3 Hardware Trojans 17
2.6.4 Physical-to-Cyber Attacks 20
2.7 Hardware Integrity Threats 22
2.8 Summary 24

3 A²EM - Architecture-Aware Electromagnetic Methodology 25
3.1 Introduction 25
3.2 The role of EMI in Safety-Critical CPS 26
3.2.1 Risk-Based EMC Management 28
3.2.2 Physical-to-Cyber Attack Model 31
3.3 A²EM: Architecture-Aware Electromagnetic Methodology 33
3.3.1 Overview 34
3.3.2 Resonant Entry Point Mapping 36
3.3.3 Electromagnetic Interference (EMI) Probing and Profiling 37
3.3.4 Simulation Approaches 41
3.4 Summary 45

4 Operationalizing A²EM 47
4.1 Introduction 47
4.2 Modeling Approach 48
4.2.1 Transduction Principles 49
4.2.2 Functional Signal-Chain Breakdown 50
4.2.3 Entry-Point Mapping & Clustering 51
4.2.4 Transfer Function Modeling 51
4.3 Case Study Demonstration - Coupled-Coil Position Sensor 52
4.3.1 Transduction Principles 52
4.3.2 Functional Signal-Chain Breakdown 54
4.3.3 Entry-Point Mapping & Clustering 55
4.3.4 Transfer Function Modeling 56
4.4 Summary 59

5 Trace Modeling **61**

5.1 Introduction . . . 61

5.2 The Challenge of Predicting Induced Currents in Real PCB Traces . . . 62

5.3 Solver Approaches for Trace-Level Electromagnetic Modeling . . . 64

5.4 Method of Moments for Thin Wires . . . 65

5.4.1 Thin-Wire Representation of PCB Traces . . . 68

5.4.2 Basis Functions . . . 69

5.4.3 Electromagnetic Coupling in the Discretized EFIE . . . 71

5.5 Trace Modeling Process . . . 73

5.5.1 Overview . . . 74

5.5.2 Trace Geometry Extraction . . . 75

5.5.3 Impedance Computation . . . 77

5.5.4 Induced Voltage Computation . . . 79

5.5.5 Resonant Optimization . . . 81

5.6 Case Study - Coupled-Coil Position Sensor . . . 83

5.6.1 Trace Geometry Extraction . . . 84

5.6.2 Impedance Computation . . . 84

5.6.3 Induced Voltage Computation. . . . 85

5.6.4 Resonant Optimization . . . 85

5.7 Summary . . . 87

6 A^{2}EM Results **89**

6.1 Introduction . . . 89

6.2 Black-box Testing . . . 90

6.3 Gray-box Testing . . . 92

6.4 White-box Testing . . . 96

6.5 Mitigation . . . 98

6.5.1 Mitigation options . . . 99

6.5.2 Implemented strategy: post-Analog-to-Digital Converter (ADC) digital filtering . 99

6.5.3 Validation: direct comparison against the pre-mitigation baseline . . . 100

6.6 Summary . . . 102

7 Conclusion **104**

7.1 Chapter Summary and Future Directions 106

7.2 Contribution to ESRG and Future Directions 109

References **110**

A Appendix **121**

List of Figures

2.1 Sensor signal chain highlighting potential attacking points. 20

2.2 Venn diagrams of safety and security overlaps: (a) Hardware Integrity Threats (b) Hardware Integrity Threat Countermeasures. 23

3.1 Overview of the Risk-Based EMC Management workflow. 29

3.2 Overview of A^{2}EM - Architecture-Aware Electromagnetic Methodology. 35

3.3 Illustatrion of hypothetical sensor resonant frequency bands. 37

3.4 Illustration of hypothetical sensor permeable frequency bands. 40

4.1 Main stages of the modeling process of a typical temperature sensor. 50

4.2 Experimental setup for the coupled-coil sensor. 53

4.3 In-phase and quadrature (I/Q) signals generated by the coupled-coil sensor over one full rotation. 54

4.4 Functional signal-chain breakdown of the coupled-coil sensor. 55

4.5 Entry-point-bounded clusters of the coupled-coil sensor. 56

4.6 Full Simulink model of the coupled-coil position sensor. 57

4.7 Simscape electrical representation of the HZBMatch block. 58

5.1 Incidence geometry and definition of the angles θ, ϕ, and γ. 72

5.2 Overview of the trace modeling process. 74

5.3 High-Level flowchart of trace geometry extraction. 76

5.4 High-Level flowchart of impedance computation. 78

5.5 Thévenin-equivalent circuits used to recover intrinsic terminal impedances from voltage transfer functions referenced to a characteristic impedance Z_0. 79

5.6 High-Level flowchart of induced voltage computation. 80

5.7 High-Level flowchart of resonant optimization. 82

5.8 Graphical rendering of extracted electromagnetic entry-point traces from sensor PCB Gerber data. 84
5.9 PCB stack-up used in trace-level electromagnetic modeling. 85
5.10 Worst-case induced-voltage envelopes $U_{\max}(f)$ for Traces 10 and 11 across the analyzed frequency range. The highlighted peaks indicate resonant frequencies identified by the local-maximum detection procedure. 86

6.1 HCoil output under nominal and disturbed conditions. 93
6.2 Spectral content of the bandpass output. 93
6.3 Spectral content of the HZMMatch output. 94
6.4 Spectral content of the HMixer output. 95
6.5 Spectral content of the HAmpLPF output. 95
6.6 Spectral content of the ADC input. 96

List of Tables

5.1 Absolute maximum induced-voltage resonance identified for each electromagnetic entry-point trace. 87

A.1 Polarisation angles associated with frequency bands presenting non-negligible deviation ($|\Delta\theta| \geq 1^\circ$) in the black-box analysis. Each row corresponds to an identified local maximum exceeding the deviation threshold. 121

A.2 Polarisation angles associated with excitation cases retained from the pre-mitigation black-box permeability table, now reporting the residual angular deviation after implementing the post-ADC digital low-pass filter. Each row corresponds to the same identified local maximum (frequency/polarisation) as in the pre-mitigation table, enabling direct one-to-one comparison, including cases whose post-mitigation deviation falls below the $|\Delta\theta| \geq 1^\circ$ screening threshold. 124

List of Algorithms

1 Black-Box Electromagnetic Vulnerability (EMV) Discovery Procedure 42

List of Acronyms

ADC Analog-to-Digital Converter.

AM Amplitude Modulation.

AMS Analog and Mixed-Signal.

ATPG Automatic Test Pattern Generation.

CAGR Compound Annual Growth Rate.

CPS Cyber-Physical System.

DfT Design for Trust.

DO Data Object.

DoS Denial-of-Service.

DPPM Defective Parts per Million.

DSB-SC Double-Sideband Suppressed-Carrier.

E/E Electrical and Electronic.

EFIE Electric Field Integral Equation.

EM Electromagnetic.

EMC Electromagnetic Compatibility.

EMCD European EMC Directive.

EMD Electromagnetic Disturbance.

EME Electromagnetic Environment.

EMI Electromagnetic Interference.

EMV Electromagnetic Vulnerability.

HiL Hardware-in-the-Loop.

HIT Hardware Integrity Threat.

IC Integrated Circuit.

ICT Information and Communication Technology.

IEME Intended Electromagnetic Environment.

IIR Infinite Impulse Response.

IMD Intermodulation Distortion.

IO Interactive Object.

IoT Internet of Things.

IP Intellectual Property.

IT Information Technology.

LED Light-Emitting Diode.

LIDAR Light Detection and Ranging.

MoM Method of Moments.

PCB Printed Circuit Board.

PEC Perfect Electric Conductor.

RB-EMC Risk-Based EMC Management.

RF Radio Frequency.

RFID Radio-Frequency Identification.

RTL Register Transfer Level.

SO Smart Object.

SoC System-on-Chip.

SWaP-C Size, Weight, Power, and Cost.

TO Trackable Object.

WSN Wireless Sensor Network.

Chapter 1

Introduction

Cyber-Physical Systems (CPS) stand as a backbone of the forthcoming technological revolution, where the integration of analog and digital circuitry enables systems to interact seamlessly with the physical world [1]. Being more and more employed in all kinds of systems, CPS enable these systems to be autonomous and independent of human intervention, or in other words, smart. With an estimated Compound Annual Growth Rate (CAGR) from 2021 to 2031 of 10.6%, which represents a forecasted value of about 23 billion US dollars by the end of 2031 [2], the CPS market value is steadily growing. This expansion is driven by a wide range of application domains, including Industry 4.0, smart grids, wearable healthcare devices, and autonomous driving, in which cyber-physical integration is a key enabler of autonomy, efficiency, and adaptive behavior.

As a matter of fact, the automotive segment's introduction of Electrical and Electronic (E/E) components in previously purely mechanical systems is a primary example of a CPS-driven market, evidenced by automotive by-Wire systems being estimated to have a CAGR from 2022 to 2027 of 16% that represents a valuation of 81.993 million US dollars by 2027 [3]. This integration of E/E components into safety-critical systems, where any malfunction might lead to potential human harm, raises certain design considerations. In this context, the electronics within these systems must consistently function without error, even in the presence of circuit faults.

Fault-tolerant electronic systems have been successfully employed in the past in areas like aerospace, aviation and military, but the attempt to massively cyberize presents new challenges. In emerging areas where new solutions are required, the expertise in fault-tolerant engineering isn't as prevalent, and even in there, existing fault-tolerant solutions might not be directly applicable to these novel scenarios [4–6]. As such, and with such intense market demand, is the engineering world ready to rapidly answer these

questions? What is the technological maturity of safety-critical CPS development, considering the Size, Weight, Power, and Cost (SWaP-C) metrics in modern industries?

1.1 Motivation

It seems despite all current efforts and concerns with functional safety metrics and overall system reliability, there are still malfunctioning products that become available in the market, often with catastrophic consequences. Famous cases are the Boeing 737 Max crashes in 2018 and 2019 which resulted in a tragic total of 348 deaths [7], and the more recent involvement of Tesla's Full Self-Driving (Beta) technology in crashes that resulted in at least 53 verified deaths [8].

CPS rely on analog electronics (front-ends) to translate the physical world into internal signals for decision-making and actuation. Industry and academic analyses consistently indicate that analog and mixed-signal front-ends account for a disproportionately large share of field failures compared to digital logic. In automotive Systems-on-Chip (SoC) in particular, Analog and Mixed-Signal (AMS) circuitry has been shown to exhibit higher vulnerability to parametric variation, noise, and environmental stress than its digital counterparts [9, 10].

Given this intrinsic sensitivity, CPS analog front-ends are particularly vulnerable to disturbances introduced at any point along the analog signal path by their operating environment. Among these disturbances, Electromagnetic Interference (EMI) stands out as an especially critical threat: interference tends to be particularly problematic for analog circuits and it can perturb signal amplitudes and operating points directly, often producing effects that can bypass the abstraction layers and fault-detection mechanisms that exist on the digital side. EMI has long been recognized as a challenge in electronic systems, dating back to early analog radio, and its relevance has only increased as CPS have become deeply embedded in safety-critical civil domains. This is reflected in standardization initiatives such as the IET's 2017 safety-critical EMI recommendations [11], and the IEEE 1848 standard [12], both of which explicitly address EMI resilience in scenarios where malfunction may lead to severe consequences, including risk to human life. Despite these standards and established design practices, instances still occur where safe operational behavior is not assured once systems are exposed to messy, dynamic electromagnetic environments or previously untested use-case scenarios [13–15].

This mismatch between compliance and actual electromagnetic robustness has been emphasized in recent work, including the lifecycle EMI-risk methodology introduced by Gkatsi et al. [16]. Their framework reinforces the ongoing shift from rule-based compatibility testing toward risk-based evaluation, in which EMI is treated as a functional exposure problem: specific electromagnetic environments may interact with specific system vulnerabilities to produce hazardous outcomes.

However, while Risk-Based EMC Management (RB-EMC) provides a system-level perspective, the concrete mechanisms by which realistic EMI conditions shape the behavior of analog sensor front-ends remain far less explored. In practice, Electromagnetic Vulnerability (EMV) remains largely ad hoc, as it is guided by known patterns, targeted bench work, and engineering intuition. This limitation motivates the development of new systematic, analysis-based methods.

These considerations frame the central challenge addressed in this thesis: understanding how electromagnetic disturbances interact with the intrinsic fragilities of analog sensor front-ends, and developing a systematic methodology to identify, characterize, and reason about the resulting unsafe behaviors in safety-critical CPS.

1.2 Research Objectives

The research work carried out in this thesis is driven by the need to strengthen the foundations of safety-critical CPS design at a time when these systems increasingly rely on complex sensing chains operating under tight SWaP-C constraints and in unpredictable electromagnetic environments. As the automotive and broader civil-safety sectors push towards higher levels of autonomy and electronic integration, the absence of systematic, circuit-level approaches for understanding electromagnetic exposure and its consequences on analog front-ends leaves a critical gap in current safety practices.

Addressing this gap requires engaging with multiple domains, namely functional safety, hardware security, and electromagnetic compatibility, because no single discipline yet provides a complete account of how vulnerabilities emerge in analog sensor chains or how they translate into system-level risk. In line with the core principles of RB-EMC, this thesis seeks not only to draw from that perspective, but also to operationalize it at the circuit and architectural level, where EMI-induced behaviors originate. The research is therefore structured around the need to consolidate insights from adjacent domains, derive actionable criteria for exposure and susceptibility, and demonstrate how such a methodology can be applied in practice.

Within this context, the work is organised around three research objectives:

1. **Cross-domain analysis of front-end vulnerabilities.** Investigate adjacent fields, including functional safety, hardware security, and electromagnetic compatibility, to identify how different research domains describe, classify, and analyse vulnerabilities in analog front-ends. This objective aims to extract the underlying conceptual and mechanistic commonalities across these disciplines.
2. **Development of a vulnerability-identification methodology.** Based on the insights obtained in Objective 1, derive a structured methodology for the identification of EMV in analog front-ends. The methodology is designed to complement the system-level RB-EMC perspective by introducing circuit-level, architecture-aware criteria for exposure and susceptibility analysis.
3. **Demonstration and evaluation using a case study.** Apply the proposed methodology to a representative analog sensor front-end in a safety-critical context. This objective includes evaluating the methodology's effectiveness in revealing EMV and, where possible, assessing the impact of architecture-driven mitigation strategies.

1.3 Associated Projects and Institutional Context

The research presented in this thesis was carried out within the broader activities of the University of Minho's Embedded Systems Research Group (ESRG). Although the thesis is not formally attached to a single R&D project, its development was deeply influenced by the technical, methodological, and infrastructural environment sustained by the group's ongoing research initiatives.

In practical terms, the continuity of R&D projects within ESRG played a decisive role in maintaining the laboratory ecosystem in which this R&D took shape. These projects enabled access to specialised equipment, experimental platforms, validation frameworks, and multidisciplinary discussions that would not have been available otherwise. As such, while the thesis stands as an independent research contribution, it is nevertheless shaped by the sustained research capacity built through these initiatives.

The projects most directly relevant to the technical and organisational context of this thesis are:

- **Sensible Car (POCI-01-0247-FEDER-037902)**— A large-scale R&D programme co-funded by FEDER through Portugal 2020 / COMPETE 2020, with a total eligible budget of € 33.780.401,74 and €21.073.806,14 in EU support. Led by Bosch Car Multimedia Portugal in collaboration with

the University of Minho, the project ran from 2018 to 2021 and targeted the development of next-generation intelligent sensing technologies for autonomous vehicles. Its scope covered Light Detection and Ranging (LIDAR) perception systems, road-condition sensors, high-precision inductive position sensors, steering-angle systems, and DC-motor control sensors, structured across four sub-projects and seventeen R&D lines.

In the context of the author's active participation in this project, the work conducted on the steering-angle sensor became a pivotal starting point. This subsystem served as the basis for the lab's early investigations into redundant sensing architectures and safety-oriented design methodologies. The insights developed during this phase subsequently informed later research efforts within the group, including work on dependable software development practices and on the integration of co-simulation approaches for analysing safety-critical sensor systems. These developments would later evolve into further lines of investigation in projects such as *neXt Sense – Safety Connected and Automated Vehicles*, as well as a dedicated PhD project on co-simulation for safety-critical sensor systems that emerged directly from this line of work and from the project-associated publication *Reliable Software Design Aided by QEMU Simulation* [17].

- **neXt Sense – Safety Connected and Automated Vehicles (SIFN-01-9999-FN-179491)**– An R&D programme in partnership with Bosch Car Multimedia Portugal, which ran from 2022 to 2025, with a budget of €14.761.998,69 and €9.856.323,39 in national public support. Its activities focused on next-generation automotive sensing technologies for connected and autonomous vehicles, organised across four sub-programmes covering MEMS sensor development, sensor reliability and safety and positioning and connectivity.

  The author's active participation in this project contributed directly to several strands of work that culminated in this thesis, with contributions primarily centred on analysing the threat space of sensor systems and developing simulation techniques for software–hardware co-development within safety-critical CPS contexts.

These projects did not determine the thesis topic, but they formed the technical and operational context in which the research was developed, particularly by providing practical exposure to automotive-grade CPS constraints and by informing methodological directions relevant to this work.

1.4 Thesis Structure

This thesis is organised into the following chapters:

Chapter 1 – Introduction

Introduces the problem addressed in this thesis by outlining the broader technological and industrial context, the motivation for the research, the main objectives pursued, and the institutional setting in which the work was developed. It also establishes the scope of the thesis and frames the main research problem addressed in the subsequent chapters.

Chapter 2 – Background and State of the Art

Reviews the theoretical and technical background required to frame the research problem. It surveys the relevant literature on CPS, analog sensor front-ends, electromagnetic disturbance mechanisms, and existing approaches to electromagnetic risk and robustness assessment, while also drawing on adjacent domains to support a broader understanding of front-end vulnerabilities.

Chapter 3 – A^{2}EM Methodology

Presents the proposed A^{2}EM methodology. This chapter introduces its conceptual foundations, defines its analytical scope, and explains how principles from adjacent domains are consolidated into an architecture-aware framework for the identification of electromagnetic vulnerabilities in analog front-ends. It further structures the analysis around explicit front-end models, resonant entry-point reasoning, and staged probing strategies under increasing levels of observability.

Chapter 4 – Operationalizing A^{2}EM

Translates the conceptual methodology into a practical workflow. It details how the proposed framework is instantiated through explicit front-end modeling, staged analysis procedures, and simulation-oriented interfaces that support its systematic application. In particular, it defines the modeling substrate adopted in this thesis and shows how the analog front-end is decomposed into functionally meaningful stages and stage-localized entry points suitable for systematic disturbance analysis.

Chapter 5 – Trace Modeling

Develops the trace-level electromagnetic modelling approach adopted in this thesis. It introduces a physically grounded thin-wire Method of Moments formulation that enables PCB interconnects to be represented as explicit electromagnetic entry points and used to derive representative excitation conditions for system-level analysis within the A^2EM workflow.

Chapter 6 – A^2EM Results

Presents the application of the developed methodology to the selected case study and analyses the results obtained. The chapter examines the identified electromagnetic vulnerabilities, their functional manifestation under disturbance conditions, and the corresponding mitigation strategy within the A^2EM evaluation framework. The analysis progresses through black-box, gray-box, and white-box regimes, linking functional screening to interpretation, attribution, and mitigation closure.

Chapter 7 – Conclusion

Concludes the thesis by synthesising the main contributions developed across the preceding chapters, revisiting how the proposed A^2EM workflow supports electromagnetic vulnerability discovery and mitigation closure in analog sensor front-ends. It also discusses the broader relevance of the thesis to ESRG's validation perspective, and outlines the most immediate directions for future research.

Chapter 2

Background

Over the past two decades, computation has become inseparable from the physical processes it governs. Algorithms have moved from passive data processing to active participation in feedback loops that shape physical behavior. They now control motion, temperature, and the physical functioning of systems, determining, in real time, how engineered systems operate across domains from manufacturing and mobility to energy, healthcare, and infrastructure.

The term CPS was introduced to describe this integration: systems in which sensing, computation, and actuation operate in continuous feedback with the physical world. They act as the coordinating intelligence of modern infrastructures, managing distributed physical processes through computation and communication while quietly orchestrating production lines, vehicles, and networks.

Yet the term itself remains contested. Some authors present CPS as the natural evolution of embedded systems; others treat it as interchangeable with the Internet of Things (IoT). To understand what CPS represent today, it is necessary to look back at their origins. Embedded systems engineering provides the conceptual and historical ground from which the CPS paradigm emerged.

2.1 Embedded Systems

Embedded systems are computing systems designed to perform a specific, pre-defined function, as opposed to general-purpose computers that support arbitrary software execution. They are characterized by goal-specific design, constrained computational resources, and dependable operation under defined environmental and temporal conditions [18]. Hardware and software are typically co-engineered to meet the requirements of a single application domain, optimizing for efficiency, reliability, and cost rather than

flexibility [19]. Many embedded systems also operate under real-time constraints, where correctness depends not only on logical outcomes but also on the precise timing of their responses [20].

The notion of an embedded system took shape during the 1970s, when microprocessors began replacing analog control circuits and electromechanical relays in industrial automation, avionics, and consumer electronics [21]. Control logic that had once been hard-wired or mechanical could now be expressed in programmable form, allowing devices to operate autonomously without reliance on central computers [22]. This transition marked a convergence between control engineering and digital computing. Early embedded controllers appeared in automotive engine management, flight-control computers, and programmable logic controllers for manufacturing [21]. Their defining engineering challenge was extreme resource limitation: processors offered only a few kilobytes of memory and minimal clock speeds. These constraints gave rise to a distinct design culture that prized efficiency, predictability, and correctness under bounded resources [22].

From the late 1980s through the 2000s, advances in semiconductor integration reshaped the landscape of embedded computing. Microcontrollers evolved into SoC devices that combined processing cores, memory, and peripheral interfaces on a single die [21]. Increasing transistor density and falling component costs enabled embedded systems to execute richer software stacks and support multitasking operating systems such as VxWorks, QNX, and Embedded Linux [18]. What had once been bare-metal control nodes now incorporated filesystems, networking, and user-level applications. During the 2000s, the introduction of multicore architectures and virtualisation support further expanded their capability, bringing them closer to general-purpose computers [23, 24]. This led to a dramatic rise in software complexity and interdependence. As embedded systems grew more capable, they also became more connected, forming distributed architectures whose collective behavior could no longer be understood or verified at the level of individual devices.

The added processing power, memory, and communication interfaces made it both practical and increasingly necessary for devices to exchange information and coordinate behavior. Control algorithms no longer had to execute locally; sensing, decision, and actuation could be distributed across several nodes. What had once been individual embedded devices became interdependent components of larger architectures, linked through deterministic buses, field networks, and eventually Ethernet-based and wireless protocols [18]. This connectivity expanded what these systems could do but also transformed the nature of their design problems: timing, causality, and reliability were no longer properties of a single controller

but of an entire network [20]. Embedded systems had evolved from isolated control units into structural elements of larger engineered systems, shaping their behavior from within.

2.2 Cyber-Physical Systems

By the mid-2000s, these systems were given a name: Cyber-Physical Systems (CPS). The term surfaced through early U.S. National Science Foundation programmes [25, 26] that sought to unite the previously distinct domains of embedded systems, control engineering, and networking into a coherent research agenda. CPS were described as engineered systems built from, and dependent upon, the synergistic integration of computation, communication, and control with physical processes, a formulation echoed in early academic papers [27, 28]. The concept gained traction because it did not propose something new so much as it named what had already emerged: systems where software and physical dynamics evolved together, joined by networks that closed the loop between sensing, computation, and actuation.

By that time, embedded platforms had long surpassed their resource-constrained origins, yet many of the abstractions and design practices inherited from those early environments remained in use, largely because they continued to work well for the problems engineers were solving. It is against this backdrop that [27] identified a deeper limitation: even the most capable hardware still relied on software abstractions that ignored the physical semantics of time and interaction. Computation was modelled as discrete and untimed, whereas physical processes were continuous and real-time, with the two domains therefore evolving under incompatible logics.

Although hardware and networking capabilities have advanced dramatically since the emergence of CPS, the theoretical issue identified by Lee [27] remains only partially resolved. In his later work, Lee reaffirms that contemporary CPS design still depends on abstractions that conceal time, concurrency, and causality behind nondeterministic software behavior. While CPS have matured technologically, their conceptual foundations remain an active field of research, where the reconciliation between logical and physical semantics continues to define the frontier of the discipline.

2.3 The Internet of Things

In parallel with the developments that led to the CPS paradigm, the Internet of Things (IoT) emerged from the networking and pervasive-computing communities. The concept first took shape in the late 1990s, when Kevin Ashton and colleagues at the MIT Auto-ID Center proposed linking everyday objects to the Internet through identification tags [29]. Their goal was commercial and logistical: to enable goods moving through global supply chains to be automatically identified, tracked, and managed in real time. These early implementations focused on what would later be called Trackable Object (TO), referring to physical entities equipped with Radio-Frequency Identification (RFID) or barcode identifiers that could be uniquely recognized and located. Although not computationally "smart," these systems introduced a new idea: that physical objects could participate in digital networks, creating continuous visibility and traceability across the material world.

During the 2000s, IoT expanded from identification to data acquisition as Wireless Sensor Networks (WSN) matured. Advances in miniaturized microcontrollers, low-power radios, and ad-hoc networking protocols such as ZigBee and 6LoWPAN enabled the rise of Data Object (DO), referring to devices that could not only be identified but also sense their environment. These systems collected parameters such as temperature, motion, and vibration, transmitting the data to remote servers for aggregation and analysis. IoT during this phase remained dominated by the networking paradigm: devices functioned as endpoints of communication rather than agents of control. Yet this stage marked a conceptual shift, as digital infrastructure began to reflect physical processes, blurring the line between information systems and the material world.

The following decade introduced direct interaction and remote control. Advances in embedded processors and cloud connectivity allowed devices to both monitor and influence physical systems, with IoT now employing the Interactive Object (IO) category, referring to endpoints that could sense their environment and actuate locally, but whose decision logic resided in the cloud. Sensor readings were forwarded upstream; actions were triggered by server-side rules, apps, or platform services and then pushed down as commands. Typical examples included remotely switched lighting, motorized blinds, smart plugs, and cloud-scheduled appliances.

A fourth stage emerged toward the late 2010s and early 2020s with the rise of Smart Object (SO). Devices gained local processing power sufficient to perform inference and decision-making close to the data source, reducing latency and dependence on the cloud while remaining networked for coordination and data

exchange. These autonomous or cognitive objects could collaborate and coordinate directly with peers, forming distributed networks of adaptive nodes. Artificial-intelligence techniques running on edge hardware amplified this tendency, giving IoT systems increasing capabilities.

Across this progression, from tagged to sensing, actuating, and autonomous objects, the IoT concept expanded from identification to interaction. What began as a communication paradigm evolved into a heterogeneous infrastructure capable of real-time monitoring and distributed control. As its technical scope widened, IoT began to overlap with the domain of CPS, adopting the same feedback and control mechanisms that underpin cyber-physical integration.

2.4 CPS vs. IoT

As the CPS and IoT domains grew out of their infancy, both drew upon the same enabling technologies, including embedded controllers, communication networks, and sensor infrastructures. As a result, industry and academia began using overlapping terms to describe systems that were technologically alike but emerged from different disciplinary lineages. The literature has since struggled to reconcile these perspectives. Industrial, academic, and governmental communities have used the terms inconsistently, reflecting distinct priorities regarding connectivity, control, and assurance. This ambiguity persisted long enough to prompt several formal attempts to define the relationship between the two domains, with such efforts appearing over the years and often reaching diverging conclusions.

One notable effort to address this ambiguity came from the National Institute of Standards and Technology (NIST) through its Special Publication 1900-202 – Cyber-Physical Systems and Internet of Things [30]. The document reviewed eleven definitions drawn from prior literature and identified four recurring criteria by which authors attempted to separate the two domains: control, platform, Internet connectivity, and human interaction. From these criteria emerged four general models of overlap between IoT and CPS:partial overlap, equivalence, CPS as a subset of IoT, and IoT as a subset of CPS. Each reflected the priorities of its originating discipline, with networking and Information Technology (IT) communities often describing CPS as a specialised segment within IoT, while control and systems-engineering fields viewed IoT as a connectivity layer within broader cyber-physical architectures.

The report concluded that none of these criteria were sufficient to define a clear boundary. Closed-loop control can appear in IoT implementations, just as CPS can depend on Internet-based communication;

similarly, the presence or absence of humans in the control loop does not distinguish the two. Observing that definitions were increasingly converging over time, NIST proposed a unified conceptual model describing both IoT and CPS as "interacting logical, physical, transducer, and human components engineered for function through integrated logic and physics". The report therefore positions CPS and IoT as fundamentally equivalent hybrid-system paradigms, differing mainly in disciplinary lineage, with CPS emerging from control and systems engineering, and IoT from networking and information technology.

A more recent review by Lesch et al. [31] revisited the same question and reached a different conclusion. After surveying more than a hundred papers, the authors observed that despite the apparent convergence noted by NIST, the literature continued to apply the terms CPS and IoT inconsistently. They classified definitions into overlapping categories and found that much of the ambiguity persisted across communities. To structure this diversity, Lesch et al. proposed a layered conceptual model depicting the historical and technical progression from IT to Information and Communication Technology (ICT), then to IoT and finally to CPS. In this view, IoT builds on IT and ICT by connecting physical entities through Internet technologies, while CPS adds the dimension of closed feedback and control between physical and cyber components.

Although this hierarchy is often presented visually, it is developmental rather than taxonomic, since each layer represents an expansion of technical capability rather than a strict inclusion relation. Conceptually, IoT provides the communication and sensing substrate upon which CPS implement coordination, control, and actuation. However, as Lesch et al. note, this distinction blurs in practice. Following their analysis of real-world use cases, the authors conclude that "the border between IoT and CPS cannot always be drawn clearly," and that the same system may be interpreted as either IoT or CPS depending on disciplinary perspective. They attribute this continuing ambiguity to the overlapping technologies and divergent origins of the two domains—networking for IoT, and control engineering for CPS.

2.5 CPS Layers

In light of the persistent ambiguity outlined in the previous section, the literature remains divided on whether CPS and the IoT should be regarded as distinct paradigms or as points along a single technological continuum. The distinction, however, becomes largely semantic in safety-critical contexts. Work on disturbance propagation, EMI-induced faults, and sensor-level security overwhelmingly adopts CPS terminology, treating these systems as engineered assemblies where computation, communication, and

physical processes interact under explicit timing and reliability constraints. For this reason, and for coherence with the assurance-oriented literature, this thesis adopts CPS as the encompassing notion.

To keep the analysis precise and technically grounded, CPS will be represented using the three-layer architecture outlined by [32]: perception, communication, and application. This abstraction is minimal but sufficient for the purposes of modeling how disturbances originate, propagate, and affect system-level behavior. It also aligns with the structure of modern embedded platforms, whose sensing, networking, and decision-making layers are increasingly implemented as separable, composable subsystems.

The perception layer holds particular relevance for this thesis. It encompasses the edge devices such as sensors, actuators, analog front-ends, and transducer interfaces, which couple the CPS to its physical environment. Despite recent increases in local processing, filtering, and even on-device intelligence, these components remain the most exposed and fragile part of the architecture. It is here that EMI, hardware non-idealities, environmental effects, and deliberate physical attacks manifest first, often as low-level signal distortions that propagate upward through the stack. Understanding the behavior of this layer is therefore essential for evaluating the robustness of safety-critical systems.

Although this thesis adopts CPS terminology, all analysis presented here applies equally to IoT systems. The perception layer is shared across both concepts: IoT and CPS differ in emphasis, not in the material substrate of edge devices. Any system built from networked sensing and actuation nodes faces the same fundamental vulnerabilities to EMI and physical perturbation, regardless of whether it is labeled as CPS or IoT. The modeling and results that follow should thus be understood as relevant to both domains.

2.6 Safety-Critical CPS Landscape

Despite being carefully designed to be fault-tolerant, CPS can still pose safety threats if lacking proper security measures, highlighting the growing interdependence of safety and security in today's technological landscape. Traditionally, continuous correct functioning of safety-critical systems was mainly achieved by relying upon redundancy. But the CPS of the future needs to be designed with architectural security concerns from the get-go, as a breach in security may cause a full system failure, regardless of how robust a system's fallback mechanisms may be [33]. More recently, academia has explored a new trend of CPS attacks capable of bypassing even complex safety- and security-oriented countermeasures, indicating a growing need for alternative design approaches [34]. These attacks, referred to as signal injection attacks,

exploit the fact that physical sensor data is not authenticated, potentially enabling complete system control in certain scenarios.

Faced with unprecedented expectations in the CPS market, current CPS development encounters mounting time-to-market pressures, along with tight Size, Weight, Power, and Cost (SWaP-C) metrics. These factors are compelling manufacturers to increasingly depend on third-party involvement across the product lifecycle, from conception and design all the way to validation and manufacturing. This may constitute a security breach, as untrusted third parties may covertly insert malicious hardware Trojans into the CPS to leak confidential information or grant unauthorized control over the system [35, 36]. On another note, these market pressures led towards various features being consolidated into a single integrated circuit, which due to the close relationship between the cyber and physical realms has resulted in the widespread adoption of AMS Integrated Circuits (ICs). These ICs are significantly more challenging to validate, as analog fault coverage isn't up to date with current needs [37]. That being said, there are active efforts to bridge this gap, with attempts to standardize analog defect modeling being made in IEEE's P2427 standard, currently in draft phase [38].

Realizing the convergence of all these topics, this thesis tackles some of the existing literature in these topics in an attempt to grasp what these seemingly unrelated topics have in common, with the main goal of producing a holistic view that would allow for an understanding of the complete scenario of all the different threats a CPS may face.

2.6.1 Analog Fault Coverage

As technological advances continue, ICs have become smaller with greater integration of silicon-implemented functionalities. This achievement is largely down to transistor size reduction and consequent lower supply voltages to keep electric fields constant. Although favorable in cost reduction and higher integration of capabilities, this scaling down of transistor technology is accompanied by greater susceptibility to radiation-induced failures, hot-spots due to increasing subthreshold leakage power values that may cause electromigration, and device degradation through aging [39, 40].

Consequently, hardware will become more and more unreliable and prone to failures as it scales down, resulting in mandatory fault-tolerance mechanisms in order to properly function. Fault-tolerant mechanisms are typically verified by means of fault injection, which is to intentionally induce faulty events to ensure that these mechanisms operate as intended and effectively withstand potential faults [41, 42]. However, as far

as analog hardware is concerned, ensuring fault injection campaigns encompass all potential faults and maintaining confidence that potential faults won't bypass fault-tolerant mechanisms remains an enduring open-ended challenge [43, 44].

Even though 80% of the circuitry in any electronic system is digital, 80% of system faults are due to analog circuit failures, with as much as 50% of testing costs being attributed to analog circuit tests [45]. A staggering 96% of in-field failures in mixed-signal automotive ICs are due to analog failures, with two thirds of these failures being attributed to test escapes [9, 10]. It then becomes evident that analog approaches fall short of digital techniques in the realm of validation and testing.

As opposed to analog circuits, digital circuits have a finite and discrete number of possible states, meaning that comprehensive fault coverage using fault injection is somewhat straightforward. In contrast, covering all fault scenarios within the analog realm proves to be nearly impossible. This is mainly due to analog circuits having a continuous input space and theoretically infinite range, along with tightly coupled behaviors across inputs and outputs. Defining what constitutes a fault in this context is also a challenge. During manufacturing, parameter deviations are commonplace, and distinguishing between a harmless parameter deviation and a defective component is not a simple task. Due to all of these factors, analog validation lacks Automatic Test Pattern Generation (ATPG) schemes [10]. Fault coverage is traditionally ensured by analyzing faulty devices returned by customers and integrating insights from these analyses into new testing protocols to detect previously undetected defects [46].

There is an increasing adoption of AMS ICs across various safety-critical domains, which elevates the expectation of achieving 0 Defective Parts per Million (DPPM) in AMS ICs [37]. Meeting the requirements of ISO 26262 mandates the use of fault injection for validation of electronic devices, but with fault coverage being such a difficult challenge in the analog segment of the ICs, it is not yet easily done. Efforts have been made to address this through initiatives like the IEEE P2427 Standard for Analog Defect Modeling and Coverage, currently in its initial draft stage [38].

2.6.2 P2427 Draft Standard for Analog Defect Modeling and Coverage

The IEEE P2427 Draft Standard focuses on Analog Defect Modeling and Coverage [38]. It is a collaborative effort involving over 30 companies and aims to provide accepted analog defect and fault models, as well as specifying information that should be included in coverage reports [47].

Defects are defined as unforeseen physical alterations in circuit elements or their connections. **Faults**, on the other hand, refer to unexpected changes in a circuit's performance.

Hard defects are defined as unexpected permanent changes of a circuit topology as a result of alterations in its physical structure. Common examples include shorts or opens that disrupt the standard flow of the circuit. These defects are straightforward in their impact, often causing clear and identifiable changes in circuit behavior.

Parametric defects, on the other hand, are characterized by variations in the operational characteristics of a circuit element. These do not necessarily involve overt structural changes but rather manifest as deviations from the expected performance parameters. Such defects might alter the electrical properties like resistance, capacitance, or the operational thresholds of components, leading to subtler shifts in the circuit's functionality. These changes can be more challenging to identify, as they subtly modify the behavior of the circuit without obvious visible signs.

Modeling these circuits is a complex task, mostly due to the continuous nature of analog signals and their dependence on physical modeling accuracy. A simple resistor has parasitic inductance and capacitance, and all values may vary with temperature, package and aging [48]. Thus, how accurate is the equivalent circuit of a component, and how all circuit input factors (such as temperature, supply voltage/noise, DC bias) are taken into account and modeled is what defines a thorough validation in the analog domain.

In 2017, Stephen Sunter initiated the revision of the ITC '97 benchmark circuits [49] for analog fault models, setting the stage for discussions that ultimately culminated in the creation of the P2427 draft standard in 2022 [38]. The collective experience of the P2427 Working Group helped set practical values for defects, such as for instance a default resistance of 200Ω for shorts and 1GΩ for opens [47]. Future opportunities to improve this standard include contributions like the work of Sadia Azam et al [50], who have proposed enhancements, notably in modeling stuck-on/off defects for MOS transistors.

2.6.3 Hardware Trojans

Amid the complexities involved in constructing a secure and dependable CPS, the development process is under substantial time-to-market and SWaP-C constraints, driven by the market's competitive demand for

new electronic solutions. This drives companies to outsource resources across various phases, encompassing design, fabrication, assembly and testing, often involving offshore entities handling any combination of these tasks. This opens up yet another possible CPS security vulnerability in the form of malicious tinkering with the outsourced parts, known as hardware Trojans [35].

A hardware Trojan constitutes an unauthorized malevolent modification or addition of circuitry to the initial design. This alteration aims to manipulate the system's functionality, thereby enabling an adversary to acquire control of the system or extract sensitive information that is safeguarded. The insertion of hardware Trojans can occur through untrusted third-party Intellectual Property (IP) providers at the Register Transfer Level (RTL) level or through fabrication tampering at untrusted foundries during the semiconductor manufacturing process. A hardware Trojan is intentionally crafted to remain dormant during potential manufacturing or production tests that might expose its existence. Typically, it remains inactive within a target system until a specific trigger prompts its activation. These alterations have been shown to exploit potential vulnerabilities, enabling the modification of system behavior and the compromise of critical information in hardware implementations through the clandestine transmission of data [35, 51–53].

Typically, hardware Trojans are engineered to remain dormant and evade detection during test phases, awakening and becoming active under a particular set of conditions once deployed in real-world applications. Trojans consist of two integral components: a **trigger** and a **payload** [54, 55].

The **trigger** serves as an activation mechanism, designed to monitor specific conditions or signals within the circuit. Its role is pivotal, as it determines the specific, often rare, circumstances under which the Trojan becomes active. This element of rarity in the trigger's design ensures that the Trojan remains dormant and undetected during standard operational testing, thus evading early detection.

The **payload**, on the other hand, is the component that executes the Trojan's harmful activity. When the trigger condition is met, the payload activates, interfacing with the circuit's original signals to perform its intended function. The effects of the payload may vary from minor alterations to the circuit's standard operations, to causing significant system breakdowns or compromising sensitive data.

Trojan triggers are usually implemented through some form of in-circuit monitoring, either via digital circuits with combinational or sequential condition matches, or through analog on-chip sensing, such as capacitor pump charges or BiCMOS temperature sensors [35].

Some Trojans don't rely on a specific trigger mechanism; instead, they are malicious alterations in the form of minor manufacturing parameter deviations. These deviations are significant enough to cause changes over time, shortening the expected lifespans of the affected ICs, and leading to unexpected early failures [56]. Although often conceptualized as security threats, these forms of Trojans pose severe safety threats in the context of safety-critical CPS.

As for payloads, a simple change in a flip-flop bit register is enough to cause massive security breaches, such as a privilege escalation attack through the overwriting of a security-critical register [51]. The implementation of payloads can vary as widely as the circuits they target, including inducing faults through crosstalk from adjacent wires [57], generating temperature rises with increased circuit activity [58], and creating alternative channels for Radio Frequency (RF) information leakage [59].

Given the severity of the problem posed by hardware Trojans, the industry and scientific literature have already adopted several countermeasure tactics [52, 53]. A common countermeasure encompasses techniques aimed at identifying hardware Trojans after the IC has been designed and fabricated. **Trojan Detection** methods leverage discrepancies in physical or operational characteristics, such as power consumption, timing, temperature, and electromagnetic emissions, to detect anomalies that suggest the presence of a Trojan [60, 61].

On the other hand, Design for Trust (DfT) strategies incorporate security considerations directly into the chip design process, aiming to prevent the insertion of hardware Trojans or facilitate their detection. These approaches include modifying designs to enhance the sensitivity of side-channel analyses, and incorporating features that obscure the chip's functionality or layout to protect against reverse engineering and tampering [62–64].

Lastly, **Split Manufacturing for Trust** approaches directly tackle the risks associated with untrusted fabrication facilities by dividing the chip manufacturing process across multiple foundries. This method involves fabricating critical layers of the chip in trusted environments, while less sensitive layers can be produced in other facilities. A notable melding of Design for Trust (DfT) and Split Manufacturing for Trust occurs in the form of splitting manufacturing of 3D IC layers, with connections between layers being purposefully obfuscated [65].

The awareness of hardware Trojans and their potential impact on safety-critical systems is not entirely new. However, the evolving nature of these threats, due to the increasing sophistication of malicious actors and

the growing complexity of CPS, keeps them an "emerging" concern in the sense that the community continues to discover new vulnerabilities and develop new countermeasures.

2.6.4 Physical-to-Cyber Attacks

Sensors play a central role in CPS and can be broadly conceptualized in distinct signal chain stages: the **transduction process** of converting physical phenomena into electrical signals, the **amplification** and **filtering** stages to enhance signal quality, and the **analog-to-digital conversion** of conditioned input signals into a readable digital format. Each of these stages of the sensor's signal processing chain exhibits distinct vulnerabilities, making them potential targets of attacks by adversaries with insider knowledge [34]. Figure 2.1 represents the possible attacks points within this signal chain.

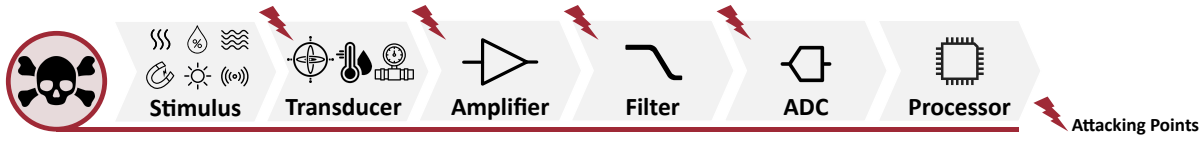


Figure 2.1: Sensor signal chain highlighting potential attacking points.

The current cybersecurity literature is progressively focusing on physical-to-cyber attacks, highlighting their importance in CPS security. This trend marks an acknowledgment of these evolving security threats, with various publications exploring how different forms of physical interference, including acoustic, electromagnetic, thermal, and optical interference, affect these sensors. Through experimental setups and theoretical analyses, these works illustrate how attackers can manipulate sensor readings, leading to erroneous outcomes in the digital system [66–70]. Some of the most common types of malicious signals employed in physical-to-cyber attacks include:

Electromagnetic Interference (EMI): EMI attacks use electromagnetic signals to disrupt the normal operation of electronic sensors by generating a strong electromagnetic field in close proximity to the target sensor. An example is the intentional interference with medical devices like pacemakers, where attackers can use EMI to alter the readings of these critical medical sensors, potentially endangering patients' lives [71].

Acoustic: Sound waves at specific frequencies are injected in acoustic attacks to interfere with or control sensors that respond to acoustic signals. Well known cases include attacks on accelerometers and gyroscopic sensors found in smartphones and drones. By emitting sound at the mechanical resonant

frequency of these sensors, attackers can disrupt their abilities to accurately measure orientation, leading to malfunctions in device stabilization and navigation [67, 72, 73].

Optical/Photo Injection: These attacks involve manipulating sensors that respond to light, like cameras or LIDAR sensors, using external light sources like lasers or Light-Emitting Diodes (LED). In the automotive industry, attackers can target LIDAR sensors used in autonomous vehicles. By shining a laser or a bright light at specific frequencies, objects or distances can be spoofed, which can confuse the target vehicle's navigation system and potentially lead to incorrect driving decisions [74, 75].

The careful crafting of malicious signals often reflects the attacker's intentions, which vary based on specific objectives like **Denial-of-Service (DoS)** or **Spoofing**.

DoS: These attacks target a sensor's normal functioning, compromising a system's availability by rendering sensors offline, essentially depriving the system of valid data inputs. An example of Denial-of-Service (DoS) can be the signal flooding of sensors, like light or infrared sensors, with an intense beam of light or infrared signals. This effectively 'blinds' the sensors and prevents them from detecting the usual range of signals they are designed to measure [76].

Spoofing: In contrast to DoS attacks, which aim to disrupt service, spoofing attacks focus on deliberate control. Here, adversaries often utilize techniques to shape the sensor measurement to specific, intended values. This approach allows them to effectively take control of systems that depend on precise physical information for their operational decisions. For example, an attacker might manipulate a navigation system that relies on gyroscope data by injecting acoustic signals, thereby changing the gyroscope's output. This leads to incorrect navigation decisions based on the altered data, allowing the attacker to influence the system's behavior by controlling its understanding of physical movement [73].

In the context of defense mechanisms, there has traditionally been a lack of focus on analog security in the development of CPS. This is largely due to the challenge of securing these physical signals, where standard digital security protocols like encryption and authentication are not applicable. As a result, the data from these sensors is often accepted without verifying its legitimacy or source. The absence of authentication and the blind trust in sensor outputs can lead to significant risks, especially when these systems are controlling critical infrastructure or machinery with closed-loops [77, 78].

While literature reviews of documented attacks have emerged in the scientific community as a direct consequence of this lack of formal security models for analog circuitry, the paper *"SoK: A Minimalist*

Approach to Formalizing Analog Sensor Security" [34] produces a comprehensive culmination of existing knowledge. It not only analyzes current threats but also introduces formalisms aimed at establishing a predictive framework in the field, namely with the introduction of a mathematical model for systematic analysis of analog sensor security. This mathematical model characterizes each sensor component with a transfer function, which describes its input-output relationship and considers the impact of unintentional noise in typical operating conditions. Building on standard approaches, the model introduces an innovative element: adversary transfer functions. This enhancement broadens the analysis to account for the effects of intentional interference, providing a more comprehensive understanding of sensor vulnerabilities. The predictive capabilities of this model are shown in the paper with a theoretical prediction of a new kind of attack.

2.7 Hardware Integrity Threats

With this multiple domain analysis, an understanding was reached that the fundamental distinction between hardware Trojans, hardware faults, and physical-to-cyber attacks is rooted in the source or motive behind the alterations in circuit behavior. This arises from the fact that, although the sources or motives behind the modifications may vary, the resulting outcomes frequently manifest in ways that render the source or motive behind them indistinguishable. This similarity blurs the lines between different causes, making it challenging to accurately link specific intentions to their tangible effects. This is represented in the Venn diagram in Figure 2.2a by its intersections, where different hardware integrity threats result in the same observable circuit alteration. For instance, a Process Reliability Trojan, characterized by the malicious deviation of a manufacturing parameter from its intended value, or a Parametric Defect, arising from unintentional manufacturing issues, can lead to identical circuit deviations. Likewise, a power supply glitch can result in the same circuit consequences as an intentional signal injection attack targeting the power supply. These similar outcomes are represented by overlapping sections within the Venn diagram. In cases where all threats converge, they collectively embody what we term as Hardware Integrity Threats (HIT), highlighting the potential for multiple sources of disruption to manifest in similar technical outcomes. An example of convergence is crosstalk on a victim trace, potentially triggered by either intentional malicious acts, like hardware Trojan insertions and physical-to-cyber signal injections, or unintentional faults. An unintentional fault could arise when a PCB trace is reconfigured to transmit a signal at a specific frequency that inadvertently causes crosstalk on a neighboring trace. In these cases, if crosstalk affecting a

victim trace could lead to system failures, it is imperative that such a trace be absolutely safeguarded, irrespective of the cause. This necessity underscores the importance of embedding detection and protection mechanisms deeply within a system's architecture.

To effectively illustrate the spectrum of potential solutions, Figure 2.2b builds upon the relationships identified in the previous HIT scheme. It presents a solution space Venn diagram, employing terminology specific to each threat's mitigation strategies as found in the literature.

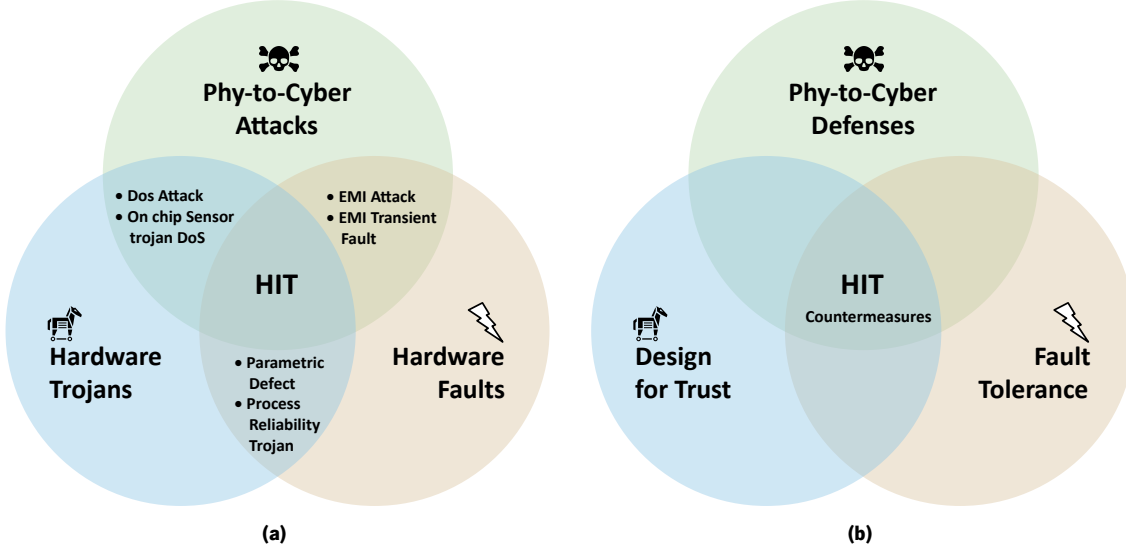


Figure 2.2: Venn diagrams of safety and security overlaps: (a) Hardware Integrity Threats (b) Hardware Integrity Threat Countermeasures.

Recognizing that different HIT often lead to similar effects, it is evident that mitigation strategies effective against a particular outcome can act as viable defenses against any threat that shares the same outcome. Therefore, investigating a comprehensive solution space that includes all these threats, while leveraging insights from across these domains, is critical for assessing and enhancing the resilience of a target architecture. From our perspective, adopting this holistic approach is essential for advancing CPS to fulfill the security and safety expectations of the current product generation. Kajol et al. [79] propose an architecture for a temperature sensor that simultaneously demonstrates resilience to faults, Trojans, and signal injection attacks. This proposal aligns with our vision of comprehensive threat mitigation and stands as an exemplary model of a HIT countermeasure.

2.8 Summary

This chapter presented a foundational background and analytical perspective to support the work developed throughout this thesis. By revisiting the evolution from embedded systems to Cyber-Physical System (CPS), and by clarifying their relationship with overlapping paradigms such as the Internet of Things, it established a consistent conceptual scope for the work. This clarification delineates the class of systems under consideration and provides the conceptual basis upon which the subsequent analysis is conducted.

Within this CPS context, the chapter examined several strands of literature in order to develop a clearer understanding of the types of threats that CPS may face. These strands include limitations in analog fault coverage and validation, physical-to-cyber attacks targeting sensor signal chains, and security risks introduced through third-party involvement and hardware Trojans. Although these topics are typically addressed within separate research communities, their joint analysis reveals common mechanisms through which disturbances originating at hardware and physical interfaces can propagate into system-level effects. This perspective is particularly relevant in the context of CPS front-ends, where electromagnetic and other physical disturbances first couple into the system.

The convergence identified through this analysis motivated the introduction of the concept of Hardware Integrity Threats (HIT), which is formally presented in this chapter. Rather than defining a new category of attacks, HIT provide an abstraction that captures heterogeneous hardware-level threats according to their origins and their impact on CPS behavior. The concept emerges as a synthesis of the analytical background developed throughout the chapter, enabling threats that are traditionally studied in isolation to be reasoned about within a unified CPS-oriented framework.

Part of the work that led to this synthesis also culminated in the peer-reviewed publication "*Let's Get Cyber-Physical: Validation of Safety-Critical Cyber-Physical Systems*" [80]. That work reflects an initial consolidation of the ideas developed here and formalizes the HIT concept in the context of CPS validation. In the broader scope of the thesis, this chapter establishes the conceptual foundation upon which the subsequent contributions are built, particularly those addressing electromagnetic vulnerability and robustness in CPS front-ends.

Chapter 3

A^2EM - Architecture-Aware Electromagnetic Methodology

3.1 Introduction

EMI has been recognized as a challenge in electronics since the early days of analog radio in the 1940s and 1950s, and it remains a persistent concern as electronic systems have become increasingly embedded in everyday life. As electronics now underpin a wide range of applications, including consumer devices, medical equipment, and modern automobiles, the impact of EMI has become particularly critical in safety-related domains, where malfunction may lead to severe consequences, including risk to human life. This growing relevance is reflected in contemporary guidance and standardization efforts, such as the IET's 2017 recommendations on EMI resilience in safety-critical systems [11] and the IEEE 1848 standard [12].

Despite the existence of these standards and practices, there are still instances where safe operational behavior of systems is not guaranteed once they are confronted with messy, dynamic Electromagnetic Environments (EME) or untested use-case scenarios [13–15].

Gkatsi et al. [16] recently proposed a risk-based approach to Electromagnetic Compatibility (EMC), named RB-EMC, as a lifecycle methodology to assess and manage EMI-related risk, reinforcing the ongoing shift from rule-based compatibility testing toward early, risk-based evaluation in domains where functional safety is at stake. It prescribes linking each identified Electromagnetic Disturbance (EMD) in the equipment's Intended Electromagnetic Environment (IEME) to the specific system vulnerability it can exploit, quantifying the likelihood and severity of each pairing, and maintaining a living risk register that records concrete mitigation actions and revisits them whenever the hardware or its installation context changes.

Yet, as Gkatsi notes [13], no fixed procedures exist for EMV discovery, leaving the EMI Identification stage highly case-dependent and reliant on known patterns, targeted bench work, and engineering intuition.

This gap widens in CPS, which make up the majority of today's E/E safety-critical platforms, when we consider the full spectrum of HIT [80] that can compromise them: hardware trojans inserted during outsourcing, latent analog-fault modes missed by coverage metrics, and physical-to-cyber attacks that inject signals directly at the sensor interface. Although these threats arise from different origins, they exploit the same physical mechanisms and therefore share a vulnerability space that current Electromagnetic (EM) checks do not systematically probe.

To address this gap, we propose A²EM, an architecture-aware electromagnetic simulation methodology to analyze CPS front-ends under EMI and systematically discover EMV. It operates by abstracting the sensor front-end into transfer-function blocks and performing controlled sweeps across the relevant signal-parameter space, allowing disturbance propagation and failure mechanisms to be explored in a structured way. While conceived to support risk-based workflows in safety-critical CPS, the approach is equally applicable to generic EMI analysis, offering a scalable framework for vulnerability discovery, resilience validation, and debugging of mitigation strategies.

This chapter presents the related EMI contextualization and background that led to the formulation of this framework, as well as the framework itself.

3.2 The role of EMI in Safety-Critical CPS

EMI enables the practical delivery of physical-to-cyber attacks because ordinary conductors, including wiring harnesses, Printed Circuit Board (PCB) traces, and sensor leads, often act as unintentional receiving structures, creating efficient coupling paths at certain frequencies. A modest waveform coupled onto these paths doesn't have to originally sit in the sensor's band to be able to affect measurements. Once inside the analog front-end, the circuitry can modify it through demodulating, rectifying, aliasing, until the conditioning chain is carrying values that look no different from the real stimulus. Interference of this kind lies outside the scope of regulatory EMC tests, which were built to police emissions and immunity for coexistence.

In the 1950s, EMI was largely treated as a nuisance to radio and television reception, and the first EMC rules were written to ensure broadcast receivers remained usable at reasonable cost. Although measurement methods and limits have grown more sophisticated since, consumer-grade EMC standards still reflect those early nuisance-mitigation principles.

Current EMC limits are drafted primarily to prevent mutual interference between products, rather than to assure robust behavior under all relevant disturbances. The European EMC Directive (EMCD) explicitly excludes safety aspects, leaving functional-safety concerns arising from EMI to be addressed by the manufacturer [81]. Harmonized immunity limits represent an economic-technical compromise, with compliance tests therefore meant to level the marketplace, granting only a presumption of conformity under EU rules. As electronics began steering cars and stabilizing power grids, it became clear that this presumption was never intended to cover safety, and that complementary safety-oriented processes were required.

Stricter EMC compliance testing, however, is not enough to solve this problem [13–15]. A digital system's behavior depends on a mix of internal states, timing offsets and stimulus combinations; even a modest microcontroller would need longer than the universe's age to step through them all ($\sim$13 billion years [82]). When real-world factors are added, such as multiple disturbances occurring simultaneously, aging components, and orientation changes, the test matrix expands further, while non-linear logic means that an untested corner can still alter the outcome. Compliance testing therefore remains essential for coexistence, but is insufficient to guarantee robust operation in safety-critical scenarios.

When real-world factors are added, such as multiple disturbances occurring simultaneously, aging components, and orientation changes, the test matrix expands further, while non-linear logic means that an untested corner can still alter the outcome.

The solution is to complement compliance testing by treating EMI as a design hazard within risk-based safety assessments. IEEE 1848 [12] has taken steps toward formalizing how EMD should be addressed in this context, yet their treatment within safety engineering remains fragmented, inconsistently applied, and often misunderstood.

Part of this stems from the fact that EMC and EMI-related safety are distinct fields with different objectives, despite both dealing with EMI. As safety discussions borrow terminology from EMC, regulatory notions such

as compliance, conformity, and immunity are often blurred together, even though each carries a distinct meaning in its original context.

Beyond this cross-domain ambiguity, EMI-related safety itself is still a recent field, and its terminology remains muddled and ill-defined. Core technical terms such as interference and disturbance, or susceptibility and vulnerability, are often used interchangeably, hindering clear discussion and slowing progress in the field

These challenges have recently prompted efforts to consolidate both perspective and vocabulary under the umbrella of risk-based safety assessments. This represents the most up-to-date attempt at unifying EMI considerations with safety engineering, although it has not yet been formalized in standards [13].

3.2.1 Risk-Based EMC Management

Over the last decade, EMI concerns in safety-critical contexts have developed along separate tracks, with distinct schools of thought emerging independently. As early as 2015, Pissoort et al. [83] presented an "EMI risk management" framing in the functional-safety context, highlighting the need for a structured approach to EMI-related hazards beyond conventional compliance boundaries. In 2017, Sabath et al. [84] outlined an ISO 31000-style workflow under the same label, aiming to bring process structure to a still loosely defined space. That same year, the IET Code of Practice on Electromagnetic Resilience [11] reframed EMI around system survivability rather than regulatory compliance and provided practical lifecycle guidance for safety-critical designs. In 2020, IEEE Std 1848 [12] introduced a resilience-oriented approach as a consensus standard, organizing lifecycle techniques and design measures for managing risks due to EMD. By the early 2020s, three near-synonymous labels were in circulation: electromagnetic resilience (IET/IEEE 1848), EMI risk management (Pissoort; Sabath), and EMC risk management / RB-EMC (EMC community/industry). These labels had partially overlapping scopes and were associated with misaligned terminology.

Gkatsi et al. [16] provided the clearest and most practically useful contribution on this issue to date, integrating and extending earlier safety-oriented EMI work into a coherent methodology, named Risk-Based EMC Management (RB-EMC), clarifying terminology, and offering the most complete risk management method for EMI currently available in the literature. Central to this unification was the resolution of long-standing semantic ambiguities, notably the formalization of four foundational terms, each defined in the EM context and usually prefixed with *electromagnetic*:

- **Disturbance (EMD)** – an external EM phenomenon (e.g., radiated field, conducted transient, or intentional emission) capable of coupling into a system.
- **Vulnerability (EMV)** – an internal system characteristic or design feature that allows a disturbance to degrade functional performance.
- **Interference (EMI)** – the resulting unintended behavior or performance degradation that arises when a disturbance successfully exploits a vulnerability.
- **Susceptibility (EMS)** – the system's measurable effects of interference under a given disturbance, reflecting how a vulnerability behaves in practice.

For example, a low-frequency radiated field from a nearby motor (the **disturbance**) may couple into a poorly shielded analog input line (the **vulnerability**), causing a shift in sensor readings (**interference**); and the system is **susceptible** only if that shifted reading crosses its decision boundaries and alters behavior (e.g., an incorrect decision or unsafe actuation).

In addition to clarifying terminology, Gkatsi sets out a multi-phased workflow methodology named Risk-Based EMC Management (RB-EMC). The key advance is a structured way of conceptualizing EM threats: EME are described in terms of their possible disturbances, victims are characterized through their vulnerability profiles, and together these form concrete disturbance–vulnerability scenarios. These scenarios become the units that engineers can model, analyze, and eventually quantify in terms of probability and consequence, giving the threat model a tangible structure that can be plugged into formal risk analysis.

Figure 3.1 summarizes the RB-EMC methodology, showing the end-to-end workflow and, for each stage, the essential inputs and resulting outputs. This methodology is structured into three main phases, which together form a complete lifecycle approach:

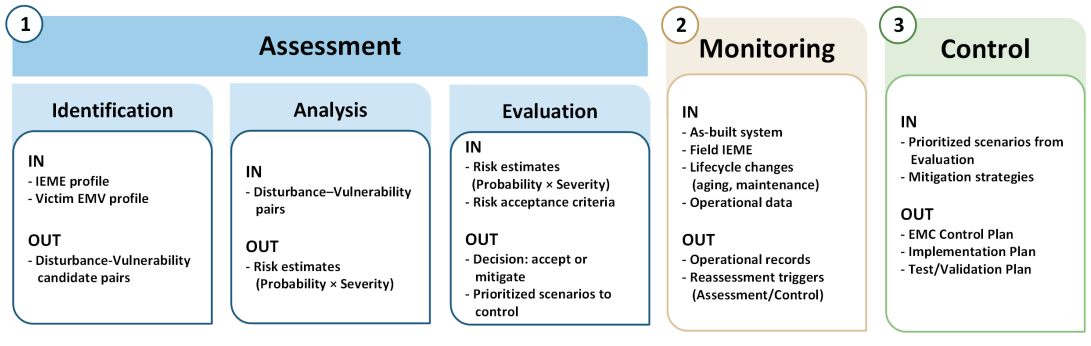


Figure 3.1: Overview of the Risk-Based EMC Management workflow.

- **EMC Assessment** – This phase identifies relevant EMD and system vulnerabilities, and assesses whether their interaction constitutes a functional hazard. This phase is itself divided in three sequential stages: Identification, Analysis, and Evaluation

  Identification – Establishes the foundation of the risk model by identifying the EMD the system may be exposed to and determining which architectural features make the system vulnerable to them. This step involves characterizing the relevant IEME, including both unintentional and intentional interference sources, and systematically mapping how these disturbances could couple into the system. Crucially, vulnerabilities are treated as architectural conditions, including layout geometry, shielding gaps, grounding topology, and signal routing, that may allow a given disturbance to interfere with signal integrity or functional behavior. The outcome of this stage is a set of disturbance–vulnerability scenarios: pairs that describe the necessary preconditions for interference to occur.

  Analysis – Estimates the probability and severity of failure for each disturbance–vulnerability scenario, assessing the likelihood that a given disturbance will be coupled in a way that activates the corresponding vulnerability and the extent to which the resulting interference could compromise the behavior of the system. The result is a set of risk values that express each scenario's potential impact.

  Evaluation – Compares the estimated risks against project-specific tolerability criteria or safety goals, determining whether each scenario is acceptable as-is or requires risk reduction; only scenarios exceeding acceptable thresholds proceed to the EMC Control phase for mitigation planning.

- **EMC Monitoring** – This phase begins after the initial EMC Assessment and spans pre-implementation, through implementation, and into in-service operation up to end-of-life. It reviews the equipment's behavior in its IEME against the assessed risks, watches for changes (environment, usage, aging), records evidence,and triggers updated controls or a return to the EMC Assessment phase whenever prior assumptions no longer hold.

- **EMC Control** – This phase turns the Evaluation decision into implemented risk reduction. It selects and applies mitigation strategies at the source, coupling path, or victim, tailored to the equipment and its environment. These strategies may include shielding, zoning, earthing/bonding, filtering, diagnostics, or architectural safeguards. The decision is documented coherently

across three plans: an EMC control plan that names the disturbance–vulnerability scenarios to be addressed, states the chosen controls and where they will be applied, and sets the acceptance criteria; an implementation plan that specifies the concrete changes, responsibilities, and schedule; and a test/validation plan that defines how success will be demonstrated, with methods, environments, setups, and pass/fail criteria aligned to those targets. After implementation and validation, responsibility returns to Monitoring, since Control and Monitoring run as a loop after Assessment.

While this methodology provides a rigorous framework for addressing the EMI problem in safety-critical systems, the field still lacks systematic approaches for EMV discovery in a truly methodical way. In practice, vulnerability identification tends to rely on architectural inspection and prior knowledge of known failure modes, rather than on formal procedures or analysis tools. The case study in Gkatsi's thesis [13] illustrates this: the vulnerabilities identified are plausible and well-justified, yet the discovery process itself remains largely ad hoc.

EMV in sensor architectures extend beyond exposure to unintentional phenomena. Adversaries with knowledge of the architecture can shape the physical environment to inject EMD at the sensor's front end, causing EMI that degrades sensing and may disrupt system operation. The dual nature of these threats, both as accidental disturbances and as vectors for deliberate attacks, suggests that a unified view of hardware threats in CPS could open the door to a broader, more comprehensive perspective of the problem. Given that these threats converge to produce similar technical outcomes, it follows that countermeasures and analytical methods designed for one type of disruption could also address others within the HIT concept.

Through this lens, we can see that while EMV discovery in safety-critical CPS still lacks dedicated analytical methods, approaches from adjacent fields can be carried over and applied. Specifically, analog cyber-security provides both a taxonomy of circuit-level vulnerabilities exploited by intentional signal injection and a transfer-function view of how adversarial EMD couples through the sensing chain [34]. Applied outside the narrow context of intentional attacks, this modeling strategy may offer a structured means of representing how disturbances couple into vulnerable architectures.

3.2.2 Physical-to-Cyber Attack Model

A common abstraction in analog cyber-security is to represent the sensing chain as a cascade of transfer functions. Legitimate inputs and injected disturbances are both subject to the same system blocks, and their propagation can be expressed mathematically. In adversarial settings, this perspective frames

injection as noise deliberately crafted to exploit the system, while in EMI contexts the same transfer paths explain how unintended disturbances may emerge as valid data. In practice, injection points fall into two broad categories: pre-transducer and post-transducer [80].

In pre-transducer attacks, the adversary targets the sensor through its physical interface. The attacker isn't limited to the modality the sensor was designed for. A transducer is inherently sensitive to its intended stimulus, but its resonances and parasitic paths make it vulnerable to other forms of energy as well. With knowledge of the transducer's physical principle, the attacker selects the signal type to which the device, or its mechanical housing, responds most efficiently.

Post-transducer injection, by contrast, targets the system after the physical signal has already been converted to voltage. At this stage, the attacker no longer interacts with the sensing interface itself, but instead aims to influence the electrical signal path directly, most commonly by coupling EMD onto traces, cables, or component leads, thereby producing EMI. Other cross-domain couplings, including optical or mechanical injection into downstream components, are theoretically possible, though far less common in the literature.

Yet even when a signal is successfully injected, there is no guarantee it will be processed as a valid measurement. Most systems include filters, gain stages, and Analog-to-Digital Converters (ADC) that reject or distort out-of-band inputs. This is especially true in post-transducer attacks, where coupling is most efficient at high frequencies, typically well outside the sensor's expected signal range. A raw RF tone or ultrasonic burst might couple in, but it won't reach the controller unless the circuit itself reshapes it into something plausible. This is where the second step of most physical-to-cyber attacks comes in: measurement shaping, the process of exploiting imperfections in the analog front-end so that the injected energy emerges from the chain as trusted data. The attacker uses the natural behaviors of amplifiers, filters, and converters to reshape the injected signal into one the system will accept.

Several measurement shaping mechanisms have been categorized in the literature [34]. Depending on the target system, attacks may exploit these effects individually or layer them together to achieve the desired output. The main ones are the following:

- **Filtering** – Ideal filters are meant to block out unwanted frequencies. However, in practice, filters response is far from ideal. Transition bands, insufficient attenuation, or degraded components can allow out-of-band energy to pass through, particularly when the injected signal is near the filter's

cutoff frequency. A carefully chosen frequency can slip just beneath the filter's practical threshold, thus bypassing the filtering stage.

- **Saturation** – When a circuit component receives more energy than it can handle, it exits its linear operating region, leading to waveform distortions, in the form of clipping and DC offset. These effects can be used deliberately to coerce the output toward a fixed value the controller interprets as real.

- **Intermodulation Distortion (IMD)** – Nonlinear circuit components can mix incoming frequencies, generating sum and difference frequencies (e.g., $f_1 \pm f_2$). These new signals (harmonics) may fall into the sensor's operational band, distorting measurements.

- **Aliasing** – When high-frequency signals are not properly attenuated before reaching the ADC, they fold back into the operational range, appearing as low-frequency artifacts. This aliasing effect can make intruding signals appear as low-frequency waveforms at the output. This is especially dangerous when the alias aligns with expected signal content.

- **Envelope Detection** – Some analog front-ends unintentionally demodulate amplitude-modulated signals. This typically occurs in stages with non-linear transfer functions, such as diode-protected inputs or early amplifier stages. An attacker can encode the desired signal onto a high-frequency carrier and rely on the circuit itself to extract the envelope, thereby producing clean, low-frequency waveforms without ever transmitting them directly.

These categories matter not only because they define systematic attack targets, but also because their presence in a circuit indicates susceptibility to EMI and, consequently, a propensity to exhibit EMV, regardless of the disturbance source. Whether the excitation is maliciously introduced or arises unintentionally, the same structural weaknesses and transfer paths tend to produce similar manifestations at the output. Under a transfer-function view, both kinds of disturbances enter as inputs that couple into the signal chain, are shaped by the same analog blocks, and can reach the controller as seemingly valid data. This adversarial modeling lens therefore extends beyond security: it offers a unified framework for reasoning about disturbances, both for uncovering latent EMV in safety assurance and for analyzing deliberate attacks. By leveraging this perspective of EMD, we are able to systematize EMV discovery in CPS front-ends.

3.3 A²EM: Architecture-Aware Electromagnetic Methodology

To address the lack of analytical, architecture-aware workflows and tools for discovering EMV in safety-critical sensor front-ends, A²EM is presented as a systematic methodology grounded in system architecture. By allowing engineers to model disturbance–vulnerability pairs directly from the system architecture, the methodology reduces the likelihood of undetected failure mechanisms and strengthens the Identification phase of RB-EMC. Rather than relying solely on historical fault data or post hoc testing, A²EM allows engineers to evaluate disturbance–vulnerability pairs proactively, directly from architectural design choices. The open-ended nature of the model further allows circuit-level expertise to complement the structured risk workflow, providing a bridge between physical signal behavior and formal safety engineering.

The primary rationale behind the proposed methodology is to leverage useful principles of physical EMC compliance tests, while trying to circumvent its impracticalities by using a simulation-based environment. By removing practical constraints inherent to physical tests, simulation enables engineers to systematically inject diverse EMD, varying parameters such as intensity, frequency content, waveform complexity, and angle of incidence. This controlled flexibility makes simulations particularly valuable in various scenarios, including identifying previously unknown EMI-induced failure modes to inform and improve risk-based assessments, as well as validating targeted countermeasures and experimenting iteratively with circuit designs.

3.3.1 Overview

A²EM models EMD as controlled voltage sources coupled at specific circuit entry points to evaluate their impact on sensor output. While adversarial noise coupling principles are adapted from [34], prior sensor security research explicitly modeled intentional malicious interference. In contrast, our methodology generalizes this concept by focusing on how diverse external EM signals physically and electrically interact with circuits, irrespective of intentionality.

Depending on the intended analytical requirements, EM signal injections can be modeled either at a more detailed circuit level, using SPICE simulations to explicitly define individual components and interconnections, or at a higher system abstraction level, using mathematical models and transfer functions to approximate aggregated circuit stages. While both abstraction levels enable the simulation of the circuit's behavior when interacting with injected signals, they differ in granularity, precision, and computational

complexity. Circuit-level modeling provides detailed insight into how specific components influence EMD interactions, whereas system-level modeling offers a computationally efficient approach with the flexibility to either abstract individual components or retain a more detailed representation of circuit stages.

To systematically assess EMI failure modes, A²EM is organized into two complementary phases: Resonant Entry Mapping and EMI Probing and Profiling. Figure 3.2 provides an overview of the overall workflow.

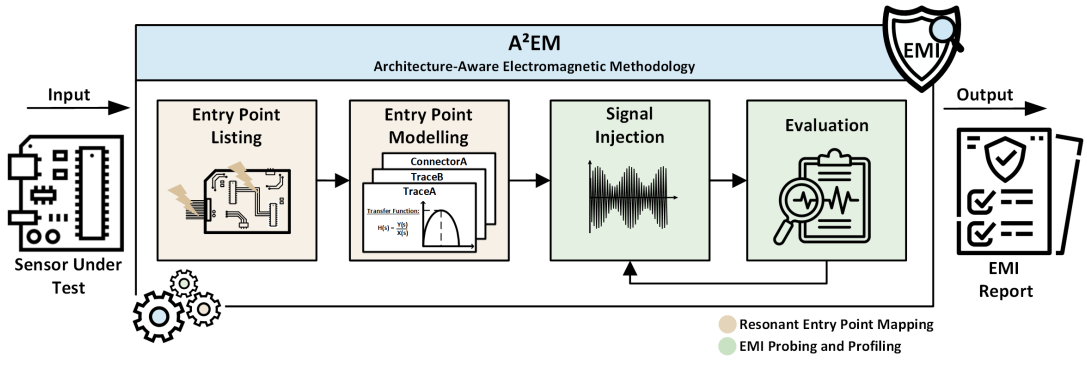


Figure 3.2: Overview of A²EM - Architecture-Aware Electromagnetic Methodology.

EMD couples into sensor circuitry at particular vulnerable entry points by interacting intricately with physical and architectural features, such as trace layout, grounding schemes, shielding structures, and component resonances. These physical properties shape how EM signals propagate, determining preferential coupling frequencies and forming distinct interference pathways. As such, in the first phase, Resonant Entry Point Mapping, circuit locations and frequency bands where EMD coupling is most efficient should be identified, ensuring that representative best-case conditions for energy transfer are considered when testing. Thus, this step encompasses analyzing the physical structure of the sensor to identify all possible entry points, as well as modeling their coupling behavior to systematically guide signal injection in the next phase.

Building on the insights gained in Resonant Entry Point Mapping, in the second phase, EMI Probing and Profiling, the main goal is to systematically inject signals at these entry points through simulation to assess whether disturbances introduced at each location disrupt the sensor's intended output. Once inside the circuit, EM signals may undergo amplification, filtering, frequency shifting, or rectification, resulting in nonlinear distortions, altered signal timing, and unintended interference with sensor measurements. As such, understanding the possible circuit transformations is crucial to successfully craft and inject signals that expose failure scenarios. By mapping how EMI transforms within the circuit, test strategies can be dynamically adjusted to better reveal weak points and evaluate their impact under different operating

conditions. The dynamic nature of this phase provides flexibility, allowing testing to be tailored to different objectives. Test strategies may be refined based on system knowledge, observed responses, and testing goals, including the identification of unknown failure modes through broad parameter sweeps and the targeted evaluation of implemented countermeasures.

3.3.2 Resonant Entry Point Mapping

Because EM signals do not couple uniformly across all circuit locations, testing all possible injection points and frequencies would be impractical and inefficient. Instead, this phase should systematically identify the spectrum of EMD frequencies most likely to enter the system by mapping circuit structures that facilitate coupling. PCB traces, vias, and cables can unintentionally serve as coupling paths for EM signals, with their vulnerability influenced by factors such as geometry, material properties, and spatial orientation. Understanding these interactions is crucial when identifying the most vulnerable circuit regions and ensuring that energy transfer is evaluated under conditions that represent the best-case exposure scenarios.

The physical phenomenon of resonance provides a mathematical basis for understanding how EMD efficiently couples into structural components of a system. When the physical length of a trace or connector aligns with either the fundamental wavelength or one of the harmonic wavelengths of the incident EM signal, maximum energy transfer occurs. Each conductor, trace, or via, depending on its physical properties, exhibits a unique set of resonant frequencies [85] that determine its vulnerability to EMD.

To quantify these interactions, transfer functions can be derived directly from these physical properties to model how external EM signals induce voltages in unintended coupling structures. To ensure comprehensive analysis, these transfer functions should incorporate worst-case incidence conditions, applying default values for polarization and angle of arrival to maximize energy transfer. This process should therefore identify circuit entry points and correlate them with corresponding transfer functions and a set/spectra of resonant frequencies.

Resonance is not confined to a single frequency point but instead extends over a continuous frequency band where energy transfer remains significant. Additionally, higher-order harmonics of the primary resonance introduce secondary peaks at integer multiples of the fundamental frequency, forming extended vulnerable regions. These harmonics define sub-bands where energy transfer, though reduced compared to the primary resonance, can still meaningfully affect circuit behavior. As a result, each resonant entry point is expected to have associated peaks in the frequency spectrum, with corresponding bandwidths

(Δf) where frequency impact is expected to remain considerable. Figure 3.3 illustrates a hypothetical sensor scenario with multiple resonant bands, each associated with a specific entry point in the circuit.

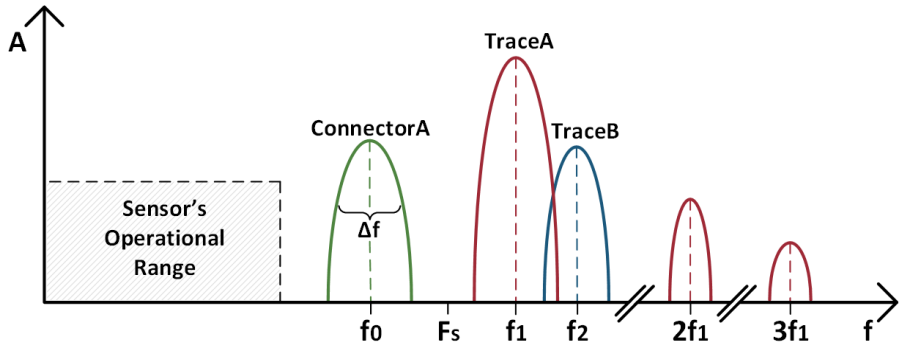


Figure 3.3: Illustatrion of hypothetical sensor resonant frequency bands.

The inclusion of harmonic sub-bands along with the main bandwidth depends on the required level of coverage. However, the extent of frequency considerations may also be refined or trimmed in the next step to optimize the assessment scope, ensuring that only the most relevant bands are analyzed based on the specific needs of the evaluation.

With a systematic identification of physical entry points, their corresponding resonant frequencies, and transfer functions that characterize EMD coupling, this phase establishes the foundation for the next step, where targeted signal crafting and diagnostic testing methodologies can be used to analyze vulnerabilities within the identified bands.

3.3.3 EMI Probing and Profiling

EMD can be both within (in-band) and outside (out-of-band) a sensor's operational frequency range, but even out-of-band disturbances are capable of affecting circuit behavior depending on their coupling pathways. After resonant entry points are identified and modeled, the next step is to investigate how signals injected at these specific locations propagate and interact with the sensor circuitry.

When an EMD successfully couples into the system through resonant entry points, it may give rise to EMI, and its effect on the circuit follows one of two scenarios. In the first, filtering circuits attenuate the interference when they are correctly designed for the susceptible frequencies, preventing it from propagating further. In the second, filters are inadequate, absent, or bypassed, allowing EMI to penetrate deeper into the circuit. Once the interference bypasses the initial filtering, it propagates into downstream circuit paths

where unintended gain or component nonlinearity can distort or demodulate it, producing interference effects that reshape the measured signal.

At this stage, the impact of EMI depends on how it interacts with the sensor's measurement process. If the interference remains confined to frequency bands outside the sensor's operational range, it may have little or no effect on sensor output. However, if EMI overlaps with relevant measurement frequencies or produces spectral artifacts that fall within them, it can degrade sensor accuracy, increase noise, or cause the sensor to produce entirely incorrect readings.

In [34], Measurement Shaping categories were originally introduced to describe how adversarial signals exploit sensor hardware to manipulate sensor measurements. Although this framework was initially developed in the context of intentional sensor attacks, it can be naturally extended to analyze how circuits unintentionally transform EMI. Whether intentional or not, electromagnetic interference can excite the same circuit behaviors leveraged by adversarial signals, resulting in measurement distortions and nonlinear transformations. While these Measurement Shaping categories are already established in the literature, our contribution lies in formulating an efficient probing methodology capable of systematically testing for EMD-induced transformations. By leveraging the structure of these categories, it becomes possible to identify the specific signal conditions required to provoke each type of circuit response, ensuring that if a given vulnerability exists, the corresponding excitation will reliably expose it. Based on this perspective, we outline below the efficient probing approaches applicable to each Measurement Shaping category.

- **Filtering** – Identifying which frequency bands remain permeable to EMD is essential to understand how disturbances propagate inside the circuit and potentially lead to interference. A single-frequency sweep at resonant entry points can empirically determine circuit permeability by verifying EMD attenuation at critical points in the signal path, rather than relying solely on predefined filter specifications. Observations at these strategic locations reveal whether residual EMD persist beyond expected levels and thus may give rise to EMI effects.

- **Saturation** – If previous analysis has identified unfiltered frequency ranges, saturation testing can be focused on the highest energy transfer points within those bands. Instead of sweeping across all frequencies, a single high-amplitude sine wave at key resonant peaks is sufficient to determine whether saturation effects emerge.

- **Intermodulation Distortion (IMD)** – Sweeping pairs of frequencies within unfiltered EMI-sensitive regions exposes what nonlinear interactions may occur. By adjusting the spacing between frequency pairs, the test ensures a virtual sweep on the band space of possible Intermodulation Distortion (IMD) artifacts.
- **Aliasing** – A systematic sweep of frequencies above $F_s/2$ reveals whether aliasing artifacts appear in the sensor's output. If aliasing remains undetected under single-frequency interference conditions, multiple frequency tests may reveal scenarios where IMD effects cause new frequency components that fall within the aliasing-prone band.
- **Envelope Detection** – Injecting Amplitude Modulation (AM) signals across unfiltered EMI-sensitive bands allows one to determine whether unintended demodulation may occur. If envelope detection is present, low-frequency artifacts will appear in the sensor's output, despite the injected signal originally existing at much higher frequencies.

Building upon the preceding probing categories, the EMI probing process is first structured around **Permeability Mapping**. This step addresses whether and where electromagnetic disturbances are able to enter the system, persist after coupling, and remain capable of affecting circuit behavior. To systematically analyze the impact that EMD may have on the circuit, it is therefore necessary to determine which frequencies remain present inside the system after the initial coupling stage. These frequencies define the bands in which EMD retain the potential to produce EMI, provided that they interact with circuit components that are sensitive or responsive to them.

Permeability, as defined in this methodology, represents the overlap between resonant entry points and insufficiently attenuated propagation paths. It identifies the subset of the resonant bandspace in which coupled disturbances are not rejected by the circuit architecture and therefore remain capable of interfering with circuit behavior. In practice, this can be assessed by remapping the bandspace derived from resonant entry point mapping using either probing results or analytical knowledge of the filtering behavior of the signal chain. This remapping reveals which portions of the initially identified resonant regions actually persist within the system.

From a permeability perspective, identifying which coupled disturbances remain present along relevant signal paths helps determine the frequency regions where EMI can produce measurable effects. These regions are represented by Δf, the darkened area in Figure 3.4, which illustrates the result of permeability

mapping as a permeable sub-region of the broader resonant bandspace previously shown in Figure 3.3. Mapping permeability therefore provides a focused view of the frequency bands in which disturbances are not sufficiently attenuated by the circuit architecture and may remain relevant for subsequent analysis.

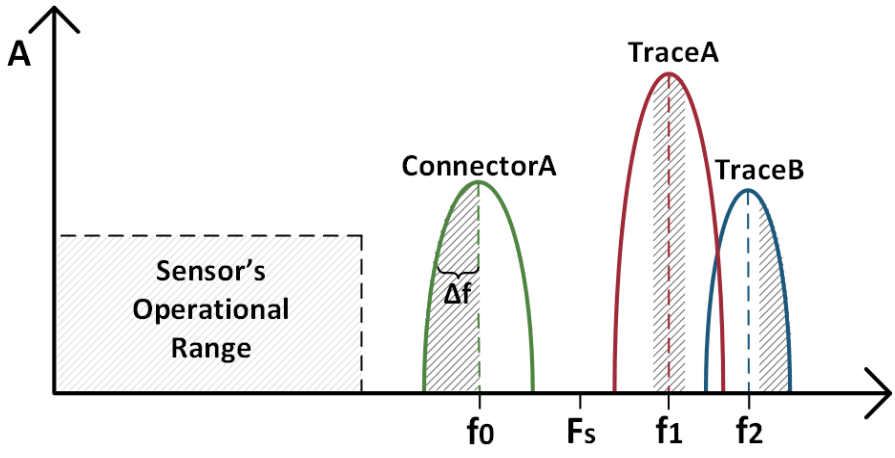


Figure 3.4: Illustration of hypothetical sensor permeable frequency bands.

The second dimension of the probing process is **Measurement Shaping**. Rather than determining whether disturbances persist inside the system, measurement shaping investigates whether the circuit architecture and its nonlinear behavior transform coupled disturbances into output deviations. This analysis can be performed directly when the objective is to test specific circuit-driven transformation mechanisms. Alternatively, when permeability mapping has already been performed, its results can be used to select the most relevant frequency regions for targeted measurement-shaping probes.

In this methodology, **measurement shaping** is organized into two fundamental categories: **Saturation** and **Unintended Demodulation**. These categories are used to translate circuit behavior patterns into specific signal injection strategies for systematic EMI probing.

Saturation remains conceptually aligned with its definition in the sensor security literature. Since saturation is an energy-driven effect, its likelihood increases when strong coupling occurs at resonant peaks, especially within permeable frequency bands and in components such as amplifiers or ADC. For this reason, saturation is most effectively probed using high-amplitude single sine-wave injections at relevant resonant frequencies.

Unintended demodulation occurs when circuit nonlinearities reshape coupled EMD, generating spectral components that were absent from the original excitation. This category consolidates IMD and envelope detection, since both arise from the same underlying mechanism of unintended frequency conversion. Even though these effects may differ in their specific manifestation, their functional consequence is similar: high-frequency disturbances can be translated into components that fall within the sensor's measurement band and therefore alter the reported output. For this reason, unintended demodulation is most effectively probed using Double-Sideband Suppressed-Carrier (DSB-SC) signal pairs, which exercise pairwise nonlinear mixing and expose potential frequency-conversion paths within the circuit.

To probe IMD, we found DSB-SC excitation to be the most effective worst-case stimulus. It fully exercises pairwise nonlinear mixing between frequency components and, in practice, also energizes the nonlinear phenomena associated with AM-type signals, thereby revealing any unintended demodulation paths in the circuit.

Aliasing does not sit neatly in either measurement-shaping category: any permeability-mapping probe that reaches the sampler will exercise content above $F_S/2$. In many platforms, the most efficient coupling paths into the analog chain occur at frequencies above the sensor's nominal band, although the relevant frequency ranges depend on the specific application. Baseband impact then usually appears only after translation, either through foldover at the sampling stage, namely aliasing, or through unintended demodulation in nonlinear front-end stages that move energy into band. Accordingly, aliasing signatures surface in direct permeability checks (when out-of-band energy reaches the sampler) and in unintended-demodulation probes (when nonlinear translation is required first).

3.3.4 Simulation Approaches

EMI-related assessments can serve different purposes depending on both the phase of the RB-EMC workflow and the level of architectural knowledge available. Within RB-EMC, *Identification* is not a one-off task but a recurring activity: it may be revisited whenever new vulnerabilities or operating conditions are suspected, with the degree of system insight at that moment determining the appropriate test strategy. Accordingly, we distinguish three complementary testing approaches:

- **Brute-Force EMI Testing (Black-Box)** – Used when entering the Identification phase with no prior system insight. The goal is exhaustive discovery of disturbance–vulnerability scenarios with minimal assumptions, ensuring that all potential failure modes are surfaced.

- **Selective Brute-Force EMI Testing (Gray-Box)** – Also part of Identification, but applied when partial knowledge of the system is already available (from schematics, layout, or prior sweeps). Instead of covering the entire disturbance space, testing is selectively directed at the most permeable bands or likely coupling paths. Importantly, because RB-EMC is cyclical, if later Monitoring or Evaluation reveals new concerns and Identification is revisited, the process naturally resumes in gray-box mode, leveraging accumulated insight rather than repeating a blind sweep.

- **Targeted Analysis and Validation (White-Box)** – In the Control phase, white-box testing moves beyond discovery to provide the diagnosis needed for mitigation. Failure modes uncovered in black- or gray-box sweeps do not themselves reveal the underlying mechanism; white-box probing deliberately recreates these scenarios with controlled signals to trace their architectural origin. Only after such diagnosis can effective mitigation strategies be designed. White-box methods are then reapplied to validate those strategies, ensuring the diagnosed vulnerabilities are closed under the same conditions that previously triggered them.

In practice, black- and gray-box testing together form the discovery backbone of RB-EMC Identification. Their objective is to systematically uncover disturbance–vulnerability scenarios, not to quantify probability or severity at this stage. White-box testing, by contrast, is tied to Control: it enables both diagnosis of the underlying mechanism and subsequent validation of mitigation strategies.

Although these approaches are anchored within the RB-EMC workflow, with black- and gray-box testing supporting Identification and white-box testing supporting Control, the same methods can also be used in a more ad hoc capacity. Developers who do not explicitly adopt RB-EMC may still employ them for exploratory testing or debugging. In that naive framing, the taxonomy simply reflects the degree of prior knowledge and precision rather than a lifecycle role. The RB-EMC integration, however, transforms these techniques from isolated test choices into elements of a repeatable process, ensuring that exploratory findings directly inform structured risk management.

Algorithm 1 illustrates how the brute-force process, which consists of two complementary sweeps, is carried out. The primary sweep injects single-frequency signals across all resonant bands to simultaneously assess permeability and test for direct interference effects, in which signals affect measurements without requiring transformation. If the sweep includes frequencies above $F_s/2$, it may also expose potential aliasing effects. The nested modulation sweep then introduces additional carefully crafted DSB-SC signals

to also probe for unintended demodulation, allowing for a systematic test of both measurement shaping effects, and full coverage of possible sensor interfering behaviors.

Algorithm 1: Black-Box EMV Discovery Procedure

Input:

$\mathcal{F}_{\text{res}}$	Set of resonant frequencies
A_{max}	Maximum test signal amplitude
Δf_{max}	Global maximum bandwidth per peak
Δf_{step}	Step size for sweeping within a band
$\Delta f_{\text{mod-max}}$	Maximum sideband spacing for DSB-SC injection
$\Delta f_{\text{mod-step}}$	Sideband spacing step size
$H_{f_{\text{res}}}(f)$	Transfer function for each resonant peak
$B(f_{\text{res}})$	Optional bandwidth per peak (default: Δf_{max})

Output: Detected EMI-induced failure modes.

Procedure:

1. **for** each $f_{\text{res}} \in \mathcal{F}_{\text{res}}$ **do**
   (a) *Bandwidth selection (peak-specific if available):* $\Delta f_{\text{test}} \leftarrow \begin{cases} B(f_{\text{res}}), & \text{if defined,} \\ \Delta f_{\text{max}}, & \text{otherwise.} \end{cases}$
   (b) **for** $f = f_{\text{res}} - \Delta f_{\text{test}}/2$ **to** $f_{\text{res}} + \Delta f_{\text{test}}/2$ **step** Δf_{step} **do**
      i. $A_f \leftarrow |H_{f_{\text{res}}}(f)| \cdot A_{\text{max}}$ *// scale amplitude using peak-specific transfer function*
      ii. Inject single-tone at f with amplitude A_f; observe sensor response.
      iii. **if** sensor output deviates **then** record the failure mode.
      iv. **for** $\Delta f = \Delta f_{\text{mod-step}}$ **to** $\Delta f_{\text{mod-max}}$ **step** $\Delta f_{\text{mod-step}}$ **do**
         A. Inject DSB-SC centered at f with sidebands at $f \pm \Delta f$; observe sensor response.
         B. **if** sensor output deviates **then** record the failure mode.

The **amplitude parameter** A_{max} defines the global EM wave strength applied to the circuit. Rather than varying amplitude at different locations, the test applies a uniform EM wave strength while systematically adjusting exposure conditions, including angles of incidence, to maximize energy coupling at each circuit entry point. The actual EMI impact at each entry point is dictated by its transfer function $H_{\text{fres}}(f)$, which inherently accounts for these exposure conditions. These transfer functions default to conditions that maximize energy coupling, ensuring that the test is conducted under best-case EMI coupling scenarios.

The **bandwidth parameter** Δf_{max} defines the range of frequencies over which the test is performed. While resonant entry mapping identifies natural EMI-sensitive regions, the actual extent of the frequency sweep is a test design choice. The cutoff point for meaningful resonance effects is inherently tied to the amplitude of the incident wave, as resonance acts as a gain function. Higher wave strengths extend the range of frequencies with meaningful coupling effects, ensuring that even low-gain resonant bands are accounted for in testing. Since different resilience levels require different stress conditions, the bandwidth selection allows for test scalability, ensuring that the sweep adapts to the desired level of analysis.

The **step resolution** Δf_{step} is used to scale appropriately with the target frequency range. In lower frequency ranges, smaller increments are necessary to ensure adequate resolution, while at higher frequencies, larger increments are sufficient to maintain efficiency without excessive data points. This adaptive scaling ensures that the sweep maintains a practical balance between test granularity and computational feasibility across different frequency ranges.

The **modulation sweep** Δf_{mod} introduces *DSB-SC* signals to probe unintended EMI transformations. While the primary sweep assesses direct EMI propagation, this secondary sweep evaluates whether interference is demodulated into unintended frequency bands. The modulation frequency Δf_{mod} is incremented in discrete steps $\Delta f_{\text{mod-step}}$, independently of the primary sweep's step resolution. Since Δf_{mod} defines the sideband spacing in DSB-SC excitation, its step size is calibrated to scale according to the resulting artifact frequency ranges that fall back into the sensor's operational range.

Failure mode detection through full-spectrum black-box sweeps provides broad coverage, but repeating such tests becomes inefficient once vulnerabilities are already known. In the RB-EMC workflow, subsequent returns to the Identification phase therefore transition naturally into a gray-box regime. Instead of re-evaluating all possible EMI interactions indiscriminately, prior evidence is leveraged to refine the test space, including permeability mapping from earlier sweeps or architectural review.

In this mode, the same brute-force procedure is applied selectively: rather than sweeping across all resonant bandwidths, signals are injected only at disturbance windows or coupling paths previously identified as permeable. This refinement improves efficiency without sacrificing systematic structure, since it maintains the principle of brute-force discovery but narrows its scope to regions where EMI is already known to propagate inside the system.

The role of gray-box testing, like black-box, remains within Identification: it uncovers disturbance–vulnerability scenarios and reveals failure modes under controlled conditions, but does not diagnose the underlying mechanisms. Failures observed in known permeable regions confirm susceptibility, while the task of tracing causes and designing mitigations belongs to subsequent white-box diagnosis in the Control phase.

Since the overall algorithmic structure is unchanged, no new process is presented. Instead, the input array of frequencies is restricted from all resonant peaks to the subset associated with permeable entry points.

At the highest level of system knowledge, Targeted Testing shifts the role of experimentation from failure detection to failure diagnosis and mitigation validation. Within the RB-EMC workflow, this corresponds to the Control phase, where disturbance–vulnerability scenarios identified earlier must be understood and closed.

Rather than broadly sweeping for potential vulnerabilities, white-box testing deliberately recreates specific scenarios under controlled conditions to uncover the underlying mechanism. Signals are crafted from prior evidence, including black- and gray-box observations, known failure patterns, and circuit analysis, so that suspected vulnerabilities can be systematically reproduced and traced to their architectural origin. This diagnostic step is essential: failure modes uncovered during Identification cannot be mitigated effectively until their causes are understood.

Once diagnosis clarifies the mechanism, mitigation strategies can be designed through filtering, shielding, layout changes, or architectural safeguards. White-box testing is then reapplied to validate these strategies, using the same tailored signals to confirm that the previously diagnosed vulnerabilities no longer manifest.

The key distinction in this phase is that signal design is guided by specific failure mechanisms rather than exploration. For example, if aliasing is suspected, controlled frequency injections above $F_s/2$ verify whether predicted artifacts appear; if saturation behaviors were indicated, signal amplitudes are deliberately varied to map coupling thresholds and confirm the effectiveness of countermeasures. By directly correlating failure behaviors with circuit responses, white-box testing transforms raw vulnerability observations into actionable diagnosis and closes the loop by validating mitigation.

3.4 Summary

This chapter addressed the problem of Electromagnetic Interference (EMI) in safety-critical Cyber-Physical System (CPS) front-ends, with a particular focus on the identification and characterization of Electromagnetic Vulnerabilities (EMV) in sensor architectures. By surveying and contextualizing the relevant literature, the chapter positioned EMI not merely as a compliance concern, but as a systemic safety hazard whose effects cannot be adequately captured by traditional EMC testing paradigms.

Building on prior research on HIT, EMI-induced failures were reframed as manifestations of a shared vulnerability space, where unintentional disturbances and deliberate physical-to-cyber attacks exploit the same architectural and physical mechanisms. This perspective enabled the consolidation of disparate threat models under a unified analytical lens and motivated the need for systematic, architecture-aware EMV discovery methods.

To address this gap, the chapter introduced and formalized the Architecture-Aware Electromagnetic Methodology (A²EM). The framework structures EMV discovery around explicit architectural models of sensor front-ends, combining resonant entry-point analysis with controlled EMI probing strategies grounded in transfer-function representations. By systematizing how disturbances couple into, propagate through, and are transformed within analog front-ends, the framework provides a concrete procedural backbone for the Identification phase of Risk-Based EMC Management (RB-EMC). In doing so, it transforms EMV discovery from an ad hoc activity into a repeatable, evidence-driven process that directly supports risk evaluation, diagnosis, and mitigation assurance.

The conceptual foundations and methodological contributions presented in this chapter culminated in the publication "*A²EM: From Vulnerability Discovery to Mitigation Assurance – An Architecture-Aware Electromagnetic Methodology for CPS*" [86], which disseminates the framework and its rationale to the broader research community.

As a primarily conceptual framework, A²EMstill leaves some aspects non-trivial, particularly its operationalization in practical workflows and its quantitative integration with formal risk metrics.. These challenges are acknowledged and intentionally deferred. Subsequent chapters build upon the framework introduced here by addressing its practical instantiation, validation strategies, and application to concrete CPS case studies, progressively bridging the gap between conceptual modeling and deployable EMI assurance practices.

Chapter 4

Operationalizing A^{2}EM

4.1 Introduction

The previous chapter introduced A^{2}EM as an architecture-aware methodology for the systematic discovery of EMV in CPS analog front-ends. By linking EMI-induced failures to concrete architectural mechanisms, A^{2}EM provides a structured alternative to ad hoc electromagnetic testing and strengthens the Identification phase of RB-EMC. While the framework establishes how vulnerabilities can be reasoned about from system architecture, it deliberately refrains from prescribing a specific modeling substrate or level of abstraction for practical execution.

Bridging this gap requires translating architectural insight into a modeling approach that remains both physically grounded and computationally tractable. In practice, systematic disturbance exploration exposes a fundamental trade-off between modeling granularity and scalability. Circuit-level environments such as SPICE are well suited to targeted analysis in white-box contexts, where detailed knowledge of the implementation is available and the objective is mechanism-level diagnosis. However, in black-box and gray-box scenarios, disturbance exploration must span wide operating conditions, multiple injection points, and large parameter spaces. Under these conditions, the computational cost of brute-force circuit-level simulation rapidly becomes prohibitive. This places conflicting demands on the analysis: models must preserve the essential signal transformations and coupling paths that govern EMI susceptibility, while remaining scalable enough to support wide parameter sweeps across injection locations, disturbance characteristics, and operating conditions. These constraints are particularly relevant when full circuit-level descriptions are unavailable or impractical to exercise exhaustively.

This chapter addresses that challenge by defining a practical workflow for instantiating A²EM as an executable modeling process. The implementation presented here should be understood as a proof of concept, intended to demonstrate feasibility and methodological coherence rather than to constitute a definitive or exhaustive modeling solution. The approach relies on transfer-function representations of analog front-ends, allowing complex circuits to be decomposed into modular stages that capture signal propagation and disturbance shaping without resorting to detailed component-level simulation. Crucially, this decomposition is not performed arbitrarily. The decomposition of the front-end into distinct modeling blocks is derived from the sensor's transduction principle, its functional signal-chain breakdown, and structural cues in the implementation that indicate plausible electromagnetic coupling surfaces. As a result, simulated injection points retain a clear physical interpretation and correspond to realistic entry paths through which disturbances may couple into the system.

The workflow introduced in this chapter is demonstrated through a PCB-integrated coupled-coil position sensor, which serves as a concrete case study illustrating how the operationalized A²EM methodology can be applied to a real analog front-end within a scalable simulation environment for systematic EMI analysis.

4.2 Modeling Approach

To operationalize the A²EM workflow, a modeling approach is required that can represent how electromagnetic disturbances couple into an analog front-end and propagate through its signal chain, while remaining computationally tractable for systematic exploration. Rather than relying on detailed component-level simulation, the approach adopted here constructs a progressive abstraction of the sensor front-end, preserving functional behavior and physically plausible coupling paths through modular representations.

The modeling workflow is structured into four successive phases, each reflecting a different level of physical and functional abstraction. The process begins with an analysis of the sensor's **Transduction Principles**, which establishes how the physical quantity of interest is converted into an electrical signal and defines the fundamental signal characteristics at the interface between the physical and electrical domains. This theoretical grounding provides the necessary context to interpret the role and intent of the subsequent signal-processing stages.

Building on this foundation, the workflow proceeds with a **Functional Signal-Chain Breakdown**, in which the main signal transformations performed by the analog front-end are identified and organized

into successive functional stages. Rather than reproducing circuit-level detail, this step focuses on the dominant operations involved in signal propagation through the front-end, including amplification, filtering, demodulation, and conversion. This functional view provides a disciplined basis for abstraction and guides the interpretation of schematic and layout information in later stages.

Based on the resulting functional structure, the third phase addresses **Entry-Point Mapping & Clustering**.Here, the abstract signal-chain representation is confronted with the physical implementation by examining schematics and PCB layout to identify geometrically extended structures, including connectors, vias, exposed pads, cable interfaces, and particularly long PCB traces, that are most susceptible to electromagnetic coupling [87, 88]. These structures are used to infer physically meaningful electromagnetic entry points, which in turn define bounded clusters of circuitry between adjacent coupling locations. This step ensures that disturbance injection points are not selected arbitrarily, but instead correspond to realistic coupling surfaces dictated by the implemented hardware.

In the final phase, **Transfer Function Modeling**, each entry-point–bounded cluster is represented using a transfer-function abstraction that captures its input–output behavior while condensing internal circuit complexity. In parallel, coupling paths associated with the identified entry points are introduced as disturbance inputs to the model. The result is a compact, scalable representation of the analog front-end that preserves the essential signal transformations and coupling relationships required for architectural EMI evaluation and disturbance-propagation analysis. Together, these four phases form the basis of the operationalized A²EM workflow developed in the remainder of this chapter.

Figure 4.1 summarizes the main stages of this modeling process for a representative temperature sensor. The figure is organized to show the progression from a high-level functional view of the signal chain to a schematic-level interpretation, and finally to a transfer-function representation suitable for system-level simulation. The Transduction Principles stage is not explicitly depicted, as it corresponds to a conceptual and analytical step that precedes circuit abstraction and serves to inform the interpretation of the subsequent modeling stages. In the sections that follow, each stage of the modeling workflow is introduced and discussed in greater detail.

4.2.1 Transduction Principles

The modeling process should begin with an analysis of the structure of the sensing element and its transduction principle, namely the mechanism by which a physical stimulus is converted into electrical signals.

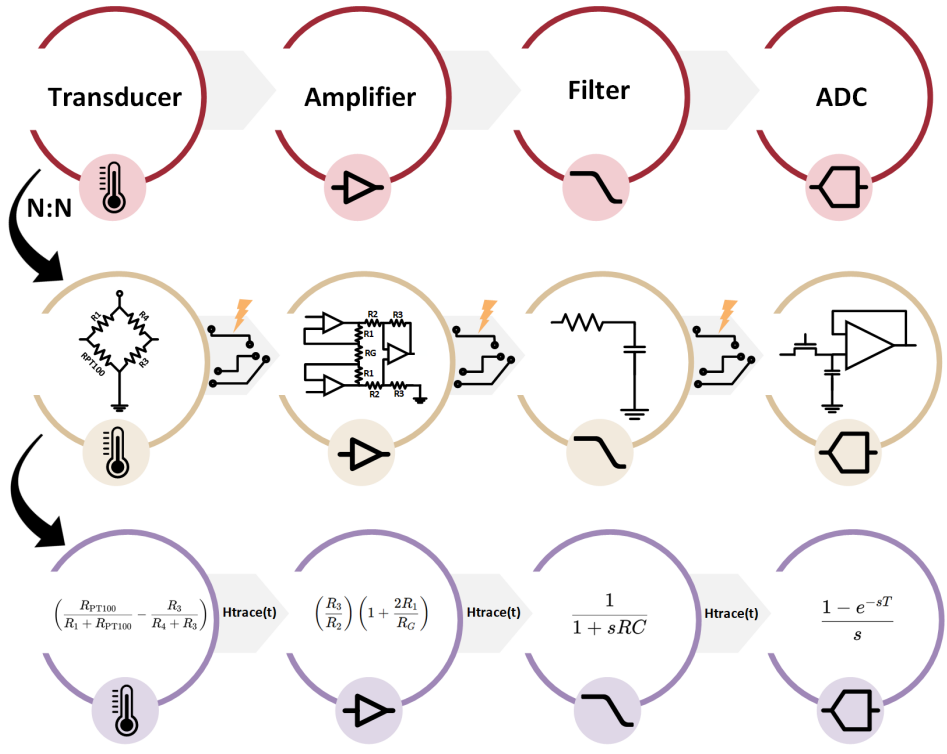


Figure 4.1: Main stages of the modeling process of a typical temperature sensor.

The emphasis at this stage is on understanding the sensing element itself: how it is built, how it interacts with the physical domain, and what form of electrical output it produces. This includes characterizing the signal's waveform, frequency content, and modulation properties. By analyzing how the transducer generates its signals, the analyst gains a fundamental picture of the sensor's operation and can anticipate the functional roles of the subsequent signal-processing stages. This foundational understanding is not only descriptive but also instrumental for later steps: it provides the information needed to emulate the transducer's output as a reusable module within the transfer-function model of the front-end.

4.2.2 Functional Signal-Chain Breakdown

The next step is a functional signal-chain breakdown, aimed at developing a fundamental understanding of how the sensor operates. Rather than providing a detailed circuit-level view, this analysis identifies the main transformations undergone by the signal, from transduction to digitization, and represents them as a sequence of functional blocks. In typical analog front-ends, this chain includes some form of amplification, filtering, and analog-to-digital conversion. Because these stages are frequently realized in distinct circuit regions, they can hint at natural partition points for abstraction and for targeted disturbance injection; however, they do not by themselves define the model partitioning. This representation establishes a

structured view of the sensor's signal chain, thereby supporting a more systematic interpretation of the schematic and PCB layout. In practice, the functional breakdown facilitates the preliminary observation of candidate clusters and, drawing on the analyst's experience and historical data, helps anticipate plausible entry points to be validated and refined in the next step. This provides the conceptual foundation that informs subsequent phases of the modeling process.

4.2.3 Entry-Point Mapping & Clustering

Building on the prior functional signal-chain breakdown, the modeling process transitions from a conceptual understanding of the sensor's role to a structural analysis of its physical implementation. The schematic and PCB layout are examined in tandem to verify anticipated disturbance entry points in geometrically extended structures such as traces, vias, exposed pads, connectors, and cable interfaces. At this stage, the definition of candidate entry points can also be refined in accordance with the specific analysis objectives. While the methodology primarily targets interconnect-level features as the most likely coupling paths, it remains flexible, allowing individual components or localized subcircuits to also be designated as entry points when justified by the system's configuration or investigation goals.

Once entry points are established, the circuit segments into groups of components and interconnects bounded by adjacent entry points. Given the chosen granularity, fewer entry points yield larger clusters aligned with broad functional stages, while more entry points produce smaller component groups that capture localized networks. In the illustrative example of Figure 4.1, this is translated as N:N relationships between the first and second stages, as each functional block may collapse or expand when compared to the schematic representation of the circuit.

4.2.4 Transfer Function Modeling

With the entry-point–bounded clusters defined, each cluster is modeled as a standalone transfer function. The resulting abstractions capture the input–output behavior of each block while condensing internal complexity into a single mathematical transformation.

In parallel, the model includes coupling paths associated with the selected entry points. For each entry point, a disturbance input with an associated coupling transfer function is introduced to represent how external electromagnetic stimuli induce voltages or currents at the corresponding structure. These coupling

functions are parameterized from layout-informed cues (i.e. interconnect length/geometry, loop area, return paths, proximity) and can be calibrated with measurements or higher-fidelity models when available. The example in Figure 4.1 also depicts this modeling stage, showing the resulting transfer functions assigned to each cluster, along with the interconnecting traces represented by their corresponding coupling functions $H_{\text{trace}}(t)$.

The result is a modular library of cluster transfer functions and entry-point coupling blocks that can be composed to assess disturbance propagation. This approach supports rapid iteration and targeted what-if analyses while keeping the model lightweight and adaptable to different analysis contexts.

4.3 Case Study Demonstration - Coupled-Coil Position Sensor

A PCB-integrated angular position sensor based on coupled-coil transduction was used to demonstrate this thesis methodology in practice. The example grounds the modeling process in a concrete system, showing how the transduction principle and functional signal-flow breakdown lead to the definition of entry-point–bounded clusters, which are then emulated through transfer-function models to enable disturbance injection and propagation analysis.

4.3.1 Transduction Principles

The sensor under study is a PCB-integrated angular position sensor that relies on inductive coupling and eddy current effects to detect rotational displacement. It consists of three planar coils etched onto a multilayer PCB: one exciter coil and two receiver coils positioned in quadrature. Figure 4.2 shows the experimental setup, where a conductive aluminum rotor is placed in close proximity to the sensor PCB to generate angular stimuli. The exciter coil is energized with a sinusoidal current, generating a time-varying magnetic field that induces eddy currents in the aluminum rotor. As the rotor turns, its angular position continually alters the distribution of these eddy currents according to its specific geometry, which in turn modulates the magnetic coupling detected by the receiver coils. According to Lenz's law, the induced currents produce opposing magnetic fields that distort the excitation pattern observed by the receivers.

As the target rotates, the mutual inductance between the exciter coil and each receiver coil changes in a way that modulates the amplitude of the induced signals. The two quadrature receiver coils capture these modulations as variations in time. The generated signals remain centered at the carrier frequency,

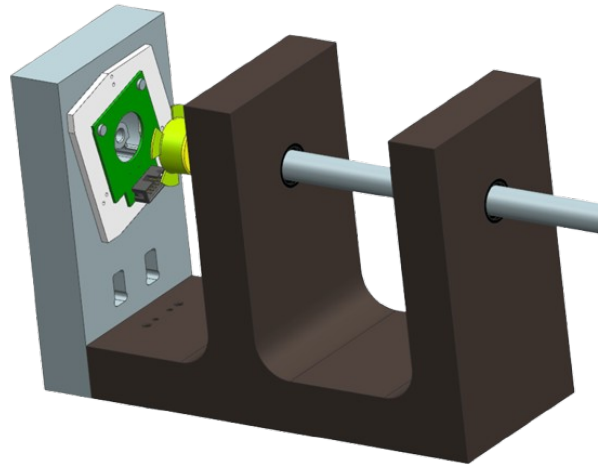


Figure 4.2: Experimental setup for the coupled-coil sensor.

while the angular information is encoded in the envelope of each waveform. In what follows, the receiver signals are denoted as `I` (in-phase) and `Q` (quadrature), corresponding to the cosine and sine components, respectively, produced by the two orthogonal coils. The mathematical expressions that represent the signals generated by both coils are as follows:

$$I(t) = D \cdot \sin(2\pi f t) \cdot [A_I \cdot \cos(\alpha(t)) + O_I] \tag{4.1}$$

$$Q(t) = D \cdot \sin(2\pi f t) \cdot [A_Q \cdot \sin(\alpha(t)) + O_Q] \tag{4.2}$$

Here, D is the excitation amplitude, f is the carrier frequency, and $\alpha(t)$ is the angular position to be retrieved. O_I and O_Q and A_I and A_Q are scaling factors of the physical coils. The A factors being the coil's gain, and the O factors being the offset.

$$I_{\text{base}}(t) \approx A_I \cdot \cos(\alpha(t)) \tag{4.3}$$

$$Q_{\text{base}}(t) \approx A_Q \cdot \sin(\alpha(t)) \tag{4.4}$$

Figure 4.3 illustrates the constant temporal evolution of these I/Q signals over a full 360° rotation, producing a sinusoidal envelope signal. The upper waveform depicts the carrier frequency while the lower ones

show modulated receiver signals *I(t)* and *Q(t)* as well as the resulting demodulated signal that is digitized. Demodulation of *I(t)* and *Q(t)* is done by synchronous multiplication with the carrier followed by low-pass filtering, yielding the baseband components:

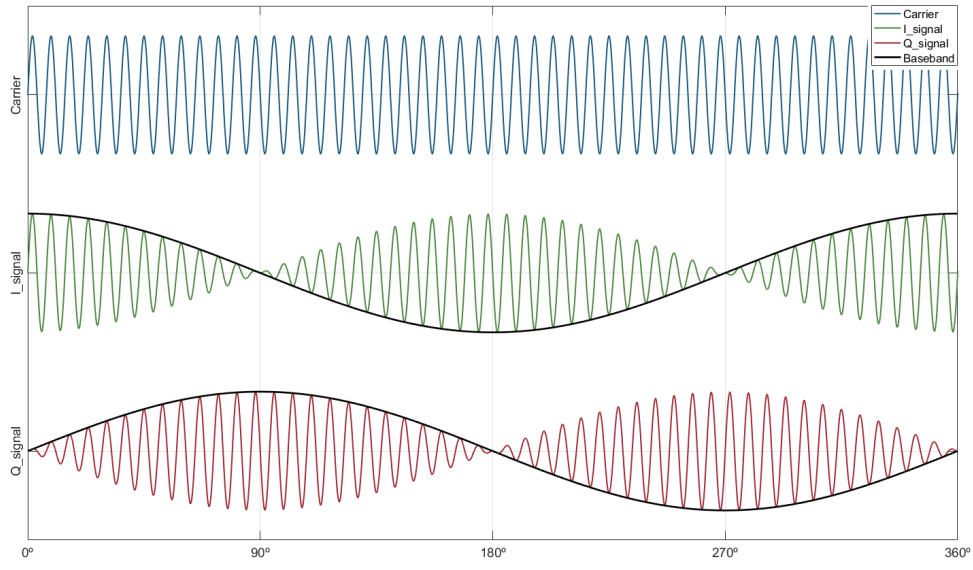


Figure 4.3: In-phase and quadrature (I/Q) signals generated by the coupled-coil sensor over one full rotation.

These clean, orthogonal baseband signals represent the final output of the analog front-end. The process of generating and extracting them is implemented in hardware and forms the structure of the front-end signal chain. The following subsection breaks this signal chain into functional blocks, each corresponding to a distinct stage in the analog processing path. Together, these blocks define how angular information is extracted from the magnetic interaction and conditioned into a form suitable for digital angle computation.

4.3.2 Functional Signal-Chain Breakdown

Figure 4.4 depicts the functional blocks abstraction of the sensor's front-end signal chain. The transduction block represents the physical sensing apparatus, which includes the exciter coil, the conductive rotor, and the two orthogonally arranged receiver coils. This block receives a sinusoidal carrier as input and produces two analog voltage signals: amplitude-modulated carriers encoding the in-phase and quadrature components of the rotor angle. These two signals are then processed through identical analog front-end paths designed to extract the underlying baseband waveforms.

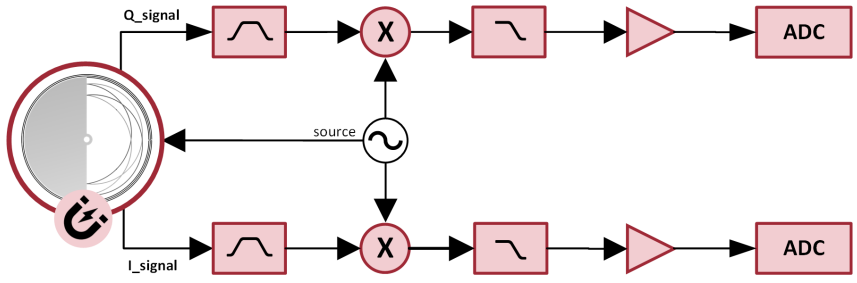


Figure 4.4: Functional signal-chain breakdown of the coupled-coil sensor.

Each signal path begins with a bandpass filter centered on the carrier frequency, attenuating spectral content outside the signal band. This rejects broadband interference, ensuring that only the modulated carrier component containing angular information proceeds to the next stage.

Demodulation is performed by a synchronous mixer that multiplies the filtered signal with a reference carrier, using the same excitation signal from the transduction block. This shifts the angular envelope from the carrier frequency down to baseband, producing amplitude-modulated signals that reflect the target's rotational position.

A low-pass filter follows, removing high-frequency components generated during mixing and isolating clean baseband waveforms. These signals are then amplified and digitized by analog-to-digital converters. The result is a pair of synchronized quadrature samples that are forwarded to the processor for angle computation.

4.3.3 Entry-Point Mapping & Clustering

Figure 4.5 shows an entry-point–aware abstraction of the analog front-end. For clarity, a single representative channel is illustrated from the transducer output to the ADC input, since the `I_signal` and `Q_signal` paths are structurally identical. Each stage is marked with a symbol corresponding to its modeled function, and orange lightning markers indicate the entry points defined for this analysis. In this case, only PCB traces were considered as entry points, with particular focus on the longest connections linking physically separated clusters. These entry points thereby set the boundaries for the clusters used in the model, providing a concrete physical anchor for the disturbance paths represented in simulation.

The signal path begins at the output of the transduction block, which represents the generation of the I/Q signals to be processed. This amplitude-modulated voltage encodes the rotor angle in its envelope and

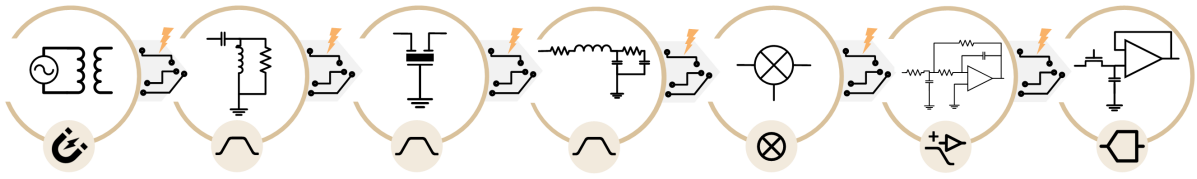


Figure 4.5: Entry-point-bounded clusters of the coupled-coil sensor.

serves as the input to the analog front-end. From this point onward, the signal is conditioned, demodulated, and digitized through a sequence of analog circuit stages.

The first processing stage is a bandpass function, modeled as a single role but implemented across three subcircuits: a high-pass filter, a ceramic bandpass filter centered at the carrier frequency, and a low-pass filter stage mainly for impedance matching. This decomposition exemplifies the N:N correspondence between functional abstraction and physical implementation: during schematic and PCB layout analysis, the initially unified cluster was refined and subdivided into three distinct blocks according to their physical realization on the PCB.

The filtered signal then feeds a mixer, implemented as a dedicated IC performing synchronous demodulation using the same carrier that drives the transduction block.

The demodulated output is routed into a stage that performs both low-pass filtering and amplification. Although these were interpreted as two separate blocks in the Functional Signal-Chain Breakdown stage, they are implemented together in a single op-amp IC with frequency-selective feedback. The collapse of these functions once again exemplifies the N:N relationship.

The final stage consists of a sample-and-hold followed by the ADC. While the converter itself is typically shielded, the analog input lines remain physically exposed. If these PCB traces are sufficiently long or poorly referenced, they can act as unintended antennas, allowing external interference to couple directly into the signal path and potentially bypass earlier analog filtering. In some cases, such signals may exploit aliasing phenomena, corrupting digital samples or inducing controllable patterns at the converter output.

4.3.4 Transfer Function Modeling

The final stage of the modeling process encodes each analog front-end block as a transfer function, forming a signal-level simulation environment built in Simulink. Figure 4.6 shows the complete model, which includes two parallel signal chains for the `I_signal` and `Q_signal` paths.

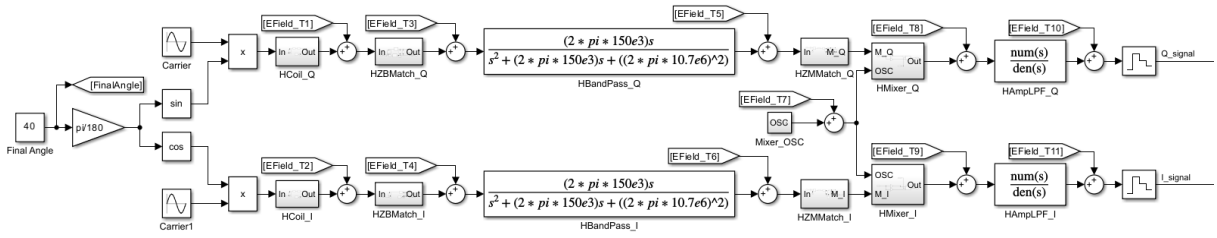


Figure 4.6: Full Simulink model of the coupled-coil position sensor.

Each block in the model carries the suffix `_Q` or `_I`, respectively, though their structure and implementation remain identical. These blocks mirror the organization of the analog front-end as defined by the entry-point–bounded clusters and are implemented as transfer functions within Simulink. Multiple modeling strategies were deliberately combined to illustrate the flexibility of this modeling approach. Some blocks are analytically derived from known circuit topologies and component values. Others are fitted to experimentally measured frequency responses, with a transfer function inserted to reflect these measurements. A particularly useful technique is using Simscape Electrical to extract transfer functions directly from the component-level structure and incorporate it into the signal-level simulation.

To analyze how EMD propagate through the analog front-end, voltage sources are injected at selected interconnects. These sources, labeled `EField_Tx` (with *x* denoting the trace index), are inserted at the entry points previously identified as plausible coupling sites. Each source can be activated individually or in combination to explore fault propagation under varied injection scenarios.

The injected signals represent external electromagnetic interference introduced directly into the signal path. In a more complete modeling setup, each disturbance's shape, strength, and frequency content would result from trace-specific coupling mechanisms derived from the PCB geometry. Modeling that coupling process is beyond the scope of this section.

The remainder of this section follows the signal path through the analog front-end, from transduction to digitization. As each stage is introduced, attention is given to the modeling strategy where relevant, to illustrate the different implementation approaches used across the chain.

Transducer Emulation and Signal Generation

The signal chain begins with the transducer, emulated by generating a target rotor angle in degrees, converting it to radians, and computing its I/Q signals. These signals are then modulated by multiplication with a carrier waveform, producing amplitude-modulated outputs that simulate the behavior of the sensor's

receiver coil. The final element of the transduction abstraction is modeled by `HCOIL_Q` and `HCOIL_I`, which capture the output impedance of the receiver coil, concluding the transducer model.

Bandpass Filtering

The next modeled section corresponds to the bandpass filter stage, which spans three blocks: `HZBMatch_I/Q`, `HBANDPASS_I/Q`, and `HZMMatch_I/Q`. These implement the functional role identified earlier in the schematic as the bandpass filtering stage. The outer blocks, `HZBMatch` and `HZMMatch`, serve as impedance matching networks that flank the central ceramic bandpass filter, while also contributing high-pass and low-pass characteristics, respectively, effectively narrowing the operational bandwidth of the stage. The matching networks are implemented using Simscape Electrical, allowing their dynamic behavior to emerge directly from the component-level description of the circuit, with their transfer functions being extracted from a schematic configuration as shown in Figure 4.7. This showcases an alternative that avoids the need for explicit analytical derivation, enabling rapid prototyping. In contrast, the central `HBANDPASS` block is modeled from a measured transfer function of a commercial filter IC, fitted to a closed-form expression. Both modeling approaches preserve compatibility with the simulation environment while demonstrating distinct paths for constructing the simulation model.

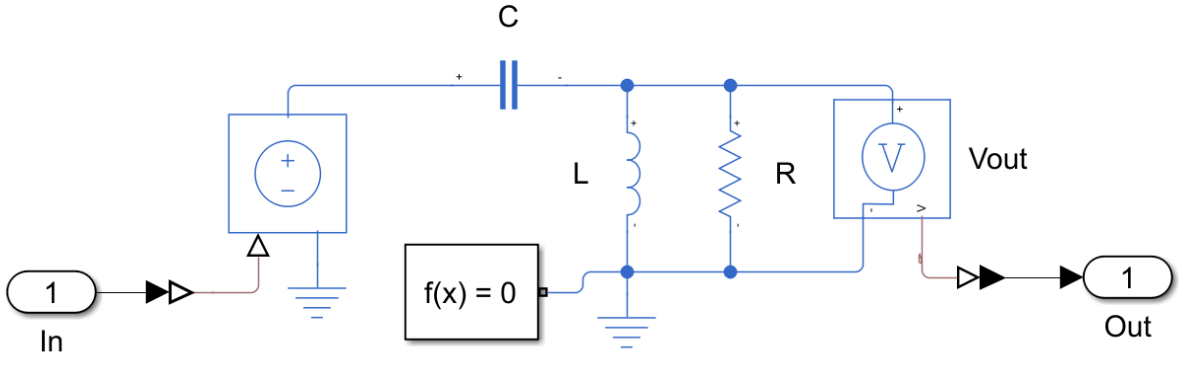


Figure 4.7: Simscape electrical representation of the HZBMatch block.

Synchronous Demodulation

The next block in the signal chain is the mixer, implemented in the model as `HMixer_Q` and `HMixer_I`. These realize an ideal switching implementation as showcased in [89]. Each mixer multiplies the incoming modulated signal by a square waveform toggling between +1 and −1, synchronized with the same carrier used to excite the transduction block. This mixing operation performs synchronous demodulation, shifting the angular envelope from the carrier frequency down to baseband. The result is a position-dependent amplitude that preserves the angular information encoded during transduction.

Baseband Conditioning

What remains after demodulation is to isolate the desired baseband envelope from the high-frequency artifacts introduced by the mixing process, specifically the sum-frequency components generated during synchronous multiplication. This task is performed by `HAmpLPF_Q` and `HAmpLPF_I`, which simultaneously apply amplification and low-pass filtering. These blocks complete the analog signal conditioning chain by attenuating undesired spectral content and scaling the signal for digitization. The modeling approach used here follows conventional practice: the circuit is analyzed and its transfer function is derived symbolically, then implemented directly as a mathematical expression using a transfer function block.

Sampling and Output Comparison

Together, these stages complete the analog front-end, culminating in `ZoH_Q` and `ZoH_I`. These final blocks capture the filtered signals into sampled representations, producing synchronized I/Q signals in discrete format, which would be theoretically appropriate for further digital processing.

To facilitate the interpretation of the simulation results, the digitized outputs are passed through an `atan2` decoder that reconstructs the angular position from the I/Q signal components, which are then converted back from radians to degrees. This estimated angle is then compared against the original input reference by direct subtraction, enabling the angular deviation induced by disturbances to be quantified in degrees. This post-processing step is not part of the modeled analog front-end itself but is included to support a more legible and easy verification of the system's functional resilience.

4.4 Summary

This chapter translated the conceptual A²EM framework into an executable modeling workflow for systematic EMI analysis of CPS analog front-ends. While Chapter 3 established the architecture-aware reasoning required to frame electromagnetic vulnerability as an identifiable and characterizable phenomenon, the present chapter focused on how that reasoning can be instantiated as a practical and scalable modeling process suitable for repeatable disturbance exploration.

The feasibility of the workflow was demonstrated through the coupled-coil position sensor case study. The example showed that a non-trivial analog front-end can be represented using modular transfer-function

models within a signal-level simulation environment, enabling structured exploration of disturbance conditions while preserving architectural meaning. The resulting implementation should be interpreted as a proof of concept, intended to validate the coherence and practicality of the operationalized A²EM workflow rather than to provide a definitive or exhaustive modeling solution.

The operationalization is anchored in architectural and functional evidence, using schematic structure, signal-chain interfaces, and stage-level frequency constraints to define the modeling substrate. This supports a functional breakdown of the analog chain aligned with the dominant front-end transformations and yields a modular transfer-function representation in which stage interfaces remain explicit. Within this representation, entry points are defined as stage-localized injection locations derived from schematic and layout evidence. For the present instantiation, only PCB interconnecting traces were exercised as explicit entry points, with their electromagnetic coupling modeling deferred to the next chapter on trace-level electromagnetic coupling.

A key implication of this operationalization is that modeling fidelity is treated as a tunable engineering dimension rather than a fixed requirement. The simulator and transfer-function models adopted here were selected to enable controllable, scalable exploration under explicit observability regimes, but they necessarily abstract away detailed circuit- and implementation-level dynamics. In contexts where higher fidelity is feasible and justified, the same workflow could be instantiated using more detailed modeling environments and richer component representations, strengthening physical representativeness and reducing reliance on idealized assumptions. Such extensions would not alter the architecture-aware structure of A²EM, but would broaden the range of front-ends and disturbance pathways for which the operational workflow can be exercised with increased confidence.

Overall, this chapter serves as the practical bridge between the architecture-aware vulnerability discovery logic introduced in Chapter 3 and the physics-informed modeling developments that follow. By formalizing an executable modeling substrate, modular stage-level representations, and a systematic entry-point definition, it enables repeatable and scalable system-level EMI analysis within the broader A²EM methodology.

Chapter 5

Trace Modeling

5.1 Introduction

The operationalization of A^2EM requires not only the identification of architecturally meaningful electromagnetic entry points, but also a physically grounded way of quantifying how external electromagnetic fields couple into the implemented hardware. While the previous chapter established how disturbances can be injected and propagated once represented within a system-level model, it deliberately stopped short of specifying how the corresponding electromagnetic coupling mechanisms should be modeled in a reusable and simulation-compatible form.

In the context of this thesis, electromagnetic entry points are restricted to PCB interconnecting traces. This choice reflects the fact that traces constitute spatially extended conductors whose geometry, routing, and electromagnetic environment make them particularly susceptible to radiated coupling. Their behavior is inherently distributed and frequency dependent, and therefore cannot be captured adequately using lumped or purely circuit-level abstractions. Under this restriction, the central problem addressed in this chapter is no longer where electromagnetic interference may enter the system, but how trace-level coupling can be computed in a manner that is both physically sound and operationally compatible with system-level simulation.

Addressing this problem requires an electromagnetic formulation capable of capturing induced current distributions along realistic trace geometries, while remaining sufficiently lightweight to support systematic exploration across multiple traces and excitation conditions. The approach adopted here builds on the thin-wire Method of Moments (MoM) formulation for PCB structures originally presented by Leone [90] and partially realized in practice through an in-house library developed by Santos [91]. Although more detailed

and higher-fidelity electromagnetic modeling approaches exist, the thin-wire formulation offers a level of physical realism that is well matched to the constraints of the considered setup and to the requirements of systematic, trace-level analysis within the A^{2}EM workflow. At the same time, the original formulation was not conceived to support trace modeling as an operational element within an architecture-driven EMV analysis methodology.

In this chapter, the thin-wire formulation is therefore extended and adapted to serve the specific requirements of A^{2}EM, introducing the final missing element required to fully operationalize the methodology: a practical and physically grounded trace model that can be incorporated directly into system-level simulation. Using standard PCB design artifacts, including layout (Gerber) data and stack-up information, the proposed approach reconstructs selected traces, discretizes their geometry, applies boundary conditions derived from the surrounding circuitry, and computes their response under external-field excitation. This response is then reformulated as an equivalent electrical excitation that can be injected into the system-level model, allowing trace-level coupling effects to be propagated through the sensor signal chain without re-solving the electromagnetic problem at run time.

In addition to introducing the trace-modeling framework, this chapter applies it to the coupled-coil position sensor used as a running case study throughout the thesis. This application provides a concrete instance of the modeling workflow and establishes trace-level coupling information that will later be exploited in system-level analyses. The interpretation of these trace-level coupling conditions in terms of functional impact is deferred to the following chapter, where they are examined using the A^{2}EM methodology.

5.2 The Challenge of Predicting Induced Currents in Real PCB Traces

When a PCB is exposed to a radiated electromagnetic environment, the external field couples directly to its copper traces [92]. These traces behave as spatially extended conductors embedded in the board layout, on which the incident field induces electrical currents. TThese currents constitute a key electromagnetic response of the PCB. Circuit-level observables, such as terminal voltages, unintended interactions between nets, and frequency-selective sensitivity to irradiation, emerge from the coupled interaction among the current distribution along the trace geometry, the electrical terminations, and the surrounding circuit environment.

Predicting that distribution is inherently challenging because a PCB trace is not a point-like element. It has finite length, bends and discontinuities, a well-defined height above a reference plane, and a cross-section that is electrically small yet not negligible in its influence on propagation and coupling. Under plane-wave illumination, different parts of the trace are excited with different phases and amplitudes [93]. At the same time, the fields produced by the induced current are not confined locally: current on one portion of the trace contributes to the field experienced elsewhere along the conductor and in its vicinity. As a result, the response at any location depends on the collective behavior of the trace and on its electromagnetic interaction with neighboring structures and terminations.

From an electromagnetic standpoint, this is the expected consequence of the field–source coupling described by Maxwell's equations: the total field is linked to sources through integral operators, and the induced current must be consistent with the field it helps create. In a full-wave formulation, this coupling naturally appears as an integral equation (with additional differential structure through current continuity), defined along the trace geometry and parameterized by frequency [93]. The implication is practical: the unknown is not a single lumped current, but a distributed function along a path, coupled to itself through the surrounding medium.

Treating the resulting continuous problem as originally formulated is not viable for realistic PCB geometries. Traces typically extend over centimeter-scale lengths, may include multiple straight and curved segments, and often involve vias, branches, and terminations. A direct handling of the continuous formulation would therefore lead to a problem with an unbounded number of degrees of freedom, which compromises its suitability as an engineering tool for systematic analysis.

This limitation motivates the need for a controlled reduction: a modeling approach that preserves the essential physics of distributed excitation, non-local coupling, and resonant behavior, while simultaneously yielding a finite problem that can be solved efficiently and whose results can be reused in system-level analyses. Rather than pursuing a single, all-encompassing electromagnetic model, the challenge becomes one of alignment: selecting a formulation whose assumptions and abstractions are consistent with the specific role that trace-level coupling plays within the A^2EM workflow. This choice determines not only which physical effects are retained, but also how naturally the resulting information can be propagated into higher-level functional models.

5.3 Solver Approaches for Trace-Level Electromagnetic Modeling

Electromagnetic solvers differ fundamentally in how they represent conductors, fields, and interactions, and in how their results can be reused beyond the electromagnetic domain. Some formulations prioritize geometric completeness, others computational efficiency, and others the ability to produce reduced descriptions that integrate naturally with circuit- or system-level simulations.

At the lightweight end, analytical *Field-to-Line (F2L)* formulations provide closed-form or semi-closed-form coupling estimates under restrictive assumptions, typically uniform plane-wave excitation, electrically simple conductor geometries (e.g., long straight runs), and idealized terminations [94, 95]. These assumptions make F2L attractive for rapid estimates, but they also limit its ability to represent the coupling mechanisms that dominate in real layouts, where geometry-driven effects concentrate around bends, discontinuities, vias, and non-uniform routing.

At the high-fidelity end, full-wave solvers such as the *Finite Element Method (FEM)* and *Finite-Difference Time-Domain (FDTD)* resolve Maxwell's equations over a volumetric discretization. They can represent detailed three-dimensional effects, but their cost is substantial: thin PCB features demand very fine meshes or grid cells, large surrounding volumes are required to model radiation boundaries, and each new frequency point or incidence condition typically requires additional computation and memory. When the analysis requires broad frequency coverage or multiple illumination configurations, the runtime and memory overhead can grow significantly, potentially making systematic, multi-trace evaluation impractical.

Between these extremes, thin-wire *Method-of-Moments (MoM)* formulations provide a pragmatic balance. By discretizing conductors primarily along their length, rather than meshing the surrounding three-dimensional volume, they retain the essential physics of distributed excitation, non-local coupling, and resonance while keeping the numerical problem tractable [90, 93]. Crucially for our use case, MoM solves directly for the induced current distribution along the trace; terminal voltages and currents then follow naturally, without the additional field-to-circuit post-processing that is often required with volumetric full-wave solvers.

MoM also aligns well with the way trace models must interface with system-level simulation: it supports explicit electrical boundary conditions at the trace terminals, including asymmetric terminations

and frequency-dependent impedances, and it accommodates spatially varying excitation along the conductor arising from incident-field phase and orientation [90]. This makes it well suited for generating trace responses that can be carried forward as electrical quantities in higher-level frameworks.

In this context, Field-to-Line approaches are often too idealized to account for layout-specific susceptibility mechanisms, while volumetric full-wave methods such as FEM or FDTD impose computational demands that hinder systematic evaluation across many traces and excitation conditions. Thin-wire MoM addresses this gap by retaining the distributed electromagnetic behavior that governs trace-level coupling, while producing results directly in electrical form. This combination makes it particularly well suited for integration into architecture-level workflows such as A^{2}EM, where physical fidelity, scalability, and reuse in system-level analysis must be balanced explicitly.

5.4 Method of Moments for Thin Wires

The thin-wire Method of Moments formulation adopted in this work, originally proposed by Leone [90], relies on a compact set of modeling assumptions. The copper trace is represented as an equivalent Perfect Electric Conductor (PEC) thin wire, meaning that the tangential electric field at its surface is assumed to vanish and that current flows without ohmic loss. This equivalent wire follows the trace centerline, is located at a fixed height above an ideal PEC ground plane, and is embedded in a homogeneous effective medium that captures the combined influence of air and substrate.

Within this representation, the unknown current is treated as a one-dimensional quantity flowing along the wire axis. The analysis is performed in the frequency domain under the $e^{j\omega t}$ convention, so that all field and current quantities are expressed as complex phasors at angular frequency ω. The use of an effective homogeneous medium is justified by a quasi-transverse electromagnetic (quasi-TEM) assumption: below an appropriate upper frequency bound, the dominant mode supported by a microstrip-like cross-section exhibits a slowly varying transverse field distribution. Under these conditions, the propagation and coupling behavior of the trace can be represented accurately using a single effective permittivity.

Within this setting, the Electric Field Integral Equation (EFIE) enforces the PEC boundary condition on the conductor: the total tangential electric field must vanish on the metal surface. At any point s along the wire contour C, the total tangential field is written as

$$E_{\text{tot},t}(s) = E_{\text{inc},t}(s) + E_{\text{scat},t}(s) \tag{5.1}$$

where $E_{\text{inc},t}(s)$ denotes the tangential component of the electric field that would exist at the wire location due to external sources and surrounding structures and $E_{\text{scat},t}(s)$ denotes the tangential field produced by the current induced on the wire. The PEC condition imposes

$$E_{\text{tot},t}(s) = 0 \qquad \Longrightarrow \qquad E_{\text{inc},t}(s) + E_{\text{scat},t}(s) = 0 \tag{5.2}$$

so that the incident and scattered tangential components must cancel along the conductor.

$$E_{\text{inc},t}(s) = -\,E_{\text{scat},t}(s) \tag{5.3}$$

The incident term $E_{\text{inc},t}(s)$ originates from a plane wave impinging on the PCB from free space. Before reaching the copper, this wave interacts with the air–substrate interface. The illumination can be decomposed into transverse-electric (TE) and transverse-magnetic (TM) components relative to the plane of incidence, each undergoing distinct reflection and transmission described by Fresnel coefficients. In what follows, $E_{\text{inc},t}(s)$ refers to the tangential field at the trace location after this interface interaction, i.e., the original plane-wave illumination filtered by transmission into the effective medium.

The scattered term $E_{\text{scat},t}(s)$ is intrinsically non-local: the field observed at a point s depends on the current distribution $I(s')$ over the entire wire. In the thin-wire EFIE, this dependence arises through inductive and capacitive coupling contributions involving the scalar Green's function G of the effective medium, together with an additional image-field contribution E_{img} associated with the presence of the PEC ground plane.

Using the PEC boundary condition to enforce the vanishing of the tangential total electric field on the surface of a perfectly conducting thin wire, the electric field integral equation (EFIE) can be written in the form

$$E_{\text{inc},t}(s) = j\omega\mu \int_C I(s')\, G\big(r(s), r(s')\big)\, \mathrm{d}s' \;+\; \frac{1}{j\omega\varepsilon}\frac{\partial}{\partial s}\int_C \frac{\partial I(s')}{\partial s'}\, G\big(r(s), r(s')\big)\, \mathrm{d}s' \;+\; E_{\text{img}}(s) \tag{5.4}$$

where $r(s)$ and $r(s')$ denote the position vectors of the observation and source points along the wire contour C, $I(s')$ is the unknown current distribution, and μ and ε are the permeability and permittivity of the effective homogeneous medium. The first two integral terms represent the tangential electric field produced on the conductor by the current flowing on the wire itself, while $E_{\text{img}}(s)$ represents the additional field contribution associated with the presence of the ground plane, required to satisfy the PEC boundary condition.

The kernel $G(r, r')$ appearing in the integral operators is the scalar Green's function of the homogeneous effective medium, defined as

$$G(r, r') = \frac{e^{-jkR}}{4\pi R}, \qquad R = \|r - r'\| \tag{5.5}$$

where k is the wavenumber in the effective medium. The quantity R denotes the Euclidean distance between the source point at r' and the observation point at r. Physically, the scalar Green's function describes the propagation of the field radiated by a unit point source in an unbounded homogeneous medium, with an amplitude decay proportional to $1/R$ and a phase variation governed by the wavenumber k.

In the presence of a perfectly conducting ground plane, the additional field term $E_{\text{img}}(s)$ can be evaluated using image theory. The effect of the ground plane is represented by an image conductor located at the mirrored position r'_{img}, obtained by geometric reflection with respect to the ground plane ($z' \mapsto -z'$), and carrying a current with the sign required by the PEC boundary condition. Rather than treating $E_{\text{img}}(s)$ as a separate external contribution, its effect can be incorporated directly into the integral operator by introducing an equivalent modified scalar Green's function,

$$G'(r, r') = \frac{e^{-jkR}}{4\pi R} - \frac{e^{-jkR'}}{4\pi R'} \tag{5.6}$$

where $R' = \|r - r'_{\text{img}}\|$ denotes the distance between the observation point and the image source. This modified Green's function accounts for the contribution of the image conductor directly at the kernel level, thereby enforcing the PEC boundary condition at the ground plane.

Using this definition, the EFIE in (5.4) can be rewritten in the equivalent form

$$E_{\text{inc},t}(s) = j\omega\mu \int_C I(s')\, G'\big(r(s), r(s')\big)\, \mathrm{d}s' \; + \; \frac{1}{j\omega\varepsilon} \frac{\partial}{\partial s} \int_C \frac{\partial I(s')}{\partial s'}\, G'\big(r(s), r(s')\big)\, \mathrm{d}s' \tag{5.7}$$

In this representation, the tangential incident field is expressed as the superposition of the fields generated by the current distribution and by its image, both evaluated using the same scalar Green's function. This form is mathematically equivalent to the formulation based on the modified Green's function in (5.6) and corresponds to the expression employed in the numerical implementation.

5.4.1 Thin-Wire Representation of PCB Traces

This analytical process requires, as a first step, an equivalent geometric representation of the trace that makes the electromagnetic formulation tractable. To this end, the copper strip is replaced by an equivalent cylindrical conductor that follows exactly the same routing as the original trace: it shares the same centreline, is located at the same height above the reference plane, and reproduces the same sequence of straight segments and bends. This abstraction is not intended to approximate the layout, but to express it in a form compatible with thin-wire electromagnetic analysis.

The trace centreline is represented as a sequence of connected straight segments forming a polyline. Each segment is characterized by three attributes: its length, the position of its midpoint, and its local tangent direction. Within the Method of Moments formulation, these segments constitute the fundamental geometric elements of the conductor. Together, they define where current flows, over what distance, and along which direction, thereby providing the spatial scaffold on which the current distribution is represented.

In parallel, the finite width and thickness of the physical copper strip are condensed into a single equivalent radius for the cylindrical wire model. This equivalent radius is chosen so that, within the quasi-TEM regime introduced earlier, the thin wire reproduces the dominant electromagnetic behavior of the original microstrip trace. In practice, this equivalence is enforced by matching the characteristic impedance of

the microstrip, which, under quasi-TEM conditions, is directly related to its per-unit-length capacitance and inductance. The thin-wire approximation further relies on a fundamental separation of scales: the equivalent radius must remain small compared to both the discretization length along the wire and the effective wavelength in the surrounding medium. These conditions ensure that the current can be treated as purely longitudinal and that the one-dimensional integral formulation remains valid.

When these conditions are satisfied, the current can be assumed to flow along the polyline axis, transverse field variations can be neglected, and the EFIE can be applied to this simplified representation without discarding the dominant coupling and resonance mechanisms of the real PCB trace.

5.4.2 Basis Functions

Once the conductor has been reduced to a chain of straight segments, the next question is how to represent the unknown current along that path. At the continuous level, the current is a function $I(z)$ defined over the entire wire; in the discretized model, it is replaced by a finite set of amplitudes attached to the geometry. Formally, the current is expanded in a set of known basis functions $f_n(z)$,

$$I(z) \approx \sum_{n=1}^{N} I_n \, f_n(z) \tag{5.8}$$

and the coefficients I_n become the electromagnetic degrees of freedom of the problem. From this point onward, the full interaction problem is recast in terms of these coefficients rather than the continuous current distribution $I(z)$.

A natural starting point is the pulse (or piecewise-constant) basis. Each basis function is associated with a single segment of the discretized wire and vanishes elsewhere. Denoting by $[z_{n-1}, z_n]$ the extent of segment n, the corresponding pulse basis function is defined as

$$f_n(z) = \begin{cases} 1, & z_{n-1} \leq z \leq z_n, \\ 0, & \text{otherwise.} \end{cases} \tag{5.9}$$

With this choice, the coefficient I_n can be interpreted directly as a constant current flowing along segment n.

This representation is simple to implement and can be acceptable for electrically short conductors. For PCB traces of practical length, however, it quickly reveals its limitations. Pulse bases allow the current to change abruptly from one segment to the next, introducing discontinuities that have no physical counterpart on a thin PEC conductor. When the trace becomes electrically long and supports standing-wave behavior, a piecewise-constant description cannot reproduce the smooth spatial structure of the true current distribution. While accuracy can in principle be recovered by refining the segmentation, this requires a large number of segments and leads to a correspondingly large MoM system.

To obtain a smoother and more physically meaningful representation, the present formulation adopts rooftop (piecewise-linear) basis functions [91]. Each rooftop basis is associated with an internal node of the discretization and spans the two segments that meet at that node. Let z_{n-1}, z_n, and z_{n+1} denote three consecutive node positions, and let $\Delta\ell_n = z_n - z_{n-1}$ and $\Delta\ell_{n+1} = z_{n+1} - z_n$ be the corresponding segment lengths. The rooftop basis function associated with node z_n is defined as

$$f_n(z) = \begin{cases} \dfrac{z - z_{n-1}}{\Delta\ell_n}, & z_{n-1} \le z \le z_n, \\ \dfrac{z_{n+1} - z}{\Delta\ell_{n+1}}, & z_n \le z \le z_{n+1}, \\ 0, & \text{otherwise.} \end{cases} \tag{5.10}$$

On the left segment, the function rises linearly from zero at z_{n-1} to a maximum at z_n; on the right segment, it decreases linearly back to zero at z_{n+1}. Outside these two segments it vanishes identically. When the current is expressed as a superposition of such rooftop functions with coefficients I_n, the resulting distribution varies linearly within each segment and remains continuous at all internal nodes. Changing a single coefficient affects the current only in a narrow neighbourhood around the corresponding node.

Physically, this behavior matches the expectations for current flowing on a thin PEC conductor in a quasi-TEM regime. Rooftop functions enforce continuity, a fundamental property of currents on conducting traces, and provide sufficient flexibility to represent gradual variations of amplitude and phase along the line. They are well suited to capturing standing-wave patterns, distributed resonances, and the smooth excitation profiles induced by an incident electromagnetic field. For a given segmentation, rooftop bases therefore deliver a more accurate and numerically stable approximation of the true current distribution than a piecewise-constant representation.

The rooftop basis has two practical advantages in this setting. First, it enforces continuity of the current by construction, which aligns with the expected behavior of current on thin PEC conductors in quasi-TEM conditions. Second, because each unknown influences only a local neighbourhood while producing a smooth spatial profile, rooftop bases typically achieve a given level of accuracy with fewer degrees of freedom than a piecewise-constant model, reducing the size of the MoM system to be solved.

The corresponding trade-off is that each unknown overlaps two segments, making matrix assembly slightly more involved: interactions must be evaluated between piecewise-linear shapes rather than constants, and near-self terms still require careful treatment due to the weak singularity of the thin-wire kernel. In practice, this additional local complexity is outweighed by the improved accuracy per unknown and the superior representation of resonant behavior along electrically long PCB traces.

5.4.3 Electromagnetic Coupling in the Discretized EFIE

Once the EFIE has been formulated for a thin-wire representation and the current has been expanded in rooftop basis functions, the next step is to convert the resulting continuous integral equation into a finite system of algebraic equations. The central idea is to approximate the distributed current along the trace with a finite set of coefficients and to enforce the EFIE in a consistent, weighted sense at a finite number of locations along the conductor.

In this discretized setting, each rooftop coefficient controls the amplitude of a piecewise-linear current contribution local to a node of the polyline. That local current distribution radiates into the surrounding effective medium, and its influence is evaluated back on the conductor through the EFIE kernel. Repeating this operation for all basis functions produces a collection of coupling terms that quantify how current flowing on one portion of the trace affects the tangential electric field experienced elsewhere.

These coupling terms assemble naturally into an impedance matrix. Off-diagonal entries represent mutual interactions between distinct portions of the conductor, capturing both near-field and radiative effects as well as the contribution associated with the ground plane through the image-current construction. Diagonal entries correspond to self-interaction, describing the response of each local current contribution to its own radiated field. In thin-wire formulations, this self-interaction requires special care because the kernel becomes singular as source and observation points coincide. In practice, the singular contribution is isolated so that the self term remains finite and numerically well-behaved, while the remaining smooth part is

evaluated numerically. The result is a dense interaction matrix that captures the non-local electromagnetic coupling along the entire discretized trace.

Excitation by the external environment is incorporated within the same framework. The incident plane wave, after interacting with the air–substrate interface, gives rise to a transmitted field at the PCB surface whose dependence on incidence angles, polarization, and dielectric properties can be expressed in closed form. This illumination is parameterized by the incidence angles (θ, ϕ) and the polarization angle γ, as illustrated in Figure 5.1. The angle θ measures the elevation of the propagation vector $\mathbf{k}_0$ with respect to the surface normal z, ϕ defines the azimuth of the projection of $\mathbf{k}_0$ onto the PCB plane (x, y), and γ specifies the orientation of the electric field vector $\mathbf{E}_0$ within the plane orthogonal to $\mathbf{k}_0$.

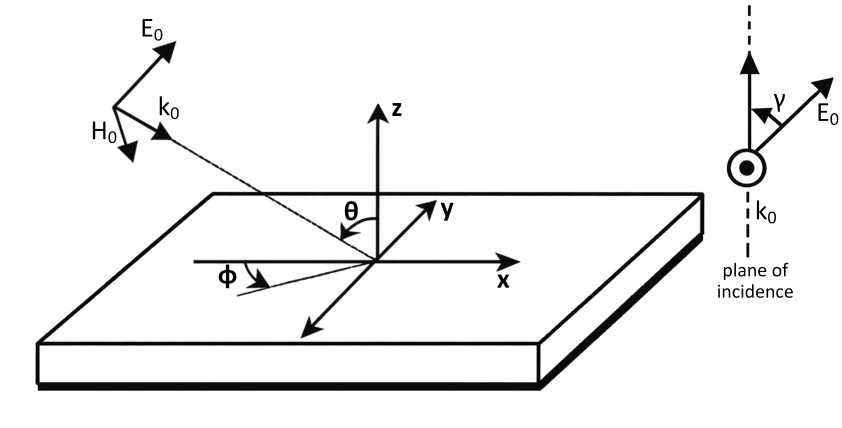


Figure 5.1: Incidence geometry and definition of the angles θ, ϕ, and γ.

Together, these three angles determine both how much of the incident plane wave is transmitted into the effective medium, as governed by the TE/TM Fresnel transmission coefficients of the grounded dielectric slab, and how the resulting field is oriented relative to the local direction of the trace. Within the thin-wire formulation, however, only a single component of this field contributes to the excitation: the component tangential to the trace at each segment centre, consistent with the assumption that current flows exclusively along the wire axis. Projecting the local illumination onto the segment tangent therefore produces a scalar excitation value V_m for each testing location, quantifying how strongly the external environment attempts to drive current along that specific portion of the conductor.

Accordingly, for each discretization element, the locally transmitted incident field is evaluated at the segment centre and projected onto the corresponding tangent direction. This projection yields a scalar excitation value for each testing location, forming the entries of the right-hand-side vector in the discretized

EFIE and quantifying how strongly the external field attempts to drive current along each portion of the trace.

After assembling both interaction and excitation contributions, the continuous EFIE is replaced by a finite linear system of the form

$$Z\,I = V \tag{5.11}$$

where I collects the unknown rooftop current coefficients, Z encodes the electromagnetic coupling along the discretized conductor (including the image contribution), and V represents the external-field excitation after interface transmission and local projection. Solving this system yields the induced current distribution along the trace, which forms the basis for computing the induced electrical quantities of interest in the subsequent stages of the modeling workflow.

5.5 Trace Modeling Process

This section describes how the thin-wire Method of Moments is embedded into a practical trace-modeling workflow and integrated into the A^{2}EM modeling infrastructure. The objective is not to introduce a new electromagnetic formulation, but to show how an established MoM approach [90] is operationalized so that PCB traces can be treated as electromagnetic entry points within a system-level analysis framework.

Viewed in a deliberately simplified and reductive way, the process takes a selected PCB trace, applies an external electromagnetic excitation, accounts for the electrical loading imposed by the surrounding circuitry, and computes the voltage induced at the trace terminals. This computation is carried out for a specific excitation condition, defined by a single frequency and a given illumination configuration, and produces an equivalent electrical response that characterizes how that trace reacts to external fields under those conditions.

Crucially, this modeling chain is not tied to a single trace or excitation. Once implemented, the same sequence of steps can be applied repeatedly to different traces and different excitation configurations as part of the construction of the A^{2}EM modeling substrate. While the numerical values change from case to case, the structure of the process remains invariant. The remainder of this section introduces the individual stages of this workflow and clarifies how they combine to transform PCB layout and circuit

information into trace-level electromagnetic models. Figure 5.2 provides an overview of the complete process before each stage is detailed in turn.

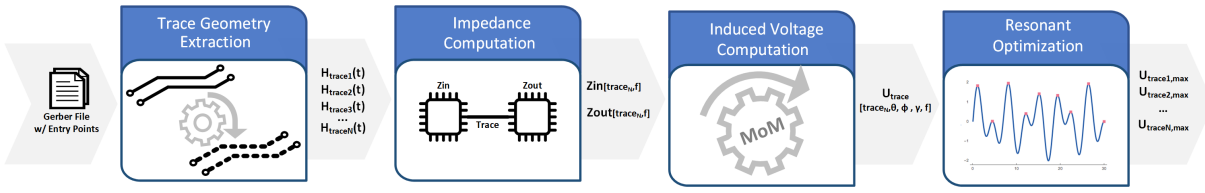


Figure 5.2: Overview of the trace modeling process.

5.5.1 Overview

The trace-modeling workflow follows directly from the thin-wire MoM formulation introduced in the previous section. At a high level, the Method of Moments requires a geometric representation on which the unknown current is defined, appropriate electrical boundary conditions at the trace terminals, and a consistent description of how the external electromagnetic environment drives current along the conductor. The stages described below provide these ingredients and use them to compute trace-level responses under external illumination.

The process begins with **Trace Geometry Extraction**, which constructs the geometric support required by the thin-wire EFIE. Starting from the PCB layout (Gerber data), and from a set of nets selected by the analyst as candidate electromagnetic entry points, each net is reconstructed as a continuous routed path on its corresponding layer. This path is then converted into a polyline representation consisting of a chain of straight wire segments with explicit node coordinates and local directions, consistent with the thin-wire abstraction used in the MoM derivation.

Next, **Impedance Computation** establishes the terminal conditions needed to close the MoM problem. Using the circuit description and the models of the blocks connected to the trace (e.g., coils, filters, front-end amplifiers), this stage determines the frequency-dependent impedances presented at the trace endpoints. These impedances act as electrical boundary conditions for the solver, ensuring that the induced current distribution is consistent not only with the incident field, but also with the way the trace is loaded by the surrounding circuitry.

With the discretized geometry and the terminal impedances available, **Induced-Voltage Computation** evaluates the electromagnetic response of each trace by assembling and solving the thin-wire MoM system in a structured and reduced manner. Rather than exploring the full space of possible illumination

conditions, this stage computes, for each frequency sample, a minimal set of baseline responses that fully characterize how the trace couples to an external field. Specifically, the incident field is evaluated along the trace, projected onto the local conductor direction, and combined with the non-local electromagnetic interactions captured by the impedance matrix to obtain the induced current distribution and the corresponding terminal voltages for these baseline cases. Taken together across the frequency grid, these results form a compact trace-level description from which the response to arbitrary field orientations can later be reconstructed without additional electromagnetic solves.

Building on this representation, **Resonant Optimization** analyses the frequency-dependent trace responses to identify the conditions under which coupling becomes most severe. Using the baseline solutions produced in the previous stage, it reconstructs the induced voltages associated with different excitation configurations and searches for the combinations that maximize coupling. The analysis retains only the dominant resonant cases, together with the excitation parameters associated with the strongest induced voltages. The outcome is a reduced, physically meaningful set of trace-level operating points that preserves the limiting electromagnetic scenarios while remaining lightweight enough for direct use in subsequent system-level simulations.

5.5.2 Trace Geometry Extraction

The trace discretization stage transforms PCB layout information into the geometric input required by the thin-wire MoM formulation. Starting from the Gerber descriptions of the conductive layers, it reconstructs the routing of each selected trace as an ordered sequence of one-dimensional segments embedded in three-dimensional space. The main operational steps of the trace geometry extraction stage are summarized in the flowchart shown in Figure 5.3.

Each copper layer is imported using a CAD/CAM parser that extracts the planar vertex coordinates of conductive features from the Gerber data. These coordinates are returned as ordered lists of (x, y) points, grouped by trace and separated by delimiter entries. All coordinates are converted to meters and expressed in a common global reference frame, with the board corner as origin and the ground plane located at $z = 0$. Within each layer, the ordering of vertices preserves the original routing continuity of each trace.

Each trace polyline is then re-sampled to enforce a quasi-uniform segmentation consistent with the chosen target element length Δl. For every consecutive pair of vertices, the segment length is evaluated: if it

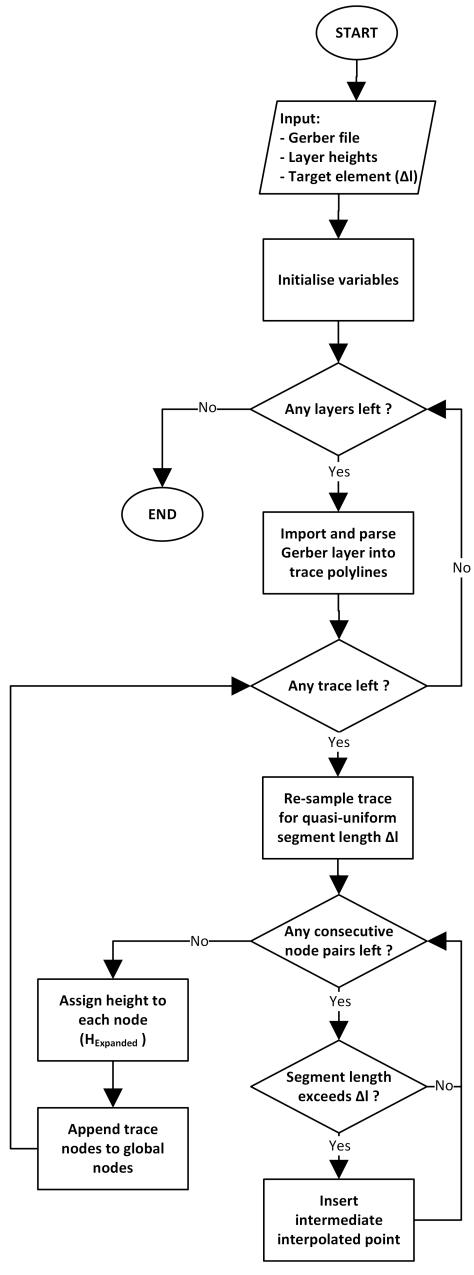


Figure 5.3: High-Level flowchart of trace geometry extraction.

exceeds the prescribed step, intermediate points are inserted by linear interpolation; otherwise, the original vertices are retained. This procedure produces an approximately uniform segmentation along the trace while preserving its geometric path, ensuring compatibility with the thin-wire MoM assumptions.

After re-sampling, three-dimensional positioning is assigned by associating each node with the height of its corresponding PCB layer. Terminal nodes may optionally be assigned a reference height (e.g., $z = 0$) when representing pads or endpoints tied to the reference plane. At this stage, a height vector aligned with the node list is constructed, assigning each non-delimiter node its appropriate vertical coordinate. This

guarantees consistency between the planar geometry extracted from the layout and the three-dimensional positioning required by the electromagnetic solver.

All discretized traces are then assembled into a single global node structure. Delimiter entries are inserted between successive traces to preserve trace boundaries within this unified representation. These delimiters allow downstream stages to identify individual traces and their terminal nodes directly from the geometry, without the need for additional indexing or bookkeeping structures.

The output of this stage is a geometry description that can be reused across multiple excitation conditions. In practice, the node coordinates and associated height information are stored in a structured file format, allowing subsequent stages of the modeling workflow to operate without repeating the layout discretization.

5.5.3 Impedance Computation

The impedance extraction stage provides the frequency-dependent terminal impedances required by the thin-wire MoM solver, evaluated on the same frequency grid used for the induced-field analysis. These impedances encode how each PCB trace is electrically terminated by the surrounding circuitry and are applied as boundary conditions when assembling the MoM system. The sequence of operations involved in the impedance computation stage is outlined in the flowchart presented in Figure 5.4.

Each circuit block connected to the PCB traces is assumed to have been characterized in a reference environment with characteristic impedance $Z_0 = 50\ \Omega$, consistent with standard RF modeling practice and with the operating conditions of most VNAs and S-parameter-based descriptions. In this setting, the small-signal voltage transfer function of a two-port can be expressed as the result of a Thévenin-equivalent voltage divider between the reference impedance and the intrinsic impedance of the block. Figures 5.5a and 5.5b illustrate the corresponding equivalent circuits used to recover the intrinsic input and output impedances from the corresponding transfer functions. This relationship allows the terminal impedances of each block to be obtained by direct algebraic inversion.

Within this workflow, the transfer functions associated with the blocks connected to each trace are treated as pre-existing circuit-level data obtained during prior system characterization. The role of this stage is therefore not to compute these transfer functions, but to convert them into the frequency-dependent impedance profiles required by the MoM formulation.

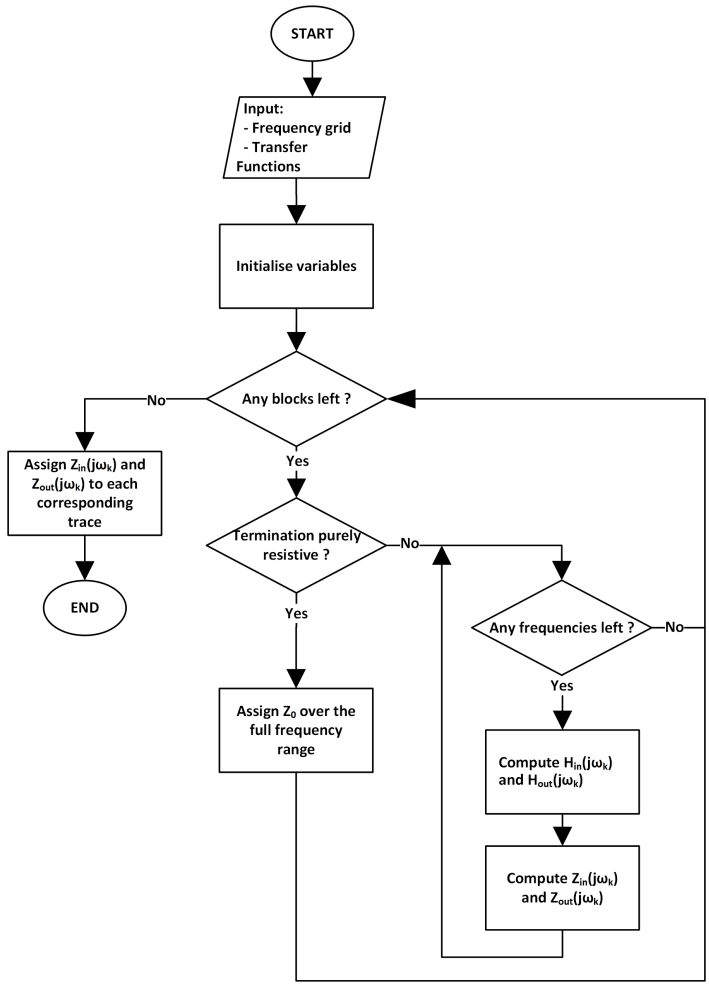


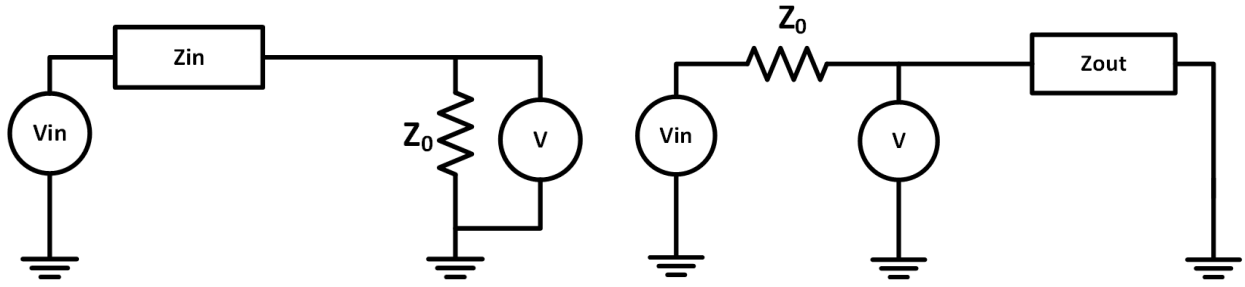
Figure 5.4: High-Level flowchart of impedance computation.

For each frequency sample f_k, with $\omega_k = 2\pi f_k$, the intrinsic input and output impedances are obtained by algebraic inversion of the Thévenin divider relations,

$$Z_{\text{in}}(j\omega_k) = Z_0 \frac{1 - H_{\text{in}}(j\omega_k)}{H_{\text{in}}(j\omega_k)}, \qquad Z_{\text{out}}(j\omega_k) = Z_0 \frac{H_{\text{out}}(j\omega_k)}{1 - H_{\text{out}}(j\omega_k)} \tag{5.12}$$

with $Z_0 = 50\,\Omega$. Circuit blocks that exhibit purely resistive terminations are assigned constant impedance values directly, while all other blocks are evaluated across the frequency grid to produce frequency-dependent profiles.

Once the intrinsic impedances of the circuit blocks have been obtained, they are mapped onto the discretized traces. For each trace, the impedance of the source-side block is applied at the first discretization element and that of the load-side block at the last, following the connectivity defined in the circuit schematic. This mapping closes the MoM problem with terminal conditions that are consistent with the



(a) Input impedance extraction (Z_{in}). **(b)** Output impedance extraction (Z_{out}).

Figure 5.5: Thévenin-equivalent circuits used to recover intrinsic terminal impedances from voltage transfer functions referenced to a characteristic impedance Z_0.

electrical environment seen by the trace.

For practical use, the resulting frequency-dependent impedance profiles can be stored and reused across simulation runs. This avoids repeated inversion of the transfer functions when the same circuit configuration is analyzed under multiple excitation conditions, and ensures consistency between the impedance data supplied to the induced-voltage computation stage.

5.5.4 Induced Voltage Computation

This stage applies the thin-wire Method of Moments (MoM) formulation introduced earlier to compute the voltages induced at the terminals of each discretized trace under external electromagnetic illumination. The equivalent-wire representation and the associated EFIE are taken as given; the focus here is on how that formulation is exercised in practice by assembling the discretized EFIE, applying circuit-level boundary conditions, and solving the resulting system across the frequency grid to obtain a physically meaningful trace-level electromagnetic response.

In principle, the electromagnetic response of a trace could be described by a multidimensional function $U_{i,k}(\theta, \phi, \gamma)$ indexed by trace index i, frequency sample k, incidence direction (θ, ϕ), and polarization γ. Exhaustively populating such a response would require solving the MoM system for a large number of illumination configurations, leading to significant computational cost and a large induced-voltage dataset. Rather than pursuing this exhaustive strategy, the present workflow exploits a key property of the formulation: linearity with respect to the incident field.

Accordingly, for each frequency sample index k, the MoM system is solved only for a reduced set of illumination conditions that form a complete basis for all in-plane excitations. Specifically, two orthogonal

in-plane polarizations, $\gamma \in \{0^\circ, 90^\circ\}$, are evaluated under a fixed broadside incidence $(\theta, \phi) = (0^\circ, 0^\circ)$. These baseline solutions are not intended to represent limiting cases by themselves; instead, they fully characterize how the trace responds to tangential electric-field components in the PCB plane at that frequency. Figure 5.6 provides an overview of the induced-voltage computation stage, highlighting the main processing steps used to assemble and solve the thin-wire MoM system under external illumination.

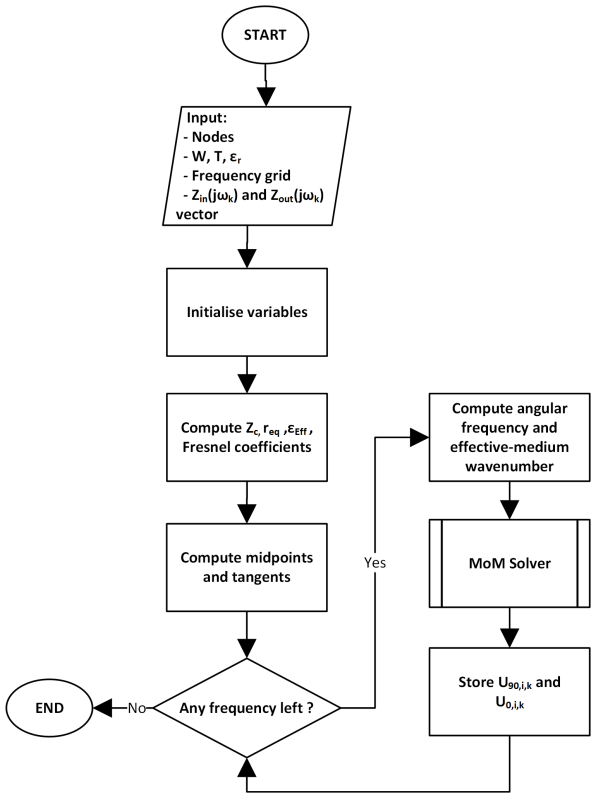


Figure 5.6: High-Level flowchart of induced voltage computation.

The solver takes as input the discretized geometry, the PCB parameters $\{W, T, H, \varepsilon_r\}$, and a set of layer-dependent quantities derived from them within this stage. These include the effective permittivity ε_{eff}, the characteristic impedance Z_c of the equivalent microstrip environment, and the equivalent wire radius a_{eq} used to regularise the thin-wire self-interaction terms. The incident-field model explicitly accounts for transmission at the air–substrate interface through TE/TM Fresnel coefficients and their grounded-slab extension, ensuring that the excitation applied to the MoM system corresponds to the field effectively present at the trace height.

Within this formulation, the EFIE is discretized using a Galerkin scheme, yielding a dense MoM system whose entries are evaluated using Gauss–Legendre quadrature, with analytic regularisation applied to the self-interaction terms. Terminal impedances derived from the surrounding circuitry are imposed as

boundary conditions at the trace endpoints, and the resulting linear system is solved for each frequency and baseline polarization to obtain the induced current distribution along the trace.

For each frequency and each baseline polarization, the MoM system is solved for the induced currents. Terminal voltages then follow directly from the complex product of the terminal currents and the corresponding port impedances. The resulting baseline induced-voltage profiles are denoted $U_{0,i,k}$ and $U_{90,i,k}$. By construction, they form a complete basis for the trace response to any in-plane polarization at that frequency. The induced voltage corresponding to an arbitrary polarization angle γ can therefore be reconstructed a posteriori by linear combination, without requiring additional electromagnetic solves.

$$U_{i,k}(\gamma) = U_{0,i,k} \cos\gamma + U_{90,i,k} \sin\gamma \tag{5.13}$$

5.5.5 Resonant Optimization

The purpose of the resonant optimization stage is to condense the induced-voltage results into a compact representation that retains only the dominant electromagnetic coupling scenarios. Rather than propagating all reconstructed voltages across frequency and polarization, this stage identifies, for each trace, the frequencies at which coupling is strongest and the corresponding polarization angles that produce those maxima. The resonant optimization procedure is summarized in the flowchart shown in Figure 5.7, which illustrates how worst-case coupling conditions are identified from the baseline trace responses.

The input to this stage consists of the two baseline induced-voltage profiles computed under broadside incidence for each trace, corresponding to the two orthogonal polarizations. Owing to the linearity of the MoM formulation with respect to the incident field, these two baseline responses fully characterize how the induced voltage varies with polarization angle and allow reconstruction for arbitrary in-plane orientations without additional electromagnetic solves.

To identify worst-case conditions, the polarization angle is discretized over a predefined set of candidate orientations $\{\gamma_m\}$. For each trace i and each frequency sample k, the induced voltage is reconstructed for all candidate polarizations, and the orientation that maximizes the voltage magnitude is identified according to

$$\gamma_{i,k}^{\max} = \arg\max_{\gamma_m} \left|U_{i,k}(\gamma_m)\right|, \qquad U_{i,k}^{\max} = \left|U_{i,k}\left(\gamma_{i,k}^{\max}\right)\right| \tag{5.14}$$

Here, $U_{i,k}^{\max}$ represents the largest induced voltage achievable on a given trace at a given frequency by rotating the polarization, while $\gamma_{i,k}^{\max}$ identifies the corresponding field orientation.

Resonant frequencies are then extracted by analyzing the worst-case induced-voltage envelope $U_{i,k}^{\max}$ along the frequency axis. A frequency index r is identified as resonant for trace i if it satisfies the local-maximum condition

$$U_{i,r-1}^{\max} < U_{i,r}^{\max} > U_{i,r+1}^{\max} \tag{5.15}$$

within the numerical tolerance associated with the frequency grid. Physically, such peaks correspond to frequencies at which the trace supports standing-wave patterns compatible with its electrical terminations and geometry, and at which coupling with the external field is particularly efficient.

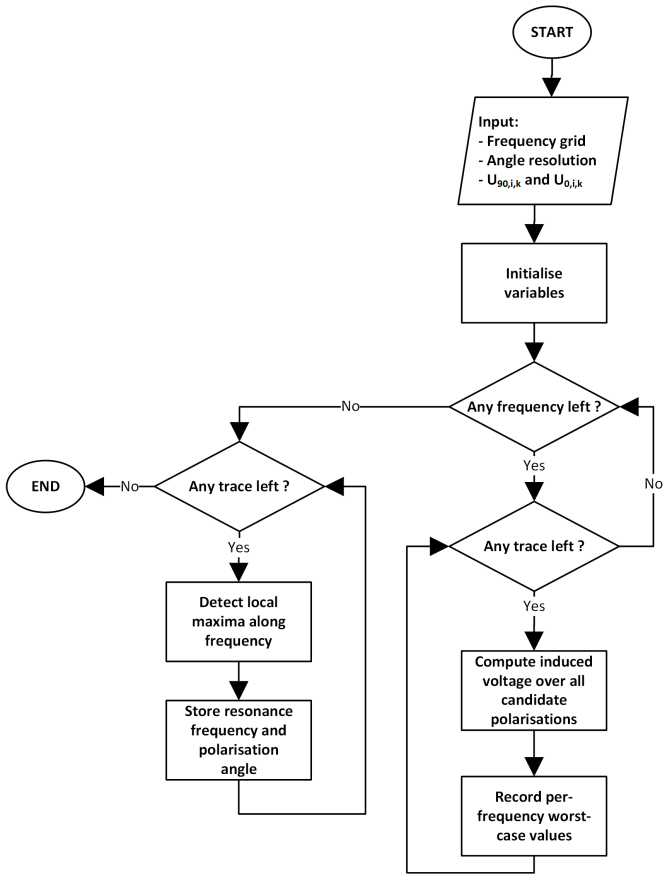


Figure 5.7: High-Level flowchart of resonant optimization.

For each detected resonance, only a minimal set of parameters is retained: the resonant frequency index r and the polarization angle $\gamma_{i,r}^{\max}$ that maximizes coupling at that frequency. Rather than storing the full

induced-voltage response across frequency and polarization, the optimization stage preserves only the resonance descriptors $\{ r,\ \gamma_{i,r}^{\max} \}$ for each trace.

This choice keeps the dataset compact while still allowing the exact worst-case excitation conditions to be recovered whenever needed. During run-time system-level simulation (i.e., in *Simulink*), the induced voltage associated with a given resonance is reconstructed on the fly from the stored resonance parameters and the baseline responses via the linear polarization combination defined in (5.13), rather than being retrieved from a large lookup table or a multidimensional response tensor.

This on-the-fly reconstruction is computationally lightweight, avoids file-access overhead, and provides a faithful yet efficient interface between the detailed electromagnetic behavior captured by the MoM solver and higher-level functional models. By retaining only physically dominant resonant conditions, the resonant optimization stage ensures that subsequent analyses focus on credible worst-case electromagnetic threats without unnecessary data volume or computational burden.

5.6 Case Study - Coupled-Coil Position Sensor

Having established the trace-level modeling framework and its integration within the A^{2}EM workflow, this section turns to a concrete application. The aim is to ground the previously introduced formulation in a representative sensor front-end and to demonstrate how the trace-modeling approach operates when applied to a real PCB implementation.

This section applies the trace-modeling framework developed in this chapter to the coupled-coil position sensor used as a case study throughout the thesis. Starting from the PCB implementation of the sensor front-end, it provides the layout- and sensor-specific information required to instantiate the modeling in practice and to demonstrate the resulting trace-level electromagnetic analysis on a real design. The electromagnetic entry points and the underlying modeling assumptions have been established earlier and are taken as given. The focus here is therefore on making explicit the practical details of the case study that are necessary to carry out the trace modeling and on illustrating how trace-level coupling is computed and characterized for the considered sensor front-end.

5.6.1 Trace Geometry Extraction

As established earlier in the thesis, the electromagnetic entry points considered in this work correspond to the PCB interconnecting traces that link clusters of functional components within the sensor front-end. This is illustrated in Figure 4.5 where a total of eleven traces were identified and selected as electromagnetic entry points for the trace-level modeling presented in this chapter.

In this step, the Gerber files of the sensor PCB are used to extract the routing of these eleven traces. The layout data provide the planar geometry of each interconnection on its corresponding copper layer and constitute the direct geometric input to the thin-wire modeling process. A graphical rendering of the extracted traces is shown in Figure 5.8 with top-layer (L1) routes in red and bottom-layer (L2) routes in blue distinguished for clarity.

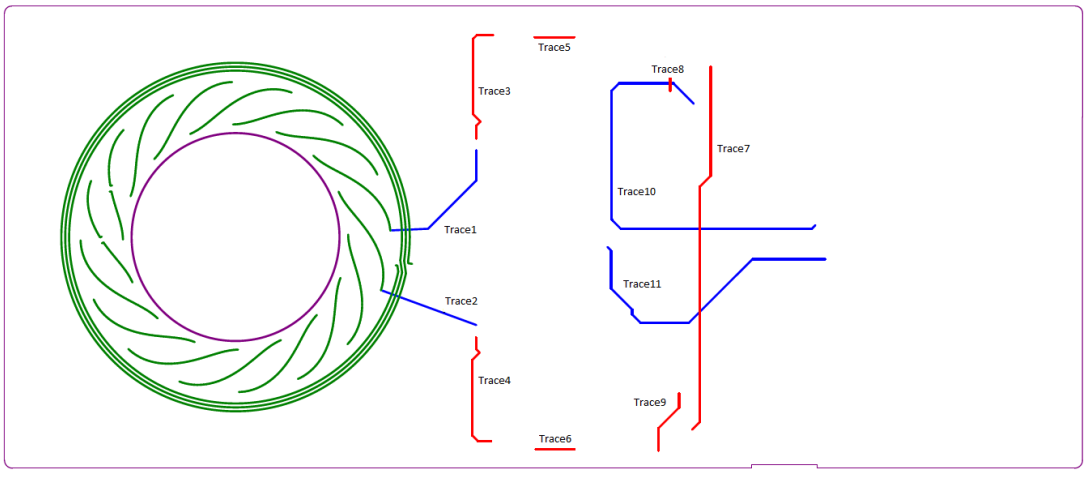


Figure 5.8: Graphical rendering of extracted electromagnetic entry-point traces from sensor PCB Gerber data.

Each extracted trace is converted into a polyline representation and discretized using a target element length of 1 mm. This discretization preserves the routing of all traces while keeping the resulting electromagnetic model computationally manageable. The resulting discretized geometry forms the basis for all subsequent trace-level computations.

5.6.2 Impedance Computation

In the impedance computation step, the electrical loading associated with each trace is determined from the schematic of the coupled-coil position sensor introduced in the operationalization of the A^2EM methodology. For each electromagnetic entry-point trace, the circuit blocks directly connected upstream and

downstream are identified, and their frequency-dependent behavior is described through voltage transfer functions obtained from prior circuit-level analysis.

These transfer functions are evaluated over the frequency range from 1 MHz to 10 GHz, using a resolution of 100 kHz. Across this frequency grid, the transfer functions are converted into equivalent intrinsic input and output impedances via standard Thévenin relations. Each trace is therefore assigned a pair of frequency-dependent terminal impedances that reflect the actual circuit environment surrounding it and that are used as boundary conditions in the subsequent trace-level electromagnetic analysis.

5.6.3 Induced Voltage Computation.

In this step, the thin-wire MoM solver is used to compute the voltages induced at the terminals of each discretized trace under external electromagnetic excitation. The physical parameters of the PCB stack-up used in the analysis are summarized in Figure 5.9. All traces are modeled with a width of 254 µm and a copper thickness of 35 µm. The dielectric height above the reference plane is set to 0.36 mm for top-layer (L1) traces and 0.20 mm for bottom-layer (L2) traces, with a substrate relative permittivity of 4.

These stack-up parameters, together with the discretized trace geometry and the frequency-dependent terminal impedances obtained in the previous steps, are provided as inputs to the thin-wire solver. The external excitation is modeled as a far-field plane wave with an incident electric-field amplitude of 20 V/m.

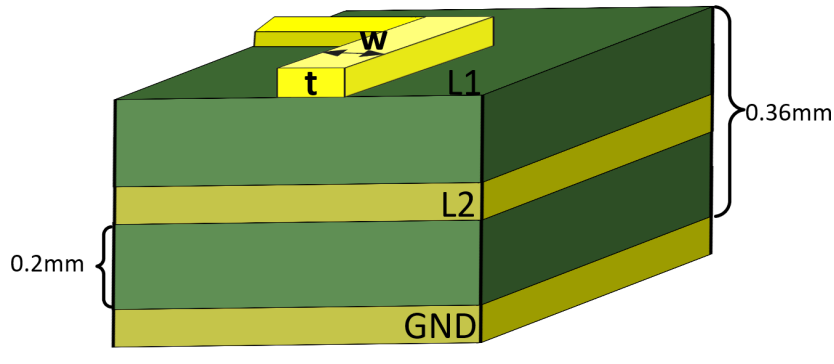


Figure 5.9: PCB stack-up used in trace-level electromagnetic modeling.

5.6.4 Resonant Optimization

The induced-voltage responses obtained for each trace are further analyzed to identify the frequency regions where electromagnetic coupling becomes dominant. This analysis focuses on extracting, from

the reconstructed induced-voltage responses, the excitation conditions associated with resonant behavior along the traces.

For each trace, the induced voltage is reconstructed as a function of polarization angle, and the maximum voltage magnitude is retained at each frequency sample. This yields, for every trace, a worst-case induced-voltage envelope $U_{\max}(f)$ together with the corresponding polarization angle that maximizes coupling at each frequency.

The resulting worst-case envelopes are then examined to identify resonant behavior. Rather than retaining the full dense frequency response, the $U_{\max}(f)$ curves are scanned to detect significant local maxima. These local peaks correspond to frequencies at which the trace exhibits enhanced sensitivity to external electromagnetic excitation. Figure 5.10 illustrates this process for two representative traces (Traces 10 and 11), showing the worst-case induced-voltage envelopes together with the automatically detected resonant peaks.

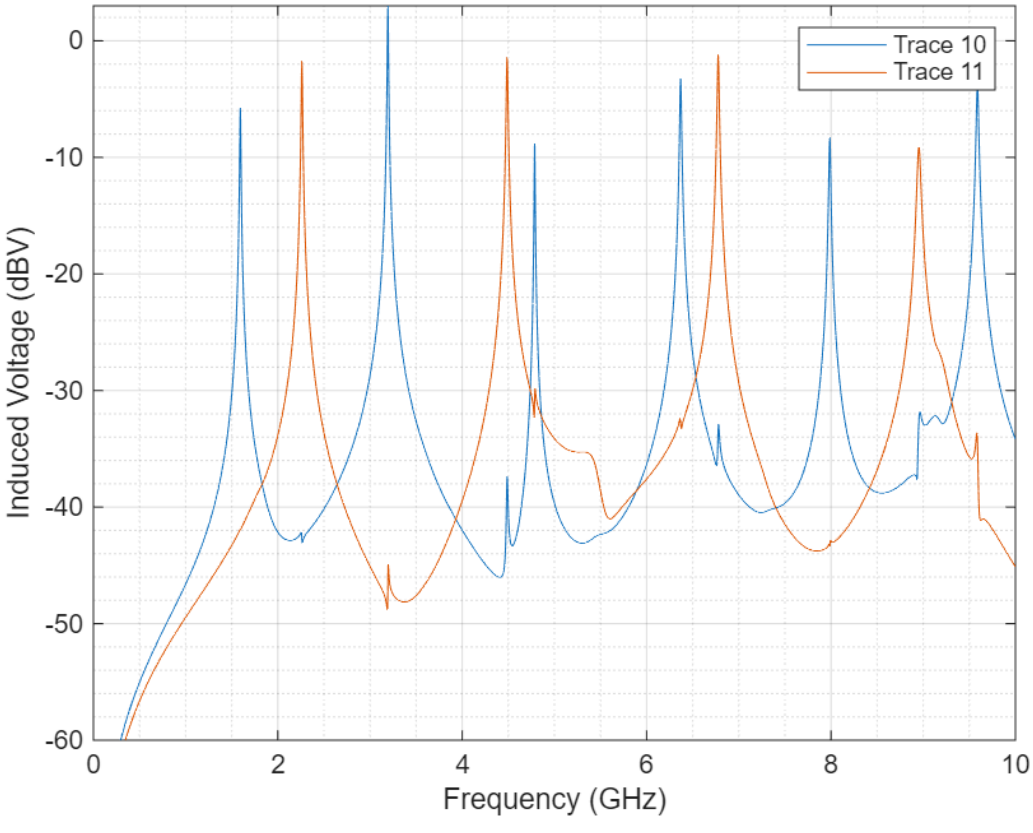


Figure 5.10: Worst-case induced-voltage envelopes $U_{\max}(f)$ for Traces 10 and 11 across the analyzed frequency range. The highlighted peaks indicate resonant frequencies identified by the local-maximum detection procedure.

From the set of resonant frequencies identified through the local-maximum analysis, the absolute maximum of $U_{\text{max}}(f)$ was also identified for each trace as an additional point of reference. This step is not required by the modeling workflow itself, but is included to provide the reader with a compact overview of the most prominent resonance per trace. Table 5.1 reports, for all eleven traces, the frequency at which the highest induced voltage occurs together with the corresponding polarization angle.

Table 5.1: Absolute maximum induced-voltage resonance identified for each electromagnetic entry-point trace.

Trace ID	Frequency (GHz)	Polarization Angle (deg)	Induced Voltage (V)
1	4.968	137	0.0465
2	6.348	21	0.0493
3	5.472	91.5	0.0650
4	5.563	88	0.0660
5	10.000	0	0.0148
6	10.000	0	0.0139
7	8.951	133	0.2038
8	3.194	155.5	0.1548
9	9.919	118.5	0.1029
10	3.194	155.5	1.4126
11	6.776	62	0.8683

5.7 Summary

This chapter developed a trace-modeling workflow that reconstructs PCB interconnecting traces from layout data, applies electrical boundary conditions derived from the surrounding circuitry, and computes the electrical response induced by external electromagnetic fields. The resulting responses are reformulated as equivalent electrical excitations, enabling direct integration into the system-level modeling environment introduced in the previous chapter.

To derive these excitation conditions in a physically grounded yet workflow-compatible form, a thin-wire Method-of-Moments formulation was adopted. Although higher-fidelity electromagnetic approaches exist, the selected formulation provides a practical balance between physical grounding and computational tractability for systematic trace-level analysis within A^2EM.

The proposed approach was applied to the coupled-coil position sensor used throughout the thesis. For the set of interconnecting traces identified as electromagnetic entry points, the modeling produced induced-voltage responses and revealed resonance-driven coupling conditions indicative of susceptibility hotspots

in the implemented PCB. The functional impact of these trace-level excitation scenarios is evaluated in the next chapter, where the corresponding equivalent excitations are propagated through the complete system model under the A^{2}EM workflow.

A natural extension of this modeling layer is to broaden the set of entry-point families exercised beyond PCB interconnecting traces. Additional structures such as vias, connectors, harness segments, and packaging or enclosure discontinuities may dominate coupling in other implementations and frequency regimes, and would require adapted geometric representations and boundary conditions. In parallel, where coupling geometries or operating frequencies exceed the practical validity range of the thin-wire MoM abstraction, higher-fidelity electromagnetic solvers and richer coupling representations could be incorporated to increase physical representativeness. These extensions would not change the role of the chapter within the A^{2}EM workflow, but would expand the range of coupling configurations and entry points for which physically derived excitation conditions can be produced and propagated at system level.

In the context of the thesis, this chapter completes the electromagnetic coupling layer needed to operationalize A^{2}EM beyond purely abstract disturbance assumptions. With trace-level coupling now explicitly represented and quantified, the methodology can be exercised using excitation conditions anchored in the physical implementation, enabling a consistent transition from architecture-aware vulnerability reasoning to physics-informed system-level evaluation and interpretation.

Chapter 6

A^2EM Results

6.1 Introduction

The previous chapters established the methodological and modeling foundations required for the operational use of A^2EM. This chapter builds on that foundation by evaluating the sensor and the analog front-end under study through the application of the *EMI Probing and Profiling* phase of A^2EM. The conceptual structure of the methodology was introduced in Chapter 3, while the modeling substrate and the identification of the electromagnetic entry points were addressed in Chapter 4. The characterization of these entry points was subsequently developed in Chapter 5, where the corresponding excitation conditions were derived.

In accordance with the abstraction levels introduced in the A^2EM chapter, EMI Probing and Profiling is conducted under progressively increasing levels of internal observability, corresponding to black-box, gray-box, and white-box analysis regimes. The objective is to progress from the identification of functionally permeable excitation conditions to a refined understanding of disturbance propagation and, ultimately, to the attribution of a structural root cause, followed by the selection, implementation, and validation of appropriate countermeasures, as previously discussed.

Although the A^2EM methodology supports the evaluation of multiple classes of functional effects under electromagnetic excitation, the present case study deliberately focuses on permeability mapping, that is, on the identification of excitation conditions that remain functionally observable at the system output. Rather than attempting to isolate or validate a specific effect class, the analysis first establishes which resonance-informed disturbance conditions produce non-negligible angular deviation. This reflects a pragmatic and conservative instantiation of the A^2EM framework, since it provides a clear basis for the

subsequent stages of interpretation, attribution, and countermeasure assessment, while remaining fully compatible with the broader scope of the methodology.

6.2 Black-box Testing

This section presents the results obtained under the black-box analysis regime, which constitutes the first level of functional evaluation within the EMI Probing and Profiling phase, shown in Figure 3.2. At this level, the objective is not to localize failure mechanisms or assign causality, but to identify electromagnetic excitation conditions that are functionally permeable, that is, capable of inducing a significant and observable deviation at the system output.

The black-box analysis builds on the resonant excitation space obtained during the Resonant Entry Point Mapping stage, which yields 98 excitation scenarios grouped into 18 distinct frequency bands. As established earlier, this stage is strictly preparatory and serves to constrain the electromagnetic excitation space based on physically grounded coupling considerations. The Resonant Optimisation step introduced at the end of the trace modeling stage then condenses the frequency-dependent coupling behavior of the identified entry points into a reduced set of excitation conditions. This eliminates the need for exhaustive frequency sweeps, such as Monte Carlo or brute-force approaches, while enabling a targeted exploration of functionally relevant regimes. In the black-box regime, excitation is applied at the circuit level, such that all identified entry points are exposed simultaneously under the same incident-field conditions, as would occur in a real operating environment.

The analog sensor front-end analysed is specified to operate with an angular error below 1° under nominal conditions. For this reason, permeability is defined as $|\Delta\theta| \geq 1^\circ$. After executing the injection campaign, applying this criterion classifies 78 scenarios as functionally permeable. Table A.1 in Appendix A lists these 78 cases, together with the corresponding trace, exact frequency, polarization angle, and observed deviation.

A first observation is the strong spectral selectivity of the functional response. Although the excitation space is constructed from the obtained resonant data, only a subset of the identified bands produces deviations beyond the 1° threshold. Within this subset, permeability is highly non-uniform. The band around **3.18–3.20 GHz** clearly dominates the error landscape, combining the broadest trace coverage (**10 out of 11 traces**) with the largest deviations (up to approximately **50**$^\circ$). This indicates that, under

black-box observability, this spectral region constitutes the most critical regime in terms of both recurrence and severity.

Beyond this dominant region, several bands exhibit broad trace coverage while remaining less severe in magnitude. In particular, bands at **1.57–1.60 GHz**, **2.24–2.27 GHz**, **6.33–6.39 GHz**, and **6.74–6.79 GHz** each induce non-negligible deviation across **seven traces**, while **7.97–8.00 GHz** affects **nine traces**. These results show that permeability is not confined to a single isolated resonance, but instead appears in multiple well-defined spectral regions that recur across a large fraction of the entry-point set. At the black-box level, this recurrence cannot be interpreted causally; rather, it indicates that the system-level output can be driven into a degraded regime under incident-field conditions that are compatible with multiple coupling paths being active concurrently.

A second observation concerns the role of polarization. Several permeable bands exhibit a wide dispersion of effective polarization angles across traces and frequencies. This is most apparent in the dominant **3.18–3.20 GHz** band, but is also visible in **1.57–1.60 GHz**, **2.24–2.27 GHz**, and **6.74–6.79 GHz**, where non-negligible deviation occurs under substantially different field orientations depending on the excitation scenario. In contrast, other bands display more constrained angular behavior. The band at **9.57–9.60 GHz**, for example, produces deviations up to approximately **19**$^\circ$ with low polarization angles. At the opposite end of the spectrum, **9.12–9.18 GHz** appears as a near-threshold permeable region, with deviations only marginally above 1° and limited trace coverage (**two traces**), and therefore represents the weakest permeable regime retained under the adopted criterion.

Finally, the set of permeable bands includes intermediate cases that remain clearly non-negligible yet more selective than the dominant regions. The band at **4.46–4.50 GHz** affects **six traces** with deviations reaching approximately **25**$^\circ$, while the narrow band at **4.78–4.79 GHz** affects **five traces** with moderate deviations (around **7.5**$^\circ$**–10**$^\circ$). Together with **8.90–8.97 GHz** (**six traces**, up to approximately **8**$^\circ$), these bands reinforce that permeability can manifest either as high-severity, broadly recurring behavior, or as more selective regimes restricted to fewer traces and/or more limited polarization subsets.

Overall, the black-box analysis identifies **11 permeable frequency bands out of 18** under the $|\Delta\theta| \geq 1^\circ$ criterion. While only a subset of these bands exhibits dominant and recurrent error-inducing behavior, all 11 constitute resonance-consistent excitation regimes capable of inducing functionally non-negligible deviation at the system output. They are therefore retained as candidate regimes for subsequent analysis

under increased levels of observability, where causal attribution and mechanism-level explanation can be pursued.

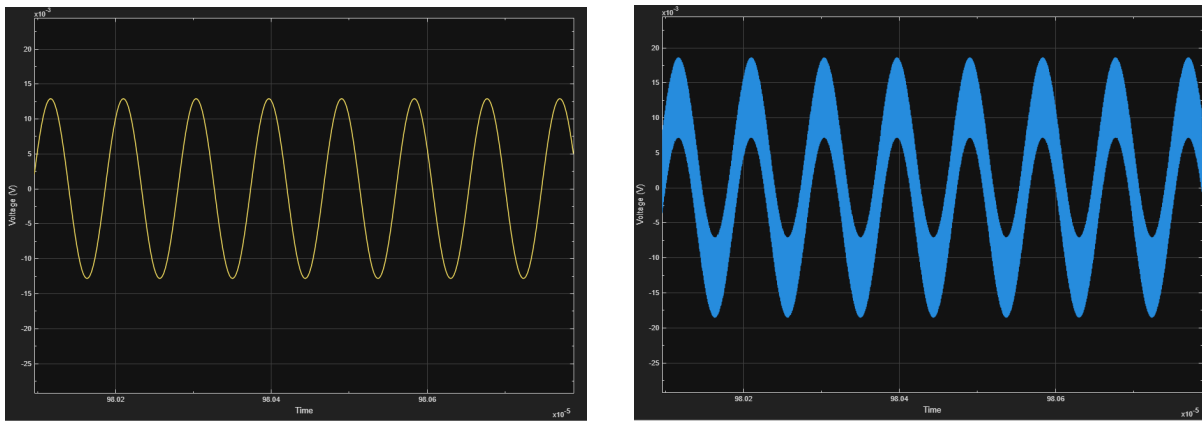
6.3 Gray-box Testing

This section describes the gray-box simulation model used to interpret the permeable excitation conditions identified in the preceding black-box regime. The objective is to clarify the functional role of each analog stage, the frequency constraints it introduces, and how these constraints shape disturbance propagation towards the sampled output. Although multiple permeable excitation conditions were identified during black-box testing, the gray-box analysis is anchored to the most permeable case, namely the single-tone disturbance at 3.1939 GHz with polarization angle $155.5°$, which produced the highest angular deviation of $50.12°$. This choice is deliberate: the purpose of the gray-box phase is not to provide exhaustive coverage of all permeable bands, but to interpret the dominant propagation and conversion mechanisms under a representative high-impact condition, thereby avoiding redundant repetition of stage-level interpretation across multiple bands that share comparable coupling characteristics. In other contexts, additional permeable cases could be examined to quantify variability across bands or to probe alternative mechanisms.

The signal chain begins with a transduction abstraction that maps the mechanical angle into two quadrature channels. The input angle drives sine and cosine branches, forming the Q and I components, respectively, as illustrated in Figure 4.6. Each component is modulated by a common carrier at 10.7 MHz, producing two carrier-centered waveforms whose quadrature relationship encodes the angular position.

The behavior at the output of the *HCoil* stage is shown in Figure 6.1. Figure 6.1a presents the nominal carrier-centred waveform in the absence of coupling, whereas Figure 6.1b shows the same node with the disturbance activated at the first downstream entry point, confirming the superposition of the injected component on the carrier envelope.

Immediately after transduction, the signal enters the first impedance-matching stage, *HZBMatch*. This block exhibits predominantly high-pass behavior with a cutoff frequency of approximately 6.65 MHz, attenuating components below this threshold while preserving the carrier region around 10.7 MHz. Since this selectivity is confined to the low-frequency end, it does not act as a rejection mechanism for disturbances in the GHz range. Consequently, the injected disturbance remains present at the *HZBMatch* output as a superposition on the nominal carrier-centred waveform.



(a) HCoil nominal output.

(b) HCoil output with injected disturbance.

Figure 6.1: HCoil output under nominal and disturbed conditions.

The next functional constraint is the ceramic bandpass filter, *HBandPass*, centred at 10.7 MHz with a pass-band of ± 75 kHz, corresponding to an effective bandwidth of approximately 150 kHz. As shown in Figure 6.2, this stage enforces narrowband selection around the carrier and attenuates out-of-band content. Consequently, GHz-range disturbances injected upstream of *HBandPass* are strongly suppressed at its output, and the remaining signal content is confined to the carrier band around 10.7 MHz. After *HBandPass*, the time-domain waveform closely resembles the nominal carrier signal, with no visible high-frequency superposition remaining.

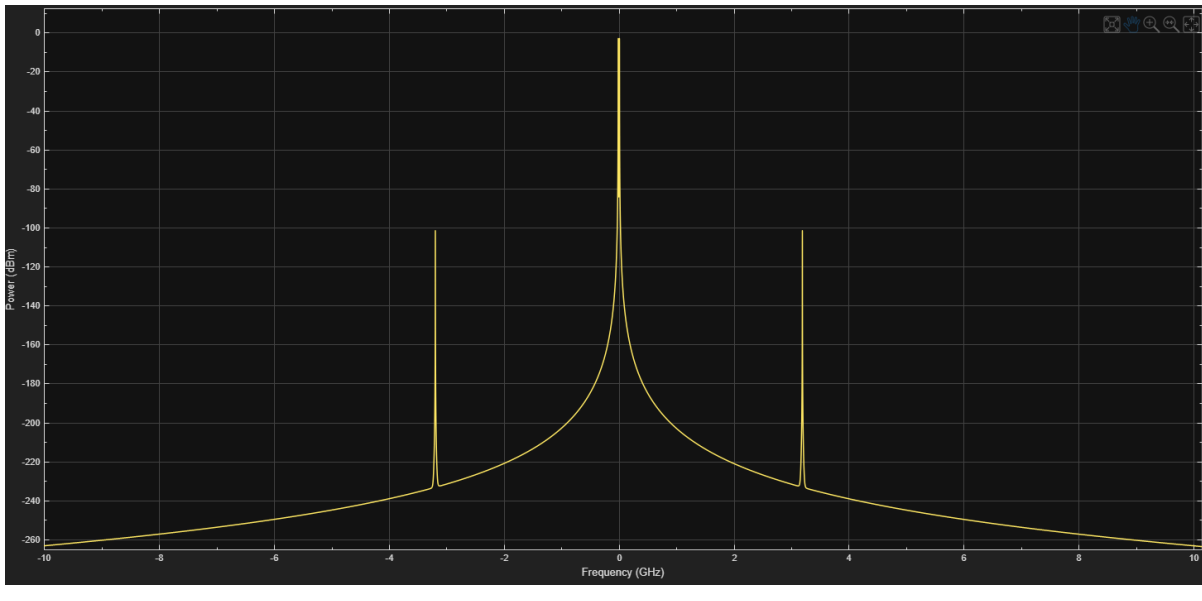


Figure 6.2: Spectral content of the bandpass output.

After the bandpass filtering, demodulation is implemented as a synchronous mixing between the conditioned signal path and a locally generated reference. This demodulation subsystem comprises three

coupled elements: the pre-mixer matching network (*HZMMatch*), the oscillator generator (*Mixer_OSC*), and the mixer itself (*HMixer*).

The signal path enters *HZMMatch* immediately after the ceramic filter. *HZMMatch* acts as a matching network with a low-pass behavior, with an effective cutoff frequency of approximately 17 MHz. Above this range, its attenuation is nearly -60 dB/decade, so that disturbance components in the GHz regime are strongly reduced before reaching the demodulation block through the signal input. This behavior is shown in Figure 6.3.

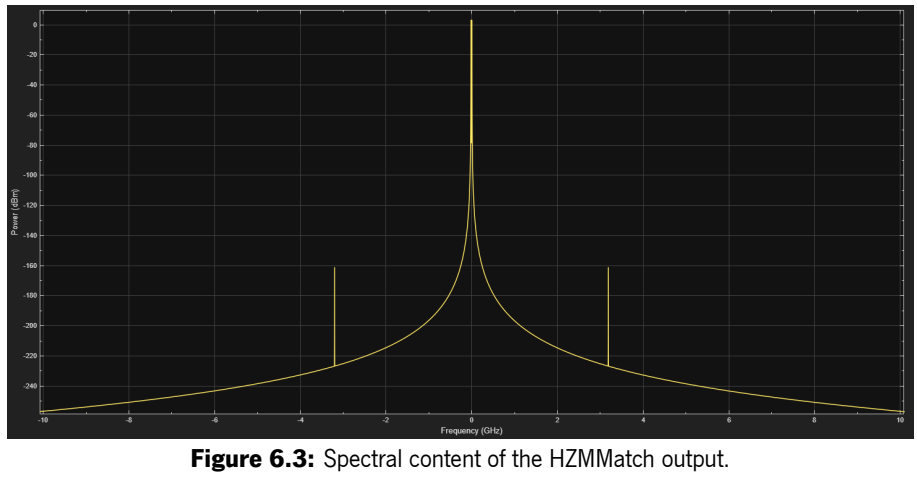


Figure 6.3: Spectral content of the HZMMatch output.

Mixer_OSC generates the local reference used for synchronous demodulation, while *HMixer* mixes this reference with the conditioned carrier-centred signal. In general terms, nonlinear operations such as mixing can, under some conditions, map out-of-band content into the low-frequency region that carries the demodulated information. For this reason, the demodulation subsection is a natural candidate to examine when considering whether interference could become observable in baseband.

In the present gray-box analysis, the observed output deviation around this subsection initially suggested that an IMD-like conversion mechanism could be involved. Stage-local inspection, however, does not support this interpretation for the excitation conditions considered. The dominant reason is the pronounced high-frequency rejection provided by *HZMMatch*, which strongly attenuates injected spectral components before they can reach *HMixer* and participate in any effective down-conversion into the baseband region. This behavior is consistent with the spectrum observed at the *HMixer* output in Fig. 6.4.

While this evidence argues against an IMD-driven explanation for the deviations observed in this case study, a dedicated IMD-oriented assessment would be required to exclude corner cases outside the present excitation set and operating conditions.

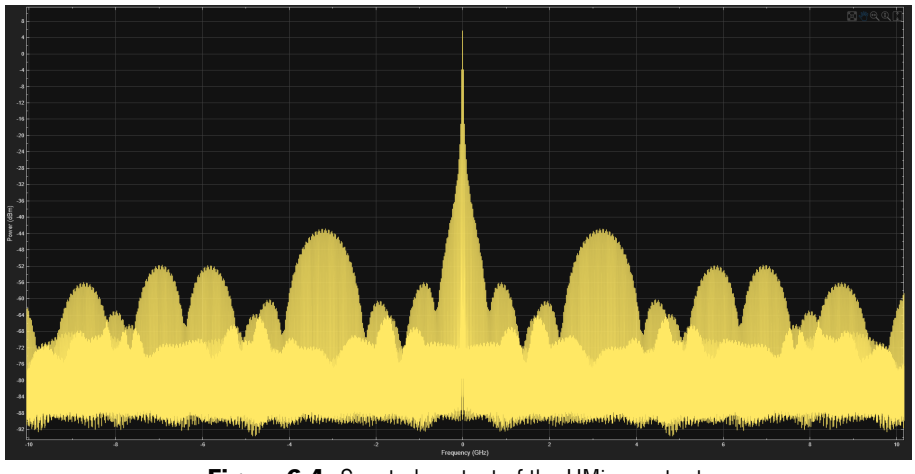


Figure 6.4: Spectral content of the HMixer output.

The final analog stage, *HAmpLPF*, implements an amplifier followed by a second-order low-pass filter with a cutoff frequency of approximately 100 Hz. It defines the effective baseband bandwidth of the demodulated I/Q channels and finalises the transition from the mixed-domain output of *HMixer* to a low-frequency signal suitable for sampling. The high-frequency component generated at the *HMixer* output by the mixing process remains present at the input of *HAmpLPF*, but it is strongly attenuated by the 100 Hz low-pass response. This behavior is consistent with the spectral content observed after *HAmpLPF*, as illustrated in Figure 6.5, where the output is confined to the low-frequency band.

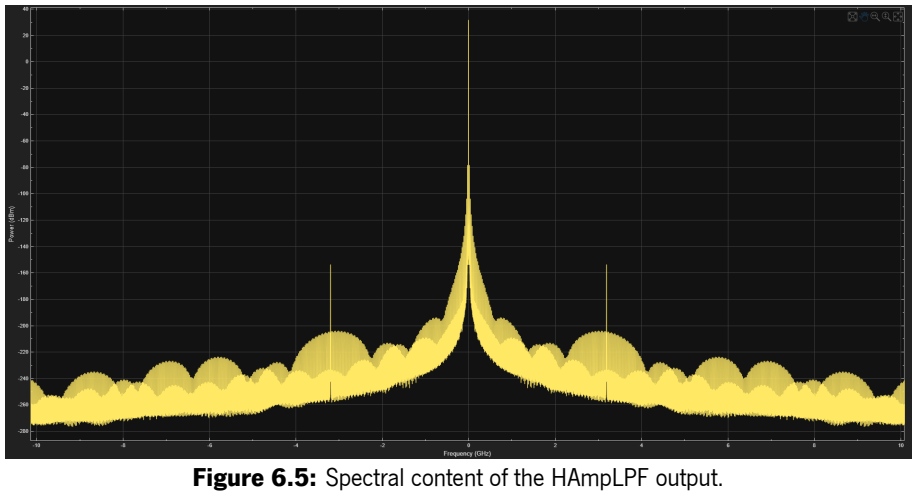


Figure 6.5: Spectral content of the HAmpLPF output.

From this point onward, the chain is constrained by a clear frequency limit: components above 100 Hz are attenuated by -40 dB/decade, whereas any disturbance that has already been converted into low-frequency content within the pass-band is preserved and conditioned. This makes *HAmpLPF* the last stage that can remove residual high-frequency content before sampling.

The last entry point is located on the final interconnect between the *HAmpLPF* outputs and the ADC inputs. Because this coupling location is positioned after the baseband conditioning block, it is not subject to the 100 Hz filtering action of *HAmpLPF*. Disturbances injected at this final interconnect therefore act directly on the sampled I/Q values and remain observable at the ADC input. This behavior is reflected in the ADC-side spectrum shown in Figure 6.6. At this stage, the sampling-interface interconnects emerge as a plausible dominant mechanism, since coupling at this location bypasses the final analog bandwidth constraint. Whether this entry point alone is sufficient to reproduce the observed system-level deviation, and thus constitutes the structural root cause relative to any residual upstream contributions, is assessed in the subsequent white-box testing phase.

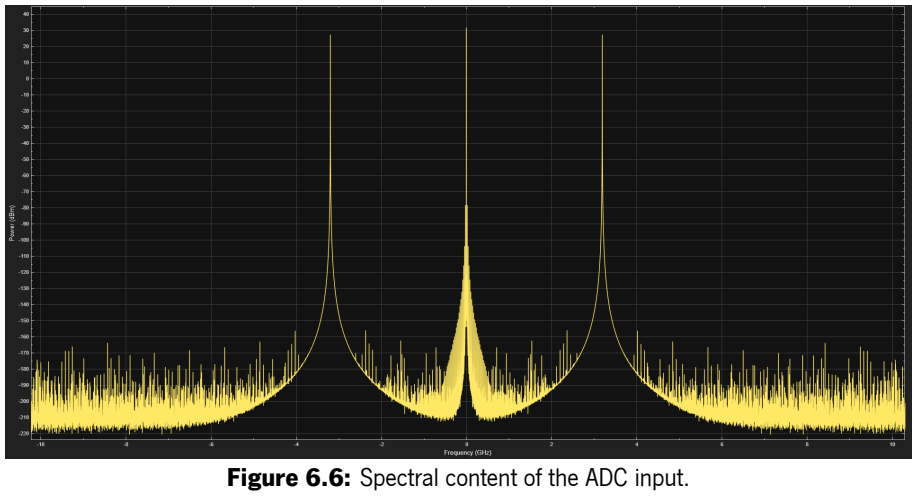


Figure 6.6: Spectral content of the ADC input.

6.4 White-box Testing

The white-box regime extends the gray-box analysis by shifting from stage-level interpretation to explicit causal attribution. While the gray-box results clarify how frequency constraints govern disturbance propagation across the analog chain, they do not determine which specific entry point dominates the observed angular deviation when multiple coupling locations are active simultaneously. White-box testing addresses this limitation by exploiting the controllability of the simulation environment to selectively enable and disable coupling sources. The objective is to verify whether coupling on the ADC-side interconnects is in fact sufficient to reproduce the observed deviation and thus constitutes the structural root cause, and to assess their relative criticality by determining whether one of the two ADC-side traces dominates the reconstructed error or whether both contribute significantly depending on the excitation condition.

It is acknowledged that this isolation strategy is not physically representative. The induced voltages computed via the Method of Moments inherently account for electromagnetic interactions between conductors; therefore, disabling entry points while retaining their originally computed induced voltages breaks that mutual coupling consistency. From a strictly physical standpoint, isolating individual traces under fixed induced-voltage conditions does not correspond to a realizable exposure scenario. Consequently, the results of this white-box phase must be interpreted in relative rather than absolute terms. The objective is not to replicate a physically exact field configuration, but to determine whether one ADC-side trace is intrinsically more critical than the other and whether coupling at the sampling interface alone is sufficient to reproduce the observed deviation.

Two excitation conditions were selected for this phase by pairing, for each ADC-side interconnect, the excitation instance that simultaneously (i) maximizes the induced voltage on that trace and (ii) produces the largest associated angular deviation in the distributed-injection baseline. The selection was further constrained to ensure that the two conditions occupy distinct frequency regions, thereby avoiding overlap between the excitation bands used to probe Trace 10 and Trace 11 and preventing ambiguity in the attribution. Under this criterion, Trace 10 is evaluated at 3.1939 GHz with polarization angle 155.5° (distributed-injection deviation 50.12°), while Trace 11 is evaluated at 6.7764 GHz with polarization angle 62° (distributed-injection deviation 26.7°). For each condition, three injection scenarios are exercised: (i) only Trace 10 enabled, (ii) only Trace 11 enabled, and (iii) both traces enabled simultaneously, enabling a direct comparison of their relative contributions and of the extent to which their combined activation reproduces the distributed case.

For the excitation corresponding to the maximum of Trace 10 (3.1939 GHz at 155.5°), activating only Trace 11 yields a deviation of 0.08816°; activating both traces yields 50.09°; activating only Trace 10 yields 50.12°. The deviation observed when only Trace 10 is active reproduces the distributed-injection case to within numerical tolerance, whereas Trace 11 alone has a negligible impact under this condition.

For the excitation corresponding to the maximum of Trace 11 (6.7764 GHz at 62°), activating only Trace 11 yields 25.95°; activating only Trace 10 yields 0.5186°; activating both traces yields 26.09°. In this case, Trace 11 alone accounts for almost the entire deviation, while Trace 10 produces only a minor contribution.

Across both excitation conditions, the combined activation of Trace 10 and Trace 11 closely matches the deviations observed under distributed injection. This supports the interpretation that, for the permeable conditions selected in black-box screening and interpreted in the gray-box stage, the dominant contribution

to the reconstructed error originates at the ADC-side interconnects. The results also enable a relative criticality assessment. Trace 10 exhibits the largest absolute deviation overall and is therefore the most critical entry point in terms of worst-case impact. Nevertheless, Trace 11 independently produces substantial deviations under its own critical excitation condition. Accordingly, mitigation cannot be limited to a single trace: both ADC-side interconnects must be addressed, since each can dominate the reconstruction error under specific excitation parameters.

Overall, the white-box analysis refines the black-box outcome and the gray-box interpretation by localizing the dominant coupling region at the sampling interface. For the excitation conditions investigated, the reconstructed angular deviation is primarily driven by the ADC-side interconnects, with upstream circuitry contributing only marginally. The ADC interconnect region is therefore established as the root cause of the observed functional permeability under the identified worst-case conditions, providing a clear target for the corrective design actions introduced in the subsequent mitigation phase.

6.5 Mitigation

This section transitions from vulnerability attribution to corrective design action. The mitigation phase progresses from a brief discussion of the countermeasure classes applicable to the identified root cause to the selection and implementation of a concrete strategy. Candidate options span physical interventions on the coupling structure itself and interface-level or signal-chain measures that reduce the susceptibility of the reconstructed output to coupled perturbations. Once a mitigation strategy is implemented, its effect is validated within the same injection and evaluation framework used during vulnerability discovery. The system is first revisited under a white-box approach to assess the countermeasure directly at the attributed coupling location and under the previously identified worst-case excitations. It is then re-evaluated in a black-box regime by reapplying the broader excitation set used during screening. This second step verifies that the mitigation generalizes across the wider set of previously permeable regimes and does not introduce unintended artifacts or degrade nominal functional behavior under other excitation conditions.

6.5.1 Mitigation options

The selection of a mitigation strategy is inherently a trade-off decision: effectiveness must be weighed against integration impact, redesign effort, qualification burden, and the extent to which a countermeasure can be applied at the identified coupling location. From the perspective of current guidance, the most direct physical countermeasures for the present vulnerability are those that harden the sampling interface itself, i.e., the ADC-side interconnect region. IEEE Std 1848-2020 frames such actions under good electromagnetic engineering and boundary control, including disciplined PCB layout and the use of shielding and filtering at controllable interfaces [96]. When instantiated for the sensor under study, this yields a small set of practically applicable options: (i) shortening (and rerouting) the ADC-side interconnects to reduce the exposed coupling structure, (ii) shielding the sampling-interface region to reduce field-to-trace coupling, and (iii) filtering as a boundary-conditioning measure (implemented either in hardware at the interface or in the digital domain after sampling).

For the present proof-of-concept, mitigation is explored along a complementary signal-oriented axis by applying filtering in the digital domain. This choice does not imply that digital filtering is universally preferable; rather, it provides a low-intrusion countermeasure that can be instantiated without modifying the coupling geometry. Depending on system constraints and lifecycle context, a different implementation may favor physical hardening at the sampling interface (shielding and layout revision) or a combination of physical and signal-oriented measures. In accordance with the described signal-oriented mitigation strategy, a post-ADC Infinite Impulse Response (IIR) low-pass filter was implemented. The corresponding implementation and evaluation is presented in the following sections.

6.5.2 Implemented strategy: post-ADC digital filtering

The mitigation implemented in this work consists of a digital low-pass filter applied immediately after the ADC. This choice is driven by three practical considerations. First, it avoids redesign of the PCB and of the ADC-side interconnects previously identified as the dominant coupling region, thereby eliminating direct impact on manufacturing changes and requalification effort. Second, it enables rapid and reproducible what-if exploration within the established simulation environment, because the mitigation can be introduced without modifying the identified coupling geometry. This preserves direct comparability with the pre-mitigation baseline: the excitation conditions and coupling model remain unchanged, and only the post-processing applied to the sampled I/Q streams is varied. Third, it acts directly in the domain where

the functional effect is observed, namely the reconstructed angle, reinforcing the separation between the application-relevant low-frequency content and any disturbance-induced components that become observable after sampling. This return to the evaluation loop, and the use of white-box controllability to support mitigation assessment, is consistent with the role assigned to each regime in the preceding results chapter (black-box for functional screening, gray-box for physically plausible stage-level interpretation, and white-box for causal attribution and mitigation evaluation).

From a modeling perspective, the block implemented in Simulink corresponds to a second-order IIR low-pass filter realized in direct form using two unit delays and constant gains derived from the ADC sampling frequency. In this work, the ADC sampling frequency is 101 kHz and the cutoff frequency (f_c) is 100 Hz. The coefficient is parameterized as

$$a = e^{-2\pi f_c / f_s} \tag{6.1}$$

where f_s denotes the sampling frequency. With this parametrisation, the filter can be written as an all-pole discrete transfer function with unit DC gain:

$$H(z) = \frac{(1-a)^2}{1 - 2a\, z^{-1} + a^2\, z^{-2}} \tag{6.2}$$

This realisation corresponds to a second-order IIR low-pass filter with a repeated real pole at $z = a$, i.e., a pole-placement design rather than a classical Butterworth or Chebyshev prototype based on complex-conjugate pole pairs. The repeated real-pole form is appropriate here because it provides a monotonic response and second-order asymptotic attenuation while retaining a simple and stable implementation for $0 < a < 1$.

The cutoff frequency is selected in the low-frequency regime, consistent with the application-relevant bandwidth. This therefore suppresses disturbance-induced components outside the useful band that, when coupled at the sampling interface, can become observable as angular deviation after sampling and estimation.

6.5.3 Validation: direct comparison against the pre-mitigation baseline

To quantify the impact of the implemented mitigation, validation was conducted through a direct one-to-one comparison against the pre-mitigation baseline. Exactly the same excitation cases (frequency bands, exact

frequencies, and polarization angles) used to populate the original permeability table were retained, even when the post-mitigation deviations fell below the 1° screening threshold. This ensures strict comparability: the effect of mitigation is assessed under identical excitation conditions, without redefining or filtering the test set.

In accordance with the mitigation workflow presented previously, validation is performed in two steps. The first step revisits the system under a white-box approach, leveraging the controllability of the simulation environment to assess the countermeasure directly at the root-cause location. The most severe pre-mitigation cases were explicitly re-evaluated. The excitation at 3.1939 GHz corresponding to the maximum deviation of Trace 10 (pre-mitigation $\approx 50^\circ$) now produces a residual deviation of 0.02061°. Similarly, the excitation at 6.7764 GHz corresponding to the maximum deviation of Trace 11 (pre-mitigation $\approx 26.7^\circ$) now produces a residual deviation of 0.04504°. The second step extends the evaluation back to a black-box regime, reapplying the broader excitation set used during the initial screening stage. This broader validation verifies that the mitigation generalizes across the full set of previously permeable conditions and does not introduce unintended artifacts or degradations under other excitation scenarios.

Beyond these two historically dominant cases, none of the remaining excitation conditions retained from the pre-mitigation permeability table produces a residual deviation above 0.05°. In other words, 0.05° constitutes the maximum observed post-mitigation deviation across the entire retained excitation set as shown in Appendix A, Table A.2. This contrasts sharply with the pre-mitigation baseline, where the same multi-GHz excitation regimes yielded deviations in the tens-of-degrees range across multiple bands.

These results indicate that the post-ADC digital filtering strategy effectively suppresses the deviations associated with the identified mechanism, not only under the previously analysed worst-case conditions, but across the full set of excitation cases retained from the black-box regime. The residual angular error remains well below the adopted functional threshold for all cases considered, confirming that the dominant vulnerability has been mitigated within the defined injection–evaluation framework.

Alternative corrective strategies may also be pursued. However, any modification to the system must be evaluated within the same disciplined assessment loop. Interventions at the physical, architectural, or algorithmic level may alter the overall system behavior and potentially introduce secondary effects. For this reason, corrective actions should always be verified under the same excitation conditions and performance criteria used for vulnerability identification, ensuring that mitigation does not inadvertently degrade nominal operation or introduce new sensitivities.

6.6 Summary

This chapter reports the end-to-end application of the operationalized A²EM workflow to the analog sensor front-end under study and demonstrates how an evaluation campaign can progress from functional observability to coupling-oriented attribution and mitigation closure. Across the observability regimes, the black-box analysis phase identified resonance-consistent excitation conditions that are functionally permeable at the output. The gray-box analysis phase allows the interpretation of these outcomes through the architectural frequency constraints of the analog chain. Finally the white-box analysis phase exploites simulation controllability to isolate coupling locations and to attribute the dominant contribution to the sampling-interface interconnect region. Taken together, the three regimes illustrate the intended role separation within A²EM, with screening, interpretation, and attribution executed under a common experimental specification.

The outcome should be interpreted as evidence that A²EM can be operationalized into a disciplined analysis loop on a non-trivial analog front-end, rather than as a global validation of the methodology. The present case study deliberately focuses on the identification of functionally permeable excitation conditions, using non-negligible output deviation as the main criterion for screening and subsequent analysis. This provides a pragmatic and conservative starting point for end-to-end closure, while remaining fully compatible with the broader scope of the methodology. A natural extension is therefore to exercise the same workflow on additional interference mechanisms and effect classes that are plausible and safety-relevant in sensor front-ends, including IMD-like conversion, aliasing and sampling artifacts, estimator sensitivity to structured interference, and other nonlinear or mixed-domain pathways.

The present results are conditioned by the architecture of the specific analog sensor front-end schematic considered (carrier-based transduction, narrowband conditioning, demodulation, baseband filtering, and sampling) and by the entry points exercised in this instantiation, namely PCB interconnecting traces. Demonstrating methodological robustness therefore requires repeating the same end-to-end A²EM instantiation across different sensor types and architectures, and exercising richer entry-point models (e.g., vias, connectors, harness segments, and packaging or enclosure discontinuities), potentially supported by higher-fidelity coupling representations. Such extensions would test whether the observed dominance of the sampling interface is a recurring architectural feature or a consequence of the present sensor and modeling assumptions.

Within these boundaries, the chapter nevertheless delivers a complete closure loop. Once the dominant coupling region was attributed at the sampling interface, mitigation options could be evaluated as trade-offs between physical hardening and low-intrusion, signal-oriented measures. For this proof-of-concept, a post-ADC digital low-pass filter with cutoff frequency f_c of 100 Hz was selected as a software-only countermeasure compatible with the established simulation workflow. Validation under strict one-to-one comparison against the pre-mitigation excitation set shows that all previously retained permeable cases are suppressed to a maximum residual deviation of $0.05°$, i.e., well below the original $1°$ functional threshold. This demonstrates that, under the mechanism localized in this instantiation, mitigation closure can be achieved without redesigning the coupling geometry, while reinforcing the broader point that A²EM functions as an engineering workflow linking vulnerability identification to actionable design decisions under explicit constraints.

Chapter 7

Conclusion

Safety-critical CPS increasingly depend on complex analog sensor front-ends to mediate the interaction between physical phenomena and computational decision-making. Because these front-ends operate under tight SWaP-C constraints, remain closely coupled to the physical environment, and sit at the boundary between sensing and computation, they constitute one of the most exposed and fragile parts of the overall system architecture.

Several adjacent domains had already identified important aspects of this fragility. Functional-safety and analog-validation literature had highlighted the persistent difficulty of achieving meaningful fault coverage in analog and mixed-signal circuitry. Hardware-security research had shown that sensor signal chains could be exploited through physical and architectural weaknesses, particularly in the context of hardware Trojans and physical-to-cyber attacks. Taken together, these strands of work suggested that analog front-ends should not be understood merely as passive interfaces, but as structurally vulnerable regions in which hardware-level disturbances may enter the system, propagate through constrained signal paths, and ultimately produce misleading or unsafe functional effects.

At the same time, the state of the art remained fragmented in the way it accounted for these vulnerabilities. The relevant knowledge was distributed across partially disconnected communities, each emphasizing different threat sources, analytical goals, and terminologies. As a result, no single perspective provided a sufficiently integrated account of how such disturbances couple into analog sensor architectures, how they propagate across front-end stages, and how they translate into system-level deviation. It is precisely within this broader space of hardware-level disturbance and propagation that the electromagnetic dimension becomes especially relevant: electromagnetic interference is not an unrelated problem, but one particularly important and realistic manifestation of the same underlying challenge, since it provides a direct means

by which physical exposure can excite architectural fragilities in analog front-ends and turn them into functional effects.

It is within these challenges that the present thesis is positioned. Rather than treating electromagnetic vulnerability as an isolated EMC issue, this work starts from the premise that the problem must be understood at the intersection of safety, security, and electromagnetic robustness, and that progress requires a more systematic way of linking physical exposure, architectural susceptibility, and functional deviation in analog sensor front-ends.

These challenges give rise to the three central research questions set out in Chapter 1: whether the fragmented insights of functional safety, hardware security, and electromagnetic compatibility can be consolidated into a common understanding of vulnerability in analog sensor front-ends; whether that understanding can be translated into a structured, architecture-aware methodology for EMV identification at circuit and signal-chain level; and whether such a methodology can be operationalized on a representative safety-critical front-end so as to support vulnerability discovery, interpretation, attribution, and mitigation closure.

This thesis addresses those questions through three connected contributions: the consolidation of fragmented cross-domain perspectives on front-end vulnerability, the formulation of A^2EM as a structured architecture-aware methodology for EMV identification, and the demonstration that this methodology can be instantiated into a complete workflow linking physically grounded disturbance conditions, stage-level interpretation, and mitigation-oriented evaluation.

The case-study results show that this approach is workable and informative on a non-trivial safety-critical sensor architecture. Through the coupled-coil position-sensor demonstrator, A^2EM supported an end-to-end progression from the derivation of disturbance conditions and the identification of functionally permeable cases to architectural interpretation, localization of the dominant coupling region, and validation of a mitigation measure within the same evaluation logic. Although this does not constitute a universal validation claim across all classes of analog front-end, it does demonstrate that electromagnetic vulnerability discovery in analog front-ends can be approached as a structured process of architectural reasoning, physically grounded analysis, and mitigation closure.

In this sense, the thesis does more than add another case study of EMI-induced malfunction. It establishes a disciplined way of analysing how physical exposure interacts with architectural fragilities in safety-critical

sensor chains, and shows that this interaction can be investigated systematically from vulnerability discovery to mitigation-oriented assessment.

7.1 Chapter Summary and Future Directions

The contributions of the individual chapters are presented below. In each case, the main work developed within the chapter is briefly revisited, together with the most immediate directions for future work that follow from it.

Chapter 1 – Introduction

The introduction framed the central problem addressed in this thesis, namely the growing reliance of safety-critical CPS on mixed-signal sensing chains exposed to tight operational constraints and unpredictable electromagnetic environments. It positioned the thesis within the broader challenges identified in the state of the art, delimited the scope of the investigation, and set out the research questions that structure the work developed in the following chapters.

Chapter 2 – Background and State of the Art

This chapter provided the conceptual background required to frame the thesis scope, revisiting the evolution from embedded systems to CPS and clarifying their relationship with adjacent paradigms such as the Internet of Things.Within this CPS context, it examined several strands of literature: limitations in analog fault coverage and validation, physical-to-cyber attacks targeting sensor signal chains, and supply-chain threats including third-party involvement and hardware Trojans. Although typically treated in separate research communities, their joint analysis revealed common mechanisms by which disturbances at hardware and physical interfaces can propagate into system-level effects.

This convergence motivated the introduction of Hardware Integrity Threats (HIT) as an aggregating concept that captures heterogeneous hardware-level threats according to their origin and functional impact. Beyond serving as a taxonomy, HIT provide a practical abstraction for cross-domain reasoning: they reduce fragmentation between safety-, security-, and hardware-oriented perspectives, enable threat scenarios to be compared under a common set of dimensions, and support the transfer of analysis and validation

practices across domains that are otherwise treated separately. Part of the work underlying this synthesis culminated in the peer-reviewed publication "*Let's Get Cyber-Physical: Validation of Safety-Critical Cyber-Physical Systems*" [80], which consolidates the HIT perspective in the context of CPS validation and motivates the thesis contributions that follow.

Chapter 3 – A^{2}EM Methodology

Building on the HIT-oriented framing, the methodology chapter introduced and formalized A^{2}EM by transferring and adapting analysis principles commonly used in sensor-security contexts to the electromagnetic robustness problem at CPS front-ends. By treating heterogeneous hardware-level threats under a shared lens and focusing on how physical disturbances translate into functional effects, A^{2}EM was defined as an architecture-aware workflow that links explicit front-end models to systematic probing strategies. In order to exercise A^{2}EM and to anchor the modeling and methodological design choices carried through the subsequent chapters, a coupled-coil position sensor was used throughout this thesis as a case study.

Within A^{2}EM, analysis is structured around architectural representations of the sensor chain and stage-level transfer-function models, so that disturbance behavior can be interpreted in terms of where coupling can occur, how it propagates through constrained analog stages, and where it becomes observable at the functional output. The framework combines controlled probing and profiling strategies and organizes evidence through increasing observability regimes (black-box, gray-box, and white-box), enabling a disciplined progression from functional screening to interpretation and attribution.

In addition to supporting mechanism-level attribution, A^{2}EM is formulated to accommodate multiple classes of EMI-driven behavior in analog front-ends, including saturation-driven deviations, sampling and aliasing artifacts, conversion-like mechanisms (e.g., IMD-style pathways), estimator sensitivity to structured interference, and other mixed-domain nonlinear manifestations.

These conceptual foundations culminated in the publication "*A^{2}EM: From Vulnerability Discovery to Mitigation Assurance – An Architecture-Aware Electromagnetic Methodology for CPS*" [86], while acknowledging that practical instantiation choices, including modeling substrate, coupling representations, and case-study exercise, are addressed in the subsequent chapters.

Chapter 4 – Operationalizing A^{2}EM

Operationalization translated the conceptual framework into an executable modeling workflow applicable to CPS analog front-ends. The chapter defined the modeling substrate, analysis stages, and interfaces required to apply A^{2}EM in practice, grounding the decomposition on architectural and functional evidence (schematic structure, signal-chain boundaries, and stage-level constraints) so that the transfer-function representation remains aligned with the dominant front-end transformations. Within this representation, entry points were operationalized as stage-localized injection locations derived from schematic and layout evidence; for the present instantiation, only PCB interconnecting traces were exercised as explicit entry points, with their electromagnetic coupling modeling deferred to the next chapter on trace-level electromagnetic coupling. The simulator and model set represent one possible operationalization, selected to enable controllable exploration and interpretable stage-level analysis; as modeling assumptions become limiting, fidelity can be increased through richer component and coupling representations and higher-fidelity simulation substrates while preserving the same workflow structure.

Chapter 5 – Trace Modeling

Trace modeling introduced the coupling layer needed to exercise A^{2}EM with disturbance conditions derived from the implemented coupled-coil sensor PCB. Using a thin-wire Method of Moments formulation, the chapter developed a workflow that reconstructs PCB interconnecting traces from layout data, applies electrical boundary conditions from the surrounding circuitry, and computes induced responses under external electromagnetic excitation. Beyond producing trace-level responses, the frequency- and polarization-dependent coupling behavior was analyzed to extract representative worst-case conditions (i.e., dominant coupling peaks) for each interconnect, yielding a reduced excitation set suitable for systematic evaluation without exhaustive sweeps. As a natural improvement path, when the thin-wire MoM abstraction becomes limiting for the trace-level coupling structures considered here, modeling fidelity can be increased by incorporating higher-fidelity electromagnetic solvers and richer coupling representations. The resulting worst-case excitation set provides the physically grounded disturbance conditions propagated and evaluated at system level in the subsequent results chapter.

Chapter 6 – A²EM Results

The results chapter exercised the operationalized A²EM workflow on the coupled-coil position-sensor front-end and demonstrated an end-to-end progression from functional screening to attribution and mitigation closure. Black-box testing identified functionally permeable excitation conditions; gray-box testing interpreted these outcomes through the architectural frequency constraints of the analog chain; and white-box testing used simulation controllability to isolate coupling locations and attribute the dominant contribution to the sampling-interface interconnect region at the ADC boundary. Within the scope of this case study, the analysis deliberately focused on the identification of functionally permeable excitation conditions, using non-negligible output deviation as the main criterion for screening and subsequent interpretation. Following attribution at the sampling interface, a post-ADC digital filtering countermeasure was implemented and validated within the same evaluation logic.

The results should be interpreted as evidence that A²EM can be instantiated into a disciplined analysis loop on a non-trivial analog front-end, rather than as a global validation claim. A direct continuation is therefore to repeat the same workflow across additional sensors and architectures and to exercise it on further interference mechanisms that may yield subtler but safety-relevant deviations, including conversion-like effects, sampling and aliasing artifacts, estimator sensitivity to structured interference, and other mixed-domain nonlinear pathways.

7.2 Contribution to ESRG and Future Directions

This thesis contributes to ESRG's work on automotive, safety-critical CPS by offering a broader perspective and a more structured way of analysing HIT. Beyond the direct contribution of A²EM as a methodology for evaluating electromagnetic vulnerabilities in analog sensor front-ends, the publication *Let's Get Cyber-Physical: Validation of Safety-Critical Cyber-Physical Systems* [80] also opened two complementary directions that extend this broader validation perspective.

One concerns rehosting-based experimentation, particularly through QEMU-based developments with fault-injection support for the analysis of reliable and redundant embedded architectures, which have already been established in previous in-lab work. Within the broader validation perspective of this thesis, this direction complements A²EM by enabling controlled system-level experimentation with injected perturbations and their functional consequences. The other concerns the development of Hardware-in-the-Loop

(HiL)-based capabilities, which remains an ongoing line of work aimed at complementing emulation- and simulation-based approaches through interaction with real sensors, actuators, interfaces, and physical coupling conditions. From the perspective developed here, this is especially relevant for the study of physical-to-cyber attacks, degraded or aged components, hardware alterations, electromagnetic vulnerabilities, and other HIT whose relevant effects depend on physical interaction or may be difficult to capture adequately in purely emulated environments.

Taken together, A^2EM, rehosting-based fault-injection environments, and HiL-oriented developments point toward a more versatile in-house framework for the validation of mixed-signal CPS. By combining cross-domain reasoning with complementary forms of experimentation and analysis, ESRG moves toward a broader HIT-oriented validation basis capable of addressing faults, physical-to-cyber attacks, hardware Trojans, electromagnetic vulnerabilities, and related hardware-level threat scenarios through complementary but connected viewpoints.

References

[1] S. R. Chhetri, A. Barua, S. Faezi, F. Regazzoni, A. Canedo, and M. A. Al Faruque, "Tool of Spies: Leaking your IP by Altering the 3D Printer Compiler," *IEEE Trans. Dependable Secur. Comput.*, vol. 18, no. 2, pp. 667–678, 2021.

[2] "Cyber Physical System Market Size - Forecast To 2031 — businessresearchinsights.com," https://www.businessresearchinsights.com/market-reports/cyber-physical-system-market-100757, [Accessed 20-03-2024].

[3] "X-by-Wire Market — futuremarketinsights.com," https://www.futuremarketinsights.com/reports/x-by-wire-market, [Accessed 20-03-2024].

[4] Y. Shen, G. Heiser, and K. Elphinstone, "Fault Tolerance Through Redundant Execution on COTS Multicores: Exploring Trade-Offs," *Proc. - 49th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Networks, DSN 2019*, pp. 188–200, 2019.

[5] V. Bandur, G. Selim, V. Pantelic, and M. Lawford, "Making the Case for Centralized Automotive E/E Architectures," *IEEE Trans. Veh. Technol.*, vol. 70, no. 2, pp. 1230–1245, 2021.

[6] H. Schirmeier, M. Hoffmann, C. Dietrich, M. Lenz, D. Lohmann, and O. Spinczyk, "FAIL□: An Open and Versatile Fault-Injection Framework for the Assessment of Software-Implemented Hardware Fault Tolerance," *Proc. - 2015 11th Eur. Dependable Comput. Conf. EDCC 2015*, pp. 245–255, 2016.

[7] J. Herkert, J. Borenstein, and K. Miller, "The Boeing 737 MAX: Lessons for Engineering Ethics," *Sci. Eng. Ethics*, vol. 26, no. 6, pp. 2957–2974, 2020. [Online]. Available: https://doi.org/10.1007/s11948-020-00252-y

[8] E. Bachman and I. Capulet, "Tesla Deaths: Digital record of Tesla crashes resulting in death — tesladeaths.com," https://www.tesladeaths.com, 2020, [Accessed 19-03-2024].

[9] K. Jurga and S. Sunter, "Measuring mixed-signal test stimulus quality," *Proc. Eur. Test Work.*, vol. 2018-May, pp. 1–6, 2018.

[10] S. Sunter and P. Sarson, "A/MS benchmark circuits for comparing fault simulation, DFT, and test generation methods," *Proc. - Int. Test Conf.*, vol. 2017-Decem, pp. 1–7, 2017.

[11] The Institution of Engineering and Technology (IET), *Code of Practice on Electromagnetic Resilience*. The Institution of Engineering and Technology, 2017. [Online]. Available: https://shop.theiet.org/code-of-practice-for-electromagnetic-resilience

[12] *IEEE Standard for Techniques and Measures to Manage Functional Safety and Other Risks with Regard to Electromagnetic Disturbances*, Std., 2020. [Online]. Available: https://standards.ieee.org/ieee/1848/7221/

[13] V. Gkatsi, "Risk-based emc engineering for susceptibility exposure: Thinking outside the box," Ph.D. dissertation, University of Twente, Enschede, The Netherlands, Mar. 2025, ph.D. thesis. [Online]. Available: https://doi.org/10.3990/1.9789036565103

[14] L. Devaraj, A. R. Ruddle, and A. P. Duffy, "System Level Risk Analysis for Immunity in Automotive Functional Safety Analyses," *Proc. 2020 Int. Symp. Electromagn. Compat. - EMC Eur. EMC Eur. 2020*, 2020.

[15] K. Armstrong and A. Duffy, "Reducing the Functional Safety Risks (and Other Risks) That Can Be Caused by EMI — New IEEE Standard 1848," *IEEE Lett. Electromagn. Compat. Pract. Appl.*, vol. 2, no. 3, pp. 53–61, sep 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9113675/

[16] V. Gkatsi, B. van den Berg, R. Vogt-Ardatjew, and F. Leferink, "Accurate Terminology and Uniform Methodology for Risk-Based EMC Engineering," *IEEE Lett. Electromagn. Compat. Pract. Appl.*, vol. 7, no. 1, pp. 5–13, 2024.

[17] R. Almeida, L. Novais, N. Naia, R. Faria, and J. Cabral, "Reliable Software Design Aided by QEMU Simulation," *Proc. IEEE Int. Conf. Ind. Technol.*, vol. 2021-March, pp. 797–804, 2021.

[18] W. Wolf, *Computers as Components: Principles of Embedded Computing System Design*, 2nd ed. Waltham, MA, USA: Morgan Kaufmann, 2012.

[19] A. Schnellbach, "Fail-operational automotive systems," Ph.D. dissertation, Graz University of Technology, 2016.

[20] H. Kopetz, *Real-Time Systems: Design Principles for Distributed Embedded Applications*, 3rd ed. Cham: Springer, 2022. [Online]. Available: https://link.springer.com/book/10.1007/978-3-031-11992-7

[21] E. H. Currie, "Introduction to embedded system design," in *Mixed-Signal Embedded Systems Design: A Hands-on Approach*. Springer, 2021, pp. 1–18. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-70312-7_1

[22] R. Zurawski, *Embedded Systems Handbook*. CRC Press, 2018. [Online]. Available: https://www.taylorfrancis.com/books/mono/10.1201/9781315218687

[23] J. Martins, A. Tavares, and M. Solieri, "Bao: A Lightweight Static Partitioning Hypervisor for Modern Multi-Core Embedded Systems," in *Next Generation Real-Time Embedded Systems (NG-RES 2020)*, ser. OASIcs, vol. 77. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, pp. 3:1–3:15. [Online]. Available: https://drops.dagstuhl.de/opus/volltexte/2020/11832/

[24] J. Martins, J. Alves, J. Cabral, A. Tavares, and S. Pinto, "rtzvisor: A secure and safe real-time hypervisor," *Electronics*, vol. 6, no. 4, p. 93, 2017. [Online]. Available: https://www.mdpi.com/2079-9292/6/4/93

[25] CPS Steering Group, "Cyber-physical systems: Executive summary," https://www.nsf.gov/news/special_reports/cyberphysical/CPS-ExecutiveSummary.pdf, 2008, national Science Foundation.

[26] National Science Foundation, "Cyber-physical systems (cps): Program solicitation nsf 08-611," https://www.nsf.gov/pubs/2008/nsf08611/nsf08611.htm, 2008, nSF Program Solicitation.

[27] E. A. Lee, "Cyber-physical systems: Design challenges," in *11th IEEE International Symposium on Object Oriented Real-Time Distributed Computing (ISORC)*. IEEE, 2008, pp. 363–369, downloaded via IEEE Xplore, accessed June 7, 2022. [Online]. Available: https://doi.org/10.1109/ISORC.2008.25

[28] R. Rajkumar, I. Lee, L. Sha, and J. Stankovic, "Cyber-physical systems: The next computing revolution," in *Proceedings of the 47th Design Automation Conference (DAC)*. ACM, 2010, pp. 731–736. [Online]. Available: https://doi.org/10.1145/1837274.1837461

[29] K. Ashton, "That 'internet of things' thing," https://www.rfidjournal.com/that-internet-of-things-thing, 1999, mIT Auto-ID Center (reproduced by RFID Journal).

[30] C. Greer, M. Burns, D. Wollman, and E. Griffor, "Cyber-physical systems and internet of things," National Institute of Standards and Technology (NIST), Tech. Rep. NIST SP 1900-202, 2019, nIST Special Publication 1900-202, March 2019. [Online]. Available: https://doi.org/10.6028/NIST.SP.1900-202

[31] V. Lesch, M. Züfle, A. Bauer, L. Iffländer, C. Krupitzer, and S. Kounev, "A literature review of IoT and CPS—What they are, and what they are not," *J. Syst. Softw.*, vol. 200, p. 111631, jun 2023. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0164121223000262

[32] Y. Ashibani and Q. H. Mahmoud, "Cyber physical systems security: Analysis, challenges and solutions," *Comput. Secur.*, vol. 68, pp. 81–97, 2017. [Online]. Available: http://dx.doi.org/10.1016/j.cose.2017.04.005

[33] D. Serpanos, "There is no safety without security and dependability," *Computer (Long. Beach. Calif).*, vol. 52, no. 6, pp. 78–81, 2019.

[34] C. Yan, H. Shin, C. Bolton, W. Xu, Y. Kim, and K. Fu, "SoK: A minimalist approach to formalizing analog sensor security," *Proc. - IEEE Symp. Secur. Priv.*, vol. 2020-May, pp. 233–248, 2020.

[35] A. Jain, Z. Zhou, and U. Guin, "Survey of recent developments for hardware Trojan detection," *Proc. - IEEE Int. Symp. Circuits Syst.*, vol. 2021-May, 2021.

[36] M. R. Monjur, S. Sunkavilli, and Q. Yu, "ADobf: Obfuscated Detection Method against Analog Trojans on I2C Master-Slave Interface," *Midwest Symp. Circuits Syst.*, vol. 2020-Augus, pp. 1064–1067, 2020.

[37] S. Sunter, "Analog Fault Simulation-a Hot Topic," *Proc. Eur. Test Work.*, vol. 2020-May, 2020.

[38] "IEEE P2427 - Draft Standard for Analog Defect Modeling and Coverage," https://standards.ieee.org/ieee/2427/7233/, 2022.

[39] S. Mukherjee, *Architecture Design for Soft Errors*, 2008.

[40] S. Borkar, N. P. Jouppi, and P. Stenstrom, "Microprocessors in the era of terascale integration," *Proc. -Design, Autom. Test Eur. DATE*, pp. 237–242, 2007.

[41] J. Arlat, Y. Crouzet, J. Karlsson, P. Folkesson, and G. Leber, "Evaluation of the MARS Architecture by means of Three Physical Fault Injection Techniques," no. i, pp. 9–11, 1999.

[42] A. Höller, G. Macher, T. Rauter, J. Iber, and C. Kreiner, "A Virtual Fault Injection Framework for Reliability-Aware Software Development," *Proc. - 2015 45th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Networks Work. DSN-W 2015*, pp. 69–74, 2015.

[43] M. Saikiran, M. Sekyere, M. Ganji, R. Yang, and D. Chen, "Low-cost defect simulation framework for analog and mixed signal (AMS) circuits with enhanced time-efficiency," *Analog Integr. Circuits Signal Process.*, pp. 0–30, 2023.

[44] S. Azam, N. Dall'Ora, E. Fraccaroli, R. Gillon, and F. Fummi, "Analog Defect Injection and Fault Simulation Techniques: A Systematic Literature Review," *IEEE Trans. Comput. Des. Integr. Circuits Syst.*, vol. PP, p. 1, 2023.

[45] R. Bhattacharya, S. H. Ragamai, and S. Kumar, "SFG based fault simulation of linear analog circuits using fault classification and sensitivity analysis," *Commun. Comput. Inf. Sci.*, vol. 711, pp. 179–190, 2017.

[46] S. Sunter, K. Jurga, and A. Laidler, "Using Mixed-Signal Defect Simulation to Close the Loop between Design and Test," *IEEE Trans. Circuits Syst. I Regul. Pap.*, vol. 63, no. 12, pp. 2313–2322, 2016.

[47] S. Sunter, "Efficient analog defect simulation," *Proc. - Int. Test Conf.*, vol. 2019-Novem, pp. 1–10, 2019.

[48] P. Horowitz and W. Hill, *The Art of Electronics: The x Chapters*. Cambridge University Press, 2020. [Online]. Available: https://books.google.pt/books?id=AcVxxQEACAAJ

[49] "Ieee p2427 working group webpage - analogue benchmark circuits," https://sagroups.ieee.org/2427/analogue-benchmark-circuits, [Accessed 25-03-2024].

[50] S. Azam, N. Dall'ora, E. Fraccaroli, A. Alberts, R. Gillon, and F. Fummi, "Investigation on Realistic Stuck-on/off Defects to Complement IEEE P2427 Draft Standard," *Proc. - Int. Symp. Qual. Electron. Des. ISQED*, vol. 2022-April, pp. 51–57, 2022.

[51] K. Yang, M. Hicks, Q. Dong, T. Austin, and D. Sylvester, "A2: Analog Malicious Hardware," *Proc. - 2016 IEEE Symp. Secur. Privacy, SP 2016*, pp. 18–37, 2016.

[52] K. Xiao, D. Forte, Y. Jin, R. Karri, S. Bhunia, and M. Tehranipoor, "Hardware trojans: Lessons learned after one decade of research," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 22, no. 1, 2016.

[53] A. Pavlidis, E. Faehn, M. M. Louerat, and H. G. Stratigopoulos, "Run-Time Hardware Trojan Detection in Analog and Mixed-Signal ICs," *Proc. IEEE VLSI Test Symp.*, vol. 2022-April, 2022.

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

[55] R. S. Chakraborty, S. Narasimhan, and S. Bhunia, "Hardware trojan: Threats and emerging solutions," *Proc. - IEEE Int. High-Level Des. Valid. Test Work. HLDVT*, pp. 166–171, 2009.

[56] Y. Shiyanovskii, F. Wolff, A. Rajendran, C. Papachristou, D. Weyer, and W. Clay, "Process reliability based trojans through NBTI and HCI effects," in *2010 NASA/ESA Conf. Adapt. Hardw. Syst.* IEEE, jun 2010, pp. 215–222. [Online]. Available: http://ieeexplore.ieee.org/document/5546257/

[57] C. Kison, O. M. Awad, M. Fyrbiak, and C. Paar, "Security Implications of Intentional Capacitive Crosstalk," *IEEE Trans. Inf. Forensics Secur.*, vol. 14, no. 12, pp. 3246–3258, 2019.

[58] P. Dash, C. Perkins, and R. M. Gerdes, "Remote activation of hardware trojans via a covert temperature channel," *Lect. Notes Inst. Comput. Sci. Soc. Telecommun. Eng. LNICST*, vol. 164, pp. 294–310, 2015.

[59] K. S. Subramani, N. Helal, A. Antonopoulos, A. Nosratinia, and Y. Makris, "Amplitude-Modulating Analog/RF Hardware Trojans in Wireless Networks: Risks and Remedies," *IEEE Trans. Inf. Forensics Secur.*, vol. 15, pp. 3497–3510, 2020.

[60] D. Agrawal, S. Baktir, D. Karakoyunlu, P. Rohatgi, and B. Sunar, "Trojan detection using IC fingerprinting," *Proc. - IEEE Symp. Secur. Priv.*, pp. 296–310, 2007.

[61] D. Forte, C. Bao, and A. Srivastava, "Temperature tracking: An innovative run-time approach for hardware Trojan detection," *IEEE/ACM Int. Conf. Comput. Des. Dig. Tech. Pap. ICCAD*, pp. 532–539, 2013.

[62] Y. Hou, H. He, K. Shamsi, Y. Jin, D. Wu, and H. Wu, "On-Chip Analog Trojan Detection Framework for Microprocessor Trustworthiness," *IEEE Trans. Comput. Des. Integr. Circuits Syst.*, vol. 38, no. 10, pp. 1820–1830, 2019.

[63] A. Chakraborty, N. G. Jayasankaran, Y. Liu, J. Rajendran, O. Sinanoglu, A. Srivastava, Y. Xie, M. Yasin, and M. Zuzak, "Keynote: A Disquisition on Logic Locking," *IEEE Trans. Comput.*

Des. Integr. Circuits Syst., vol. 39, no. 10, pp. 1952–1972, oct 2020. [Online]. Available: https://ieeexplore.ieee.org/document/8852678/

[64] A. Vijayakumar, V. C. Patil, D. E. Holcomb, C. Paar, and S. Kundu, "Physical Design Obfuscation of Hardware: A Comprehensive Investigation of Device and Logic-Level Techniques," *IEEE Trans. Inf. Forensics Secur.*, vol. 12, no. 1, pp. 64–77, jan 2017. [Online]. Available: http://ieeexplore.ieee.org/document/7546854/

[65] T. D. Perez and S. Pagliarini, "A survey on split manufacturing: Attacks, defenses, and challenges," *IEEE Access*, vol. 8, pp. 184 013–184 035, 2020.

[66] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, "Light commands: Laser-based audio injection attacks on voice-controllable systems," *Proc. 29th USENIX Secur. Symp.*, pp. 2631–2648, 2020.

[67] T. Trippel, O. Weisse, W. Xu, P. Honeyman, and K. Fu, "WALNUT: Waging Doubt on the Integrity of MEMS Accelerometers with Acoustic Injection Attacks," *Proc. - 2nd IEEE Eur. Symp. Secur. Privacy, EuroS P 2017*, pp. 3–18, 2017.

[68] C. Yan, G. Zhang, X. Ji, T. Zhang, T. Zhang, and W. Xu, "The Feasibility of Injecting Inaudible Voice Commands to Voice Assistants," *IEEE Trans. Dependable Secur. Comput.*, vol. 18, no. 3, pp. 1108–1124, 2021.

[69] Y. Tu, A. Rodriguez, S. Rampazzi, K. Fu, B. Hao, and X. Hei, "Trick or heat? Manipulating critical temperature-based control systems using rectification attacks," *Proc. ACM Conf. Comput. Commun. Secur.*, pp. 2301–2315, 2019.

[70] Y. Park, Y. Son, H. Shin, D. Kim, and Y. Kim, "This ain't your dose: Sensor spoofing attack on medical infusion pump," *10th USENIX Work. Offensive Technol. WOOT 2016*, 2016.

[71] D. F. Kune, J. Backes, S. S. Clark, D. Kramer, M. Reynolds, K. Fu, Y. Kim, and W. Xu, "Ghost talk: Mitigating EMI signal injection attacks against analog sensors," *Proc. - IEEE Symp. Secur. Priv.*, pp. 145–159, 2013.

[72] J. Zhang, A. R. Beresford, and I. Sheret, "SensorID: Sensor Calibration Fingerprinting for Smartphones," in *2019 IEEE Symp. Secur. Priv.*, vol. 2019-May. IEEE, may 2019, pp. 638–655. [Online]. Available: https://ieeexplore.ieee.org/document/8835276/

[73] Y. Son, H. Shin, D. Kim, Y. Park, J. Noh, K. Choi, J. Choi, and Y. Kim, "Rocking drones with intentional sound noise on gyroscopic sensors," *Proc. 24th USENIX Secur. Symp.*, pp. 881–896, 2015.

[74] C. Yan, "Can You Trust Autonomous Vehicles : Contactless Attacks against Sensors of Self-driving Vehicle," 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:27264520

[75] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and LiDAR," *Blackhat.com*, pp. 1–13, 2015. [Online]. Available: https://www.blackhat.com/docs/eu-15/materials/eu-15-Petit-Self-Driving-And-Connected-Cars-Fooling-Sensors-And-Tracking-Drivers-wp1.pdf

[76] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 10529 LNCS, pp. 445–467, 2017.

[77] A. Barua and M. A. Al Faruque, "Sensor security: Current progress, research challenges, and future roadmap," *IEEE/ACM Int. Conf. Comput. Des. Dig. Tech. Pap. ICCAD*, 2022.

[78] K. Fu and W. Xu, "Risks of trusting the physics of sensors," *Commun. ACM*, vol. 61, no. 2, pp. 20–23, jan 2018. [Online]. Available: https://dl.acm.org/doi/10.1145/3176402

[79] M. A. Kajol, M. M. R. Monjur, and Q. Yu, "A Circuit-Level Solution for Secure Temperature Sensor," *Sensors*, vol. 23, no. 12, p. 5685, 2023.

[80] L. Novais, N. Naia, J. Azevedo, and J. Cabral, "Let's get cyber-physical: Validation of safety-critical cyber-physical systems," *IEEE Access*, vol. 12, pp. 142 569–142 581, 2024.

[81] European Commission, "Guide for the emc directive (directive 2014/30/eu)," European Commission, Brussels, Belgium, Tech. Rep., Dec. 2018, clarifies that the EMCD is not a safety-related directive; consolidated guidance. [Online]. Available: https://ec.europa.eu/docsroom/documents/33601/attachments/1/translations/en/renditions/native

[82] Planck Collaboration, "Planck 2018 results. vi. cosmological parameters," *Astronomy & Astrophysics*, vol. 641, p. A6, 2020.

[83] D. Pissoort, A. Degraeve, and K. Armstrong, "EMI Risk Management: A necessity for safe and reliable electronic systems!" *5th IEEE Int. Conf. Consum. Electron. - Berlin, ICCE-Berlin 2015*, pp. 208–210, 2016.

[84] F. Sabath, "A systematic approach for electromagnetic interference risk management," *IEEE Electromagn. Compat. Mag.*, vol. 6, no. 4, pp. 99–106, 2017.

[85] R. Morrison, *Fast Circuit Boards: Energy Management*. Wiley, 2018. [Online]. Available: https://www.wiley.com/en-us/fast-circuit-boards%3A-energy-management-p-9781119413905

[86] L. Novais, N. Naia, P. Lages, V. Silva, and J. Cabral, "A2EM: From vulnerability discovery to mitigation assurance—an architecture-aware electromagnetic methodology for cpss," *IEEE Transactions on Electromagnetic Compatibility*, 2026, early access article; final pagination not yet assigned.

[87] M. Mehri, N. Masoumi, and J. Rashed-Mohassel, "Trace Orientation Function for Statistical Prediction of PCB Radiated Susceptibility and Emission," *IEEE Trans. Electromagn. Compat.*, vol. 57, no. 5, pp. 1168–1178, oct 2015. [Online]. Available: http://ieeexplore.ieee.org/document/7078834/

[88] H. C. Hsieh, C. N. Chiu, M. S. Lin, C. H. Wang, and C. H. Chen, "An equation-based hybrid method for predicting radiated susceptibility responses of RF/microwave circuits," *IEEE Trans. Electromagn. Compat.*, vol. 53, no. 2, pp. 339–348, 2011.

[89] Analog Devices, Inc., "Chapter 4: RF/IF Circuits," in *Basic Linear Design*, H. Zumbahlen, Ed. Norwood, MA: Analog Devices, Inc., 2007, pp. 4.1–4.73, accessed: 2025-10-22. [Online]. Available: https://www.analog.com/media/en/training-seminars/design-handbooks/basic-linear-design/chapter4.pdf

[90] M. Leone, "Radiated susceptibility on the printed-circuit-board level: Simulation and measurement," *IEEE Trans. Electromagn. Compat.*, vol. 47, no. 3, pp. 471–478, 2005.

[91] P. J. da Silva Santos, "Wave-based sensor, actuator and optimizer," Ph.D. thesis, Universidade do Minho, Escola de Engenharia, Braga, Portugal, oct 2023.

[92] Paul and Clayton R, *Introduction to Electromagnetic Compatibility {Wiley Series in Microwave and Optical Engineering; 2nd Ed.}*, 2006.

[93] R. F. Harrington, *Field Computation by Moment Methods*. IEEE Press, 1968.

[94] M. Leone and H. L. Singer, "On the coupling of an external electromagnetic field to a printed circuit board trace," *IEEE Trans. Electromagn. Compat.*, vol. 41, no. 4 pt 2, pp. 418–424, 1999.

[95] F. Rachidi, "A review of field-to-transmission line coupling models with special emphasis to lightning-induced voltages on overhead lines," *IEEE Trans. Electromagn. Compat.*, vol. 54, no. 4, pp. 898–911, 2012.

[96] *IEEE Standard for Techniques and Measures to Manage Functional Safety and Other Risks with Regard to Electromagnetic Disturbances*, Institute of Electrical and Electronics Engineers (IEEE) Std. IEEE Std 1848-2020, Jun. 2020, approved June 2020; Published April 2021; PDF ISBN.

Appendix A

Appendix

Table A.1: Polarisation angles associated with frequency bands presenting non-negligible deviation ($|\Delta\theta| \geq 1°$) in the black-box analysis. Each row corresponds to an identified local maximum exceeding the deviation threshold.

Band (GHz)	Trace	f (GHz)	γ (deg)	Deviation (deg)
1.57–1.60 ($\Delta f = 30$ MHz)	Trace 3	1.5899	109.0	6.16
	Trace 5	1.5717	179.5	2.15
	Trace 6	1.5921	0.0	8.46
	Trace 7	1.5911	85.0	11.12
	Trace 8	1.5933	51.5	14.88
	Trace 9	1.5784	115.0	1.83
	Trace 10	1.5933	52.0	14.88
2.24–2.27 ($\Delta f = 30$ MHz)	Trace 5	2.2492	0.0	9.40
	Trace 6	2.2429	0.0	6.28
	Trace 7	2.2597	159.0	21.43
	Trace 8	2.2677	38.0	10.76
	Trace 9	2.2587	114.0	7.94
	Trace 10	2.2537	124.0	7.46
	Trace 11	2.2611	178.0	23.84
3.18–3.20 ($\Delta f = 20$ MHz)	Trace 1	3.1877	125.5	24.43
	Trace 2	3.1956	21.5	30.47
	Trace 3	3.1899	104.5	21.53

Continued on next page

Table A.1 – continued from previous page

Band (GHz)	Trace	f (GHz)	γ (deg)	Deviation (deg)
	Trace 5	3.1902	179.0	35.22
	Trace 6	3.1961	179.5	40.82
	Trace 7	3.1944	150.0	49.60
	Trace 8	3.1939	155.5	50.12
	Trace 9	3.1984	113.5	24.71
	Trace 10	3.1939	155.5	50.12
	Trace 11	3.1974	11.5	30.90
4.46–4.50 ($\Delta f = 40$ MHz)	Trace 6	4.4875	179.0	4.14
	Trace 7	4.4860	101.0	24.66
	Trace 8	4.4970	128.0	11.97
	Trace 9	4.4617	113.5	5.68
	Trace 9	4.4953	113.0	15.15
	Trace 10	4.4878	100.5	24.65
	Trace 11	4.4869	102.0	24.86
4.78–4.79 ($\Delta f = 10$ MHz)	Trace 7	4.7870	82.5	9.77
	Trace 8	4.7878	100.5	8.59
	Trace 9	4.7879	113.0	7.54
	Trace 10	4.7857	79.0	10.06
	Trace 11	4.7896	95.0	7.49
6.33–6.39 ($\Delta f = 60$ MHz)	Trace 2	6.3482	21.0	6.72
	Trace 2	6.3811	20.5	9.54
	Trace 5	6.3338	0.0	3.12
	Trace 5	6.3850	1.0	5.93
	Trace 7	6.3697	52.0	20.75
	Trace 8	6.3686	44.5	21.23
	Trace 9	6.3617	115.0	4.81
	Trace 10	6.3686	44.5	21.23
	Trace 11	6.3626	62.0	15.98

Continued on next page

Table A.1 – continued from previous page

Band (GHz)	**Trace**	f **(GHz)**	γ **(deg)**	**Deviation (deg)**
6.74–6.79 ($\Delta f = 50$ MHz)	Trace 4	6.7765	99.5	19.41
	Trace 6	6.7522	179.5	2.95
	Trace 7	6.7759	61.5	26.65
	Trace 8	6.7810	81.5	21.26
	Trace 9	6.7439	116.0	3.00
	Trace 9	6.7848	112.0	10.35
	Trace 10	6.7799	50.0	23.70
	Trace 11	6.7764	62.0	26.70
7.97–8.00 ($\Delta f = 30$ MHz)	Trace 1	7.9834	14.0	8.82
	Trace 3	7.9849	65.5	9.27
	Trace 5	7.9951	0.5	5.73
	Trace 6	7.9756	0.0	4.85
	Trace 7	7.9851	44.5	10.39
	Trace 8	7.9866	40.0	10.54
	Trace 9	7.9834	115.5	3.37
	Trace 10	7.9864	40.0	10.54
	Trace 11	7.9785	142.5	1.44
	Trace 11	7.9955	112.5	1.93
8.90–8.97 ($\Delta f = 70$ MHz)	Trace 4	8.9602	127.0	7.50
	Trace 6	8.9521	0.0	5.87
	Trace 7	8.9512	133.0	8.18
	Trace 9	8.9266	117.5	4.30
	Trace 10	8.9008	38.5	1.02
	Trace 10	8.9638	166.5	5.89
	Trace 11	8.9531	136.5	8.23
9.12–9.18 ($\Delta f = 60$ MHz)	Trace 7	9.1783	64.0	1.28
	Trace 10	9.1294	21.5	1.03
9.57–9.60 ($\Delta f = 30$ MHz)	Trace 5	9.5913	0.0	18.48

Continued on next page

Table A.1 – continued from previous page

Band (GHz)	Trace	f (GHz)	γ (deg)	Deviation (deg)
	Trace 6	9.5790	0.0	14.83
	Trace 7	9.5841	26.0	16.53
	Trace 8	9.5881	3.5	19.27
	Trace 10	9.5880	3.5	19.26
	Trace 11	9.5821	4.0	16.97

Table A.2: Polarisation angles associated with excitation cases retained from the pre-mitigation black-box permeability table, now reporting the residual angular deviation after implementing the post-ADC digital low-pass filter. Each row corresponds to the same identified local maximum (frequency/polarisation) as in the pre-mitigation table, enabling direct one-to-one comparison, including cases whose post-mitigation deviation falls below the $|\Delta\theta| \geq 1^\circ$ screening threshold.

Band (GHz)	Trace	f (GHz)	γ (deg)	Deviation (deg)
1.57–1.60 ($\Delta f = 30$ MHz)	Trace 3	1.5899	109.0	0.0007011
	Trace 5	1.5717	179.5	0.001208
	Trace 6	1.5921	0.0	0.002557
	Trace 7	1.5911	85.0	0.002224
	Trace 8	1.5933	51.5	0.006457
	Trace 9	1.5784	115.0	0.000481
	Trace 10	1.5933	52.0	0.006456
2.24–2.27 ($\Delta f = 30$ MHz)	Trace 5	2.2492	0.0	0.004866
	Trace 6	2.2429	0.0	0.001419
	Trace 7	2.2597	159.0	0.008637
	Trace 8	2.2677	38.0	0.003055
	Trace 9	2.2587	114.0	0.002681
	Trace 10	2.2537	124.0	0.002376
	Trace 11	2.2611	178.0	0.002002
3.18–3.20 ($\Delta f = 20$ MHz)	Trace 1	3.1877	125.5	0.009823
	Trace 2	3.1956	21.5	0.007723
	Trace 3	3.1899	104.5	0.01733

Continued on next page

Table A.2 – continued from previous page

Band (GHz)	**Trace**	f **(GHz)**	γ **(deg)**	**Deviation (deg)**
	Trace 5	3.1902	179.0	0.03081
	Trace 6	3.1961	179.5	0.008391
	Trace 7	3.1944	150.0	0.01658
	Trace 8	3.1939	155.5	0.02061
	Trace 9	3.1984	113.5	0.002072
	Trace 10	3.1939	155.5	0.02061
	Trace 11	3.1974	11.5	0.004656
4.46–4.50 ($\Delta f = 40$ MHz)	Trace 6	4.4875	179.0	0.001507
	Trace 7	4.4860	101.0	0.01535
	Trace 8	4.4970	128.0	0.006911
	Trace 9	4.4617	113.5	0.003437
	Trace 9	4.4953	113.0	0.02145
	Trace 10	4.4878	100.5	0.006059
	Trace 11	4.4869	102.0	0.009158
4.78–4.79 ($\Delta f = 10$ MHz)	Trace 7	4.7870	82.5	0.000085
	Trace 8	4.7878	100.5	0.02572
	Trace 9	4.7879	113.0	0.02771
	Trace 10	4.7857	79.0	0.007495
	Trace 11	4.7896	95.0	0.005513
6.33–6.39 ($\Delta f = 60$ MHz)	Trace 2	6.3482	21.0	0.003007
	Trace 2	6.3811	20.5	0.000441
	Trace 5	6.3338	0.0	0.0003904
	Trace 5	6.3850	1.0	0.004613
	Trace 7	6.3697	52.0	0.003532
	Trace 8	6.3686	44.5	0.001369
	Trace 9	6.3617	115.0	0.005708
	Trace 10	6.3686	44.5	0.001369
	Trace 11	6.3626	62.0	0.0408

Continued on next page

Table A.2 – continued from previous page

Band (GHz)	Trace	f (GHz)	γ (deg)	Deviation (deg)
6.74–6.79 ($\Delta f = 50$ MHz)	Trace 4	6.7765	99.5	0.04619
	Trace 6	6.7522	179.5	0.001409
	Trace 7	6.7759	61.5	0.02405
	Trace 8	6.7810	81.5	0.006438
	Trace 9	6.7439	116.0	0.001651
	Trace 9	6.7848	112.0	0.0004257
	Trace 10	6.7799	50.0	0.01013
	Trace 11	6.7764	62.0	0.04504
7.97–8.00 ($\Delta f = 30$ MHz)	Trace 1	7.9834	14.0	0.0003303
	Trace 3	7.9849	65.5	0.001733
	Trace 5	7.9951	0.5	0.001301
	Trace 6	7.9756	0.0	0.002652
	Trace 7	7.9851	44.5	0.00216
	Trace 8	7.9866	40.0	0.004648
	Trace 9	7.9834	115.5	0.0002748
	Trace 10	7.9864	40.0	0.004187
	Trace 11	7.9785	142.5	0.005657
	Trace 11	7.9955	112.5	0.0002326
8.90–8.97 ($\Delta f = 70$ MHz)	Trace 4	8.9602	127.0	0.00727
	Trace 6	8.9521	0.0	0.0007665
	Trace 7	8.9512	133.0	0.003236
	Trace 9	8.9266	117.5	0.003816
	Trace 10	8.9008	38.5	0.0004837
	Trace 10	8.9638	166.5	0.0007127
	Trace 11	8.9531	136.5	0.0007063
9.12–9.18 ($\Delta f = 60$ MHz)	Trace 7	9.1783	64.0	0.0005676
	Trace 10	9.1294	21.5	0.001651
9.57–9.60 ($\Delta f = 30$ MHz)	Trace 5	9.5913	0.0	0.0006729

Continued on next page

Table A.2 – continued from previous page

Band (GHz)	**Trace**	f **(GHz)**	γ **(deg)**	**Deviation (deg)**
	Trace 6	9.5790	0.0	0.001771
	Trace 7	9.5841	26.0	0.02955
	Trace 8	9.5881	3.5	0.005867
	Trace 10	9.5880	3.5	0.006165
	Trace 11	9.5821	4.0	0.0001093