Data Exfiltration

Vítor Hugo Silva
Master in Information Security
Department of Computer Science
Faculty of Sciences of the University of Porto
2026

Data Exfiltration

Vítor Hugo Silva

Dissertation carried out as part of the Master in Information Security

Department of Computer Science

2026

Supervisor
Rolando Martins, Full Professor, Faculty of Sciences of the University of Porto

External Host Supervisor
João Soares, Full Professor, Faculty of Sciences of the University of Porto

“ Giving power to prejudiced persons is giving a loaded gun to a baby ”

unknown

Acknowledgements

I would like to acknowledge my professor, Rolando Martins, for allowing me to work on a subject that drove me, and Professor João Soares for the time spent helping me accomplish this great work. Without them, this would not have been possible. I would also like to thank my family for their support, along with my friends and work colleagues. Finally, I would like to thank Koerber Porto for allowing me to fulfill my dream of completing my master's thesis with great patience and flexibility.

Resumo

Configurações air-gapped e outros sistemas de alta segurança, especialmente em contextos militares e de inteligência, são tipicamente considerados resistentes ao roubo remoto de dados devido à separação rigorosa das redes e ao controlo das interfaces físicas. Normas como TEMPEST e a disciplina mais ampla de Segurança de Emanações procuram reduzir sinais físicos eletromagnéticos, acústicos ou de outra natureza que possam expor dados sensíveis. Ainda assim, trabalhos anteriores demonstraram que canais encobertos podem explorar efeitos físicos subtis para transferir dados através destas barreiras de isolamento.

A maior parte da investigação anterior examinou canais laterais eletromagnéticos, acústicos, magnéticos, óticos ou térmicos, enquanto as características da qualidade da energia elétrica receberam muito menos atenção sistemática como possíveis portadoras de informação. Em particular, cargas eletrónicas de potência não lineares que interagem com a rede podem induzir alterações detetáveis na distorção harmónica e no fator de potência, métricas que raramente são consideradas como canais de sinalização.

Esta tese investiga se a variação intencional da carga computacional pode produzir alterações mensuráveis nas amplitudes harmónicas e no fator de potência no ponto de ligação à rede de sistemas que utilizam fontes de alimentação comutadas (SMPS). Primeiro, desenvolve-se o enquadramento teórico, cobrindo os componentes do fator de potência sob cargas não lineares e os mecanismos de geração de harmónicos em SMPS com e sem correção ativa do fator de potência. A dependência dos harmónicos de baixa ordem em relação à carga é identificada como o principal mecanismo físico de sinalização.

É utilizada uma montagem experimental para avaliar se os dados podem ser codificados através de alterações controladas da carga computacional. A avaliação apresenta os resultados de descodificação observados, a taxa bruta, o goodput calculado e o efeito da topologia da fonte de alimentação e da posição do recetor nas condições testadas. Cada configuração foi adquirida uma única vez ($N = 1$); consequentemente, o estudo constitui uma prova de conceito exploratória e não uma caracterização estatística da fiabilidade do canal.

Os resultados fornecem evidência de prova de conceito de que a variação da carga computacional pode causar alterações detetáveis nas métricas de qualidade de energia no ponto de ligação à rede. Em configurações selecionadas, estas alterações permitiram recuperar mensagens e material criptográfico através da infraestrutura elétrica. Apesar das baixas taxas de dados e da dependência da topologia da fonte e do ponto de medição, o trabalho acrescenta

um canal encoberto ao nível da infraestrutura e contribui para a modelação de ataques e para a monitorização defensiva em ambientes de alta segurança.

Palavras-chave: Ciência de Computadores, Segurança, Segurança Informática, Cibersegurança, Exfiltração de Dados

Abstract

Air-gapped configurations are commonly adopted in high-security systems, particularly in military and intelligence contexts, and are often perceived as resistant to remote data theft because of their strict network separation and tightly controlled physical interfaces. Standards such as TEMPEST, together with the broader discipline of Emanations Security, seek to reduce electromagnetic, acoustic, and other physical signals that can expose sensitive data. Nonetheless, earlier studies have demonstrated that covert channels can exploit subtle physical effects to transfer data across barriers to isolation.

Previous research has extensively examined electromagnetic, acoustic, magnetic, optical, and thermal side channels, whereas electrical power quality characteristics have received far less systematic attention as possible information carriers. In particular, nonlinear power-electronic loads interacting with electrical mains can induce detectable changes in harmonic distortion and power factor, quantities that are rarely considered as signaling channels.

This dissertation examines whether intentional variations in computational workload can produce measurable changes in harmonic amplitudes and power factor at the grid connection of systems using switched-mode power supplies. It also examines how power-factor-correction designs affect the observability of these features across electrical infrastructure. The theoretical framework covers the power-factor components under nonlinear loads and harmonic-generation mechanisms in switched-mode power supplies with and without active power-factor correction. The load dependence of low-order harmonics is identified as the principal physical signaling mechanism.

An experimental setup was used to evaluate whether the data could be deliberately modulated through controlled workload changes. The evaluation reports the observed decoding outcomes, raw bit rate, calculated payload goodput, and effects of the PSU topology and receiver placement under the tested conditions. Each configuration was acquired once ($N = 1$); therefore, this study is an exploratory proof of concept rather than a statistical characterization of channel reliability.

The findings provide proof-of-concept evidence that computational load variation can cause detectable changes in power quality metrics at the grid connection. In the selected configurations, these changes supported the recovery of messages and cryptographic material through the electrical infrastructure. Despite the low data rates and dependence on PSU topology, the selected observable, receiver placement, and electrical environment, this study adds an

infrastructure-level covert channel and informs both attack modeling and defensive monitoring in high-security settings.

Keywords: Computer Science, Security, Information Security, CyberSecurity, Data Exfiltration

Table of Contents

List of Figures ix
Glossary x

1. Introduction 1
   1.1. Motivation 1
   1.2. Research Objectives 2
   1.3. Research Questions 3
   1.4. Threat Model and Assumptions 3
   1.5. Contributions 4
   1.6. Thesis Outline 5
2. Background 7
   2.1. Understanding the Attack Vectors 8
   2.2. Power Grid 10
      2.2.1. Power Factor and Reactive Power in Linear and Nonlinear Loads 11
      2.2.2. Power Supply Units and Harmonic Generation in Switched-Mode Power Supplies 14
3. Related Work 19
   3.1. Air-gap Security 19
      3.1.1. Audio Waves 21
      3.1.2. Magnetism 24
      3.1.3. Electromagnetic Waves 25
4. Power Factor and Harmonics As Covert Channel 33
   4.1. Experimental Setup and Instrumentation 33
      4.1.1. Testbed Overview 34
      4.1.2. Instrumentation and Acquisition Parameters 35
      4.1.3. Devices Under Test 36
      4.1.4. Workload Modulation Methods 37
      4.1.5. Experimental Design and Scope 38
   4.2. Measurement Scenarios and Receiver Locations 38
   4.3. Symbol Encoding Model 40
   4.4. Measured Quantities and Feature Extraction 42
      4.4.1. Power and Power Factor Metrics 43
      4.4.2. Relative and Absolute Harmonic Metrics 43
      4.4.3. Windowing and Time Resolution 44
      4.4.4. Measurement Window and Symbol Duration 45

4.4.5. Measurement and Estimation Uncertainty ... 45
4.5. Baseline Characterization ... 46
4.6. Load Sweep and Scaling Behaviour ... 48
4.7. Transient Load Modulation (Pulse Ramp) ... 51
4.8. On-Off Keying ... 53
4.8.1. Symbol Detection Model ... 54
4.9. Receiver Placement and Electrical-Path Effects ... 56
4.9.1. Meaning of Distance in a Conducted Power-Line Channel ... 56
4.9.2. No PFC Results ... 57
4.9.3. Passive PFC Results ... 59
4.9.4. Active PFC Results ... 61
4.10. Summary of Receiver Placement and PSU Effects ... 64
5. Discussion ... 66
5.1. Discussion with Respect to the Research Questions ... 66
5.1.1. RQ1: Can workload modulation produce measurable and temporally structured changes in power factor and harmonic distortion? ... 66
5.1.2. RQ2: Can these measurable patterns support end-to-end data recovery, and what decoding performance is observed under the evaluated conditions? ... 67
5.1.3. RQ3: How can the power supply topology and receiver placement affect the observability and practical viability of the proposed covert channel? ... 68
5.2. Comparison with Existing Power-Line Covert Channels ... 69
5.3. Defensive Implications ... 70
5.4. Limitations ... 70
6. Conclusion ... 72

Bibliography ... 74
A. Decode results ... 79
A.1. No Power Factor Correction Power Supply ... 79
A.2. Passive Power Factor Correction Power Supply ... 82
A.3. Active Power Factor Correction Power Supply ... 85

List of Figures

2.1. Comparison between PSU without Active PFC (left) and with Active PFC (right). Without the PFC, the current is drawn in pulses near the voltage peaks, producing a high harmonic distortion. Active PFC shapes the input current to follow the mains voltage waveform, reducing THD and improving power factor. 17
2.2. Input current waveforms for SMPS without PFC, with passive PFC and with active PFC, evaluated under both no-load (idle) and loaded operating conditions. 18

4.1. General Layout 34
4.2. Office measurement location and loads 35
4.3. Representative office floor plan and electrical path for the office-level measurement scenario. Distances are approximate and the drawing is not to scale. 39
4.4. Harmonic heatmap of the input current during idle operation using an SMPS without power factor correction. 47
4.5. Harmonic heatmap during idle operation using a passive PFC PSU. 47
4.6. Harmonic heatmap during idle operation using an active PFC PSU. 48
4.7. Power factor measured during the CPU load sweep experiment (Active PFC PSU). 49
4.8. Absolute and relative harmonic content heatmap of the input current during the CPU load sweep experiment (Active PFC PSU). 50
4.9. Absolute and relative harmonic content heatmap of the input current during the CPU load sweep experiment (Passive PFC PSU). 50
4.10 Absolute and relative harmonic content heatmap of the input current during the CPU load sweep experiment (Non-PFC PSU). 51
4.11 Power factor evolution during transient CPU load modulation. 52
4.12 Harmonic response in absolute and relative. 53
4.13 Heatmap and timeseries for each PSU design measured directly to the PSU. 54
4.14 Comparison of measurement at PSU vs office measurements for the no-PFC PSU - ASCII 58
4.15 Comparison of measurements at the PSU and office locations for the no-PFC PSU - AES key 58
4.16 Comparison of measurements at the PSU and office locations for the no-PFC PSU - ECDSA 59
4.17 Comparison of Passive PFC PSU measurements at PSU and office locations - ASCII 60
4.18 Comparison of Passive PFC PSU measurements at PSU and office locations - AES key . 60
4.19 Comparison of Passive PFC PSU measurements at PSU and office locations - ECDSA ... 61
4.20 Comparison of Active PFC PSU measurements at PSU and office locations - ASCII 62
4.21 Comparison of Active PFC PSU measurements at PSU and office locations - AES key 62
4.22 Comparison of Active PFC PSU measurements at PSU and office locations - ECDSA 63

Glossary

AES	Advanced Encryption Standard
ASCII	American Standard Code for Information Interchange
CIA	Confidentiality, Integrity, and Availability
CD	Compact Disc
CPU	Central Processing Unit
DDR	Double Data Rate
DIMM	Dual In-line Memory Module
DNS	Domain Name System
DoH	DNS over HTTPS
DPI	Deep Packet Inspection
DRAM	Dynamic Random Access Memory
DVD	Digital Versatile Disc
ECDSA	Elliptic Curve Digital Signature Algorithm
EM	Electromagnetic
EMSec	Emanations Security
FM	Frequency Modulation
GB	Gigabyte
GSM	Global System for Mobile Communications
HDD	Hard Disk Drive
HDMI	High-Definition Multimedia Interface
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
ICMP	Internet Control Message Protocol
IoT	Internet of Things
IPS	Intrusion Prevention System
kbps	kilobits per second

LOS	Line of Sight
MB	Megabyte
Mbps	megabits per second
NATO	North Atlantic Treaty Organization
NS	Name Server
OOK	On-Off Keying
PFC	Power Factor Correction
PF	Power Factor
PSU	Power Supply Unit
QAM	Quadrature Amplitude Modulation
RAM	Random Access Memory
RF	Radio Frequency
RSA	Rivest–Shamir–Adleman
SIM	Subscriber Identity Module
SIPRNet	Secret Internet Protocol Router Network
SMPS	Switched-Mode Power Supply
SMS	Short Message Service
SSH	Secure Shell
SSD	Solid-State Drive
TEMPEST	TEMPEST
TLS	Transport Layer Security
TPM	Trusted Platform Module
TXT	Text
USB	Universal Serial Bus
VCD	Video CD
VGA	Video Graphics Array
Wi-Fi	Wireless Fidelity
ZTA	Zero Trust Architecture

1. Introduction

Sensitive data are a crucial aspect of information security. The protection of sensitive data remains a persistent challenge in information security. Organizations that handle such data, whether commercially, governmental, or critical sovereign information, must guarantee compliance and risk at an acceptable level. The risk related to information leaks or exfiltration is growing, and critical infrastructure is particularly targeted during nationwide conflicts.

These risks are practically difficult to address, especially in an environment where digital footprints increase, integrations grow, and society, whether civil, military, or governmental, expands into data-driven organizations. This increases the landscape for attacks to occur, increases the complexity of monitoring, and prevents exfiltration attempts. Even with robust components such as endpoint protection, network detection and response solutions, and strict access controls, research continues to reveal that unconventional channels can bypass traditional safeguards.

High security profiles and organizations implement architectures with zero trust. The level of paranoia is so great that some architectures are totally air-gapped from any other system because the exfiltration risk is very high owing to the very high impact, even for low likelihood. In these environments, it is normal to follow the standards for EMSec and TEMPEST to suppress unintentional channels, such as acoustic emissions and electromagnetic radiation, and other covert channels. However, prior work has demonstrated that physical aspects can create covert channels that are difficult to address, directly or indirectly, unless a very tight monitoring capability and a high awareness of these attacks are in place.

This dissertation examines an emerging problem in emanation security and TEMPEST, in particular with electrical observables, harmonic distortion, and power factor deviation produced by nonlinear power supplies, which can create covert channels for exfiltrating data from an air-gapped system.

1.1. Motivation

Systems with air-gap configurations are normally treated as fortresses against data exfiltration or other attacks because they lack network connections and have very strict physical controls. However, history has proven that this is not always the case. Physical isolation does not eliminate all communication channels. Research has shown that covert channels can still exist

through physical byproducts, either by variations in computational workload, physical components with audio-induced aspects, memory accesses and writes that induce electromagnetic waves that radiate off isolation, and many others.

One basic observation motivates this study is that all digital systems require electrical power for operation. Whether air is gapped or not, this requires continuous power delivery to keep the system running. Because power typically comes from the utility grid, there is a landscape of potential unexplored scenarios that can act as covert channels. The utility grid is not designed with secrecy in mind, nor are the electrical building wiring designs.

Many modern computers rely on switch-mode power supplies. These switched-mode power supplies draw current in sharp nonlinear bursts, which inject harmonic noise into the power grid. The variation in the computational load, either increasing or decreasing, influences the intensity and character of this harmonical fingerprint and power factor aspects. By deliberately changing the workload of these power supplies, an attack can draft clear and decodable patterns into the electrical grid.

This represents not only an offensive opportunity but also a defensive and preventive one. If a distinct electrical signature is trackable by the system's operation, power quality monitoring can become a powerful new sensor to detect such anomalies and strengthen defenses in air-gapped environments.

1.2. Research Objectives

This dissertation focuses on the central question: Can a computer workload be intentionally modulated to produce measurable changes in harmonic distortion and power factor at its grid connection, and can those changes be used as a covert channel?

To answer this central question, we developed an electrical model for nonlinear loads on a grid, examined how the power factor decomposes into separate components, and investigated the generation of harmonics in switched-mode power supplies. We also considered circuit designs that attempt to decrease harmonic distortion and power factor deviation, known as power factor correction circuits. We analyzed power supplies with active, passive, and no power-factor correction circuits.

To this end, this study proposes an experimental framework for measuring and analyzing electrical observables at power supply connections to the grid and assesses whether workload-induced variations produce measurable and temporally structured behavior that supports payload recovery in recorded transmissions.

Thereafter, we examine the observed decoding performance, nominal raw rate, calculated payload goodput, and influence of the PSU topology and receiver placement. These results were compared with those of existing air-gap exfiltration methods, placing this work within the broader landscape of covert channels and outlining strategies for detection, monitoring, and mitigation. Because each experimental configuration was acquired once, this objective was exploratory and did not include a statistical characterization of reliability, noise tolerance, or packet success probability.

1.3. Research Questions

This dissertation aims to address the following three questions:

RQ1. Can workload modulation produce measurable and temporally structured changes in the power factor and harmonic distortion?

RQ2. Can these measurable patterns support end-to-end data recovery, and what decoding performance is observed under the evaluated conditions?

RQ3. How can the power supply topology and receiver placement affect the observability and practical viability of the proposed covert channel?

1.4. Threat Model and Assumptions

The attacker has already compromised an air-gapped or otherwise isolated system and can execute code that intentionally increases or reduces CPU activity on demand. No conventional network path is available, and the threat model excludes the physical removal of data through removable media. Instead, malware encodes data by varying the CPU workload, thereby changing the system's power demand and the electrical quantities observable at the grid interface.

The receiver is connected to the same electrical infrastructure, either directly next to the PSU or upstream of the building wiring. It samples the voltage and current and derives features such as the active and apparent power, power factor, and low-order harmonic amplitudes. These two vantage points correspond to the experimental setups: a near-PSU probe and an office-level receiver on the building's distribution network.

The intended use case is the low-rate exfiltration of compact secrets rather than bulk transfer. The attacker aims to transmit small but critical values, such as short messages, cryptographic keys, or signatures, over a low-bandwidth channel. This aligns with the payloads evaluated in this thesis: an ASCII string, an AES-128 example key, and an ECDSA-related key material.

This study evaluates recoverability and rate; however, it does not experimentally quantify stealth or detectability.

The channel depends on the PSU design, receiver placement, background loads, and measurement noise. In the recorded captures, the active PFC reduced the low-order harmonic observability used by the proposed receiver. Upstream receiver placement may also reduce observability by mixing the target contribution with other loads and electrical-path effects, although the observed BER was feature- and environment-dependent and was not always worse at the office level. Therefore, the threat model assumes a skilled but constrained adversary operating under favorable conditions, not a universally reliable exfiltration technique.

A more capable adversary could attempt to increase channel observability through hardware or supply chain tampering, for example, by replacing, degrading, or bypassing the PFC stage. Such manipulation could increase harmonic distortion and make workload-induced changes easier to observe. These actions were not evaluated experimentally and were included only as threat model extensions.

1.5. Contributions

This study extends the frontier of covert channels in air-gapped systems by using changes in electrical observables as the basis of modulation.

We begin with a theoretical introduction to the power factor, its behavior under nonlinear AC loads, the production of harmonic distortion, and the role of switched-mode power supplies. We analyze how changes in computational load affect the power factor and introduce measurable harmonic content at the grid interface. We then evaluated whether these changes could be organized into temporal patterns and decoded by a receiver.

Accordingly, we established an experimental framework that treats workload-induced electrical variation as a signal and reports the observed power factor and harmonic measurements, SNR, BER, valid Reed–Solomon outcomes, raw bit rate, and calculated payload goodput.

The mechanism was evaluated at two receiver placements, including an office-level point containing contributions from other nonlinear loads. Because each configuration was acquired once, this comparison demonstrates exploratory feasibility under the tested conditions rather than signal robustness, repeatability, or expected success probability.

Finally, this study places the method within a broader set of air-gapped exfiltration methods. We compared the proposed method with existing covert channels and refined the taxonomy of

electrical-level covert channels to examine how this class of threats can inform detection and mitigation strategies in air-gapped systems.

This dissertation investigates a distinct set of electrical observables relative to the closest power-line covert channels. PowerHammer[1] demonstrated that CPU load modulation can produce decodable variations in line current, with emphasis on the achievable transmission rate at line-level and phase-level measurement points. NoDE[2] instead exploits the high-frequency voltage ripple associated with the switching operation of active power factor-correction circuits and receives this signal from another outlet on the same electrical infrastructure.

In contrast, this study investigates whether the derived low-frequency power-quality quantities, particularly the true power factor and individual current-harmonic RMS amplitudes, can serve as information-bearing features. It evaluates this mechanism across power supplies without PFC, with passive PFC, and with active PFC, and at both near-PSU and office-level receiver locations. Therefore, the contribution is not merely the observation that computational workload affects mains current, but also proof-of-concept evidence that routinely measurable power quality metrics can support end-to-end payload recovery under selected conditions and that PSU topology is a central determinant of channel observability.

1.6. Thesis Outline

This dissertation is organized into six primary chapters.

Chapter 1 outlines the research problem, motivates the exploration of power quality parameters as potential covert channels, and specifies the research goals and contributions of this study.

Chapter 2 provides a theoretical foundation. It reviews the principles of data exfiltration and attack vectors. The chapter then develops the electrical engineering background for this study, covering power-factor decomposition, harmonic distortion in nonlinear loads, and switched-mode power supplies with and without active power-factor correction.

Chapter 3 reviews related work on air-gap covert channels across the acoustic, vibrational, magnetic, electromagnetic, optical, thermal, and power-line domains. It compares the main trade-offs in detectability, bandwidth, receiver requirements, and practical deployment, thereby situating this dissertation within the existing literature.

Chapter 4 details the proposed methodology and experimental study. The system configuration, measurement framework, workload modulation scheme, and harmonic analysis techniques are described. This chapter assesses the viability of encoding data via variations in the power factor

and harmonic distortion, quantifies the achievable performance, and considers the detection challenges and real-world constraints.

Chapter 5 discusses the experimental findings with respect to the research questions, compares the proposed method with existing power-line covert channels, and considers the defensive implications and limitations of this study.

Chapter 6 concludes by summarizing the main findings, highlighting the principal contributions and limitations of the study, and outlining possible directions for future research.

2. Background

Data exfiltration is a security weakness that allows an attacker to move information from a targeted system, over which they have limited influence, to another system under the attacker's control. Copying data into a removable drive, uploading it to cloud storage, burning a DVD, and taking copies to a printer are basic and trivial examples of data leakage. Data exfiltration typically implies a more sophisticated method for leaking data. Methods that require network connectivity, such as DNS tunneling, HTTPS tunneling, and other protocol abuses, are more sophisticated. Despite being well-known methods, these are addressed by the cybersecurity industry using various security controls. Such controls include blocking or restricting shared media, monitoring file shares, limiting URLs based on site category to prevent the use of cloud services, controlling or disabling printers, blocking DVD drives and USB storage, and implementing TLS inspection or blocking unwanted protocols that can be abused.

As attack techniques progress, detection and prevention capabilities have improved. Next-generation firewalls are capable of deep traffic inspection, TLS inspection, advanced behavioral analytics in endpoint detection, and response agents that can identify, alert, and stop exfiltration attempts by monitoring sensitive information regarding endpoints and behavior patterns. Although there are more mature and complex defenses, certain exfiltration methods are so advanced that they are difficult to detect or prevent; thus, they must be explicitly and rigorously managed in environments with very high security requirements.

A high-security environment is normally associated with military systems that handle secrets, intelligence agencies that collect valuable and sensitive information, and cutting-edge research facilities that use novel processes to innovate, making it highly valuable. Information in these contexts is sensitive, and the consequences of a leak can be severe, resulting in a low likelihood of high risk due to its very high impact if the risk materializes. Therefore, all potential scenarios must be addressed by physical controls, network-based technologies, or even air-gapped architectures.

Physical techniques involve attackers in close proximity to a target. The exfiltration of information by this means can either be by emanations or component theft. For example, the extraction of hard drives or storage components, cold boot attacks that explore in-memory plaintext cryptographic keys to decrypt hard drives, and TPM sniffing to extract boot keys are discussed in Chapter 3.

Networked exfiltration is linked to attackers who have direct remote control or access to the device and use a network with a channel of exfiltration. These types of methods rely on an unsupervised or poorly monitored network architecture to be unable to detect and prevent network exfiltration techniques. Such methods can rely on protocol abuse, such as DNS tunneling, DNS over HTTPS, SSH, and practically any other encrypted channel, which makes exfiltration difficult to detect. Network administrators often rely on network detection methods to detect anomalies, with or without AI-driven classification of behaviors, such as high-risk domain connections, peak volume of data transfer to external destinations, and others. These controls are normally costly but are available with a comfortable budget for companies and other organizations that want to reduce this risk. TLS inspection to monitor data on encrypted channels introduces legal challenges due to privacy concerns.

Air-gapped exfiltration techniques are applied to isolated systems in which the network is unavailable or physical access is restricted or limited. Attackers often rely on unconventional covert channels that lie under the monitoring radar. Conventional channels have been extensively studied and documented, whereas unconventional covert channels are often unforeseen owing to their high complexity and low execution probability in real environments.

Data exfiltration is not a new concept. It has existed for thousands of years. Currently, it is a key subject in industrial espionage, military and strategic intelligence, workplace threat intelligence, and information services. In this thesis, we present an overview of data exfiltration techniques, introduce the concept of covert channels, and discuss side channels across several domains including physical, network-based, and air-gapped exfiltration methods. Our proposed technique is novel, has not yet been published, and operates specifically on air-gap systems.

2.1. Understanding the Attack Vectors

The feasibility of a data exfiltration technique depends on the attacker's position, the interfaces available on the target system, and the level of isolation imposed by the surrounding environment. An adversary with direct physical access faces different constraints from those operating remotely through a network. Likewise, attacks against isolated systems require communication paths that do not depend on conventional network connectivity. Therefore, it is necessary to clearly define these conditions before evaluating a particular exfiltration method.

Physical exfiltration requires the adversary or a cooperating insider to interact directly with the target system or its surrounding infrastructure. Possible techniques include removing storage devices, copying information to removable media, printing sensitive documents, introducing

malicious hardware, and accessing exposed system interfaces. Although these methods may provide comparatively high transfer rates, they generally require close proximity and may produce observable actions, such as peripheral usage, hardware modifications, system restarts, or unauthorized access to restricted areas. Therefore, their feasibility depends on the physical controls, monitoring procedures, and peripheral restrictions deployed around the target.

Network-based exfiltration assumes that a compromised system has access to a communication network. Data may be transferred directly to an external destination or concealed within legitimate-looking traffic by abusing protocols such as DNS, HTTP, HTTPS, SSH, or ICMP. Attackers may also use proxy and domain-fronting techniques to direct malicious communication towards an authorized service or domain [3]. In mature environments, these methods can be mitigated through network segmentation, protocol filtering, traffic inspection, destination restrictions, and behavioral monitoring. Consequently, network-based exfiltration can provide substantially greater bandwidth than physical covert channels; however, it is also exposed to well-established preventive and detective controls.

Air-gapped systems remove conventional network paths and usually impose stricter physical controls. Examples frequently associated with isolated or highly restricted classified infrastructure include SIPRNet [4], the Joint Worldwide Intelligence Communications System [5], and NATO CRONOS [6]. However, the absence of ordinary network connectivity does not guarantee that a system cannot be compromised or that information cannot leave it. Malware may still be introduced through removable media, maintenance procedures, malicious insiders, hardware tampering, or supply chain compromises.

Supply chain attacks are particularly relevant, because a system may be compromised before it reaches its final operational environment. An adversary may modify software, firmware, or hardware during manufacturing, integration, or distribution. Reported examples include the compromise of SIM-card-related infrastructure [7], unauthorized code discovered in Juniper networking equipment [8], and firmware implants targeting hard drives [9]. Once malicious code is executed on an isolated system, the absence of a network connection prevents conventional exfiltration but does not eliminate alternative physical communication paths.

A compromised air-gapped system may encode information regarding the physical effects generated during normal operation, including acoustic, optical, thermal, magnetic, electromagnetic, vibrational, and electrical emissions. These effects may be intentional when malware actively

controls a hardware component or unintentional when sensitive computation produces observable physical leakage. For example, a previous work demonstrated cryptographic key extraction using electromagnetic and power-related measurements taken near computing equipment [10]. Likewise, acoustic emissions generated during cryptographic operations have been shown to reveal sensitive information [11].

These unconventional paths are commonly treated as covert or side channels because they communicate through physical quantities rather than through an explicitly authorized interface. Their practical effectiveness depends on the emitting hardware, environmental noise, receiver placement, modulation method, and the attacker's ability to control or observe the relevant physical phenomenon. Although these channels generally offer less bandwidth than network-based methods, they may still be sufficient to exfiltrate compact and valuable information, such as credentials, cryptographic keys, short messages, or configuration data.

This dissertation considers a compromised air-gapped computer on which malware can deliberately modify the computational workload. Although the target has no usable network connection, it remains connected to the electrical supply. Variations in computational activity alter the power demanded from the power supply unit and may consequently affect the electrical quantities observable at the grid interface. This motivates the investigation of electrical infrastructure as a potential covert communication path. The following section introduces the power grid and power quality concepts required to understand this mechanism.

2.2. Power Grid

A power grid is an essential component of any digital infrastructure. It is safe to assume that the vast majority of systems require electrical power. Data centers, factories, and offices depend on electricity, whether it comes from the public grid or is self-generated by the facility. Some industrial sites, such as oil refineries, may use a hybrid setup (on-site generation) to optimize refining margins by comparing the co-generated energy from natural gas pipelines with power from the public grid. The same principle applies to air-gap systems. Although these systems are isolated with no network connections or external peripherals, they still require power to operate. The electrical supply enters the system from the outside. Therefore, it is reasonable to consider the possibility of using power lines as communication channels because there is actually wiring from the external environment and the isolated system. Certain techniques for doing so are straightforward to implement but are also easy to detect, whereas others are more complex to perform but far more difficult to identify.

2.2.1 Power Factor and Reactive Power in Linear and Nonlinear Loads

The power factor is fundamental for describing the behavior of electrical loads with respect to the power grid. This is the key to understanding the covert channel proposed in this work. In sinusoidal steady-state operation, voltage and current can be expressed as

$$v(t) = V_m \cos(\omega t + \theta_v), \quad i(t) = I_m \cos(\omega t + \theta_i),$$

with the associated phasor representations $V\angle\theta_v$ and $I\angle\theta_i$. The complex power S is given by

$$S = P + jQ = VI^*,$$

where:

- P expresses the active (real) power in watts (W),
- Q expresses the reactive power in volt-ampere reactive (VAr),
- $|S| = \sqrt{P^2 + Q^2}$ is the apparent power.

The power factor is then defined as

$$PF = \frac{P}{|S|}.$$

In addition, this power factor does not consider harmonic distortions. The power consumption for a given load is not purely resistive. For capacitive and inductive loads, such as those that require inrush currents to charge capacitors or inductive elements, this power is consumed and subsequently returned to the grid. This affects the phase angle between the voltage and current waveforms.

Power Factor in the Sinusoidal Case

When both voltage and current are ideal sinusoids, the power factor reduces to

$$PF = \cos\varphi,$$

where $\varphi = \theta_v - \theta_i$ denotes the phase-angle difference between the voltage and current. In this case, any departure from the unity power factor is caused exclusively by this phase shift, which is typically introduced by inductive or capacitive elements. Inductive loads cause the current to lag behind the voltage ($Q > 0$). For capacitive loads, causes the current to lead to voltage ($Q < 0$).

Here, the increase in apparent power results from the periodic storage and release of energy by the reactive components, capacitors for capacitance reactive power, and inductors for inductive reactive power.

Power Factor Under Non-Sinusoidal Conditions

In electrical systems, mainly those that incorporate power eletronic converters such as switched-mode power supplies, the current waveform is normally nonlinear when not equipped with power factor correction circuits.

Any periodic but distorted current waveform can be represented as a Fourier series:

$$i(t) = I_1 \sin(\omega t + \phi_1) + \sum_{n=2}^{\infty} I_n \sin(n\omega t + \phi_n),$$

where I_1 is the RMS value of the fundamental component and I_n denotes the RMS values of higher-order harmonic components.

The total RMS current is then given by

$$I_{\text{rms}} = \sqrt{I_1^2 + \sum_{n=2}^{\infty} I_n^2}.$$

The active power is mainly determined by the fundamental component:

$$P \approx V_1 I_1 \cos \phi_1.$$

However, these harmonic currents increase the RMS current without providing a corresponding increase in usable (real) power. Consequently, the apparent power increased, whereas the real power remained almost constant, which reduced the power factor.

For non-sinusoidal currents, the power factor can be expressed as

$$PF = \cos\phi_1 \times \frac{I_1}{I_{\text{rms}}},$$

This can be rewritten using the total harmonic distortion as follows:

$$PF = \cos\phi_1 \times \frac{1}{\sqrt{1 + THD^2}},$$

where:

$$THD = \frac{\sqrt{\sum_{n=2}^{\infty} I_n^2}}{I_1}.$$

This formulation reveals two separate causes of power-factor degradation:

1. **Displacement power factor** ($\cos\phi_1$) - arising from the phase shift between the fundamental voltage and current.
2. **Distortion factor** ($1/\sqrt{1 + THD^2}$) - resulting from the harmonic distortion in the current waveform.

For linear inductive or capacitive loads, the power factor is mainly affected by the phase displacement. In contrast, for rectifier–capacitor input stages (e.g., non-PFC SMPS), harmonic distortion is generally the primary factor that lowers the power factor.

Measurement Considerations

Power analyzers or smart meters normally calculate the power factor from defined integration intervals. In a nonlinear load, the calculated power factor reflects the phase shift or harmonic distortion. The objective of these analyzers or smart meters is either to report power quality for compliance or as billing for large consumers of reactive power. Therefore, the calculated parameters, are limited to the intervals defined, and any change in the measured quality or reactive power is limited to this interval, for example, 15 min as per OMIE in the Iberian energy market[12].

Understanding these differences between harmonic distortion and phase shift is important for evaluating the power factor of installations, as harmonic distortion could be the driver or just a phase shift due to reactive loads.

2.2.2 Power Supply Units and Harmonic Generation in Switched-Mode Power Supplies

The electrical energy supplied was delivered as an alternating current. In a 230V (RMS), 50 Hz system, the instantaneous line voltage follows a sinusoidal waveform, swinging between the ±325V peak. On the other hand, most digital electronics, such as computers and networking hardware, operate internally from low-voltage direct current rails (e.g., 3.3V, 5V, and 12V). As a result, a power supply unit that rectifies AC to DC and transforms the voltage output is required.

Currently, computing hardware almost invariably uses switched-mode power supplies. Although these converters are compact and highly efficient, their behavior at the grid interface differs substantially from that of an ideal linear load. Specifically, the rectification and energy storage circuitry at the SMPS input stage causes the current drawn from the mains to be nonlinear, which is the dominant contributor to the harmonic distortion injected into the supply network.

AC–DC Rectification and Nonlinear Load Characteristics

A standard SMPS front-end typically consists of an electromagnetic interference filter, a full-wave bridge rectifier, and a bulk electrolytic capacitor that forms a high-voltage DC bus. After the AC input is full-wave rectified, the original sinusoidal waveform becomes a pulsating DC signal. The bulk capacitor charges only during the intervals when the instantaneous value of the rectified input voltage is higher than the capacitor's present voltage, that is, when

$$V_{\text{rect}}(t) > V_{\text{cap}}.$$

Consequently, the mains current is drawn in brief bursts near the peaks of each half-cycle. The actual power draw interval was significantly less than 180°, producing a narrow and high-amplitude current spike.

This pulsed and intermittent behavior of switched-mode power supplies functions as a nonlinear load, unlike an ideal resistive load, where the current is linearly related to the voltage. This nonlinearity is the primary cause of harmonic distortion.

Harmonic Decomposition and Total Harmonic Distortion

A non-sinusoidal waveform can be decomposed into a Fourier series related to the fundamental and integer harmonics:

$$i(t) = I_1 \sin(\omega t) + \sum_{n=2}^{\infty} I_n \sin(n\omega t + \phi_n),$$

where I_1 is the amplitude value of the fundamental and I_n is the value of the n-th harmonic.

Taking harmonic distortion into account, we can express the Total Harmonic Distortion by

$$\text{THD} = \frac{\sqrt{\sum_{n=2}^{\infty} I_n^2}}{I_1}.$$

In switched-mode power supplies without power-factor correction, the THD at the rated load can exceed 60%. The main distortion components are typically odd-order harmonics (3rd, 5th, 7th, etc.) owing to the symmetry of the rectifier.

These power supplies generate two main components into the spectral harmonics. The first is the low-order harmonics, which are multiples of the fundamental frequency, such as 50 Hz, and the high-frequency switching harmonics. The low-order harmonics are caused by the non-linearity of the load, as discussed earlier. The high-frequency switching harmonics are caused by the nature of the power supply being switched. This is because its internal MOSFETs switch rapidly to maintain a constant DC voltage output, instead of using linear power supplies that require large and costly transformers. Therefore, power supplies are required to have EMI filters to suppress this noise in the grid.

Active Power Factor Correction

To minimize this harmonic distortion and improve power quality, modern power supplies implement power-factor correction circuits. In this circuit, a regulated boost converter is then inserted between the bridge rectifier and the bulk storage capacitor.

The active PFC circuit shapes the input current to follow mains voltage waveform. Ideally, this relationship can be written as

$$i(t) \propto v(t),$$

With this, the power factor can increase and reduce the magnitude of low-order harmonics by reaching values typically above 0.95 and a THD below 10%. In comparison, without active PFC, the input current falls back to the pulsed waveform, causing increased harmonic distortion and a reduced power factor, typically falling to values of 0.6 to 0.75.

Dependence on Downstream Load

The instantaneous power required by the downstream loads determines the amplitude of the input current and, in turn, the magnitude of the corresponding harmonic components, even if harmonic distortion is introduced at the rectifier front end.

The DC/DC switching stage of an SMPS controls the output voltage in accordance with the system specifications. The energy collected by the front-end stage from the AC mains must be adjusted in proportion to changes in the computing workload, which alters the DC power drawn from the high-voltage bus.

Therefore, the amplitudes of both the fundamental and harmonic components are scaled with the output power, even if the rectifier architecture and PFC control system largely determine the shape of the current waveform. Understanding how the internal system activity can indirectly affect the measurable electrical properties at the grid interface depends on this load dependence.

The topologies of the switched-mode power supply with and without an active PFC are compared in Figure 2.1. A pulsed input current focused around the voltage peaks was produced in the non-PFC arrangement by connecting the bulk capacitor and bridge rectifier directly to the mains. By forcing the input current to follow the instantaneous mains voltage, the regulated boost converter in the active PFC configuration, on the other hand, lowers harmonic distortion and increases the power factor.

The next DC/DC conversion stage was essentially the same in both topologies, as shown in Figure 2.1. Therefore, the amplitude of the current drawn at the front end, which modulates the harmonic content observed at the grid level, is primarily impacted by changes in the internal computational load.

From the perspective of the proposed channel, the relevant distinction in Figure 2.1 is how a change in the downstream computational demand reaches the grid-side current. In the non-PFC topology, an increase in the load changes the amplitude and conduction interval of the rectifier current pulses, which can alter several low-order harmonic amplitudes simultaneously. In the active PFC topology, the control loop counteracts part of this effect by shaping the input current; therefore, the same workload variation is expected to produce a weaker low-order harmonic signature.

The current waveforms for various power-factor correction designs under both idle and loaded

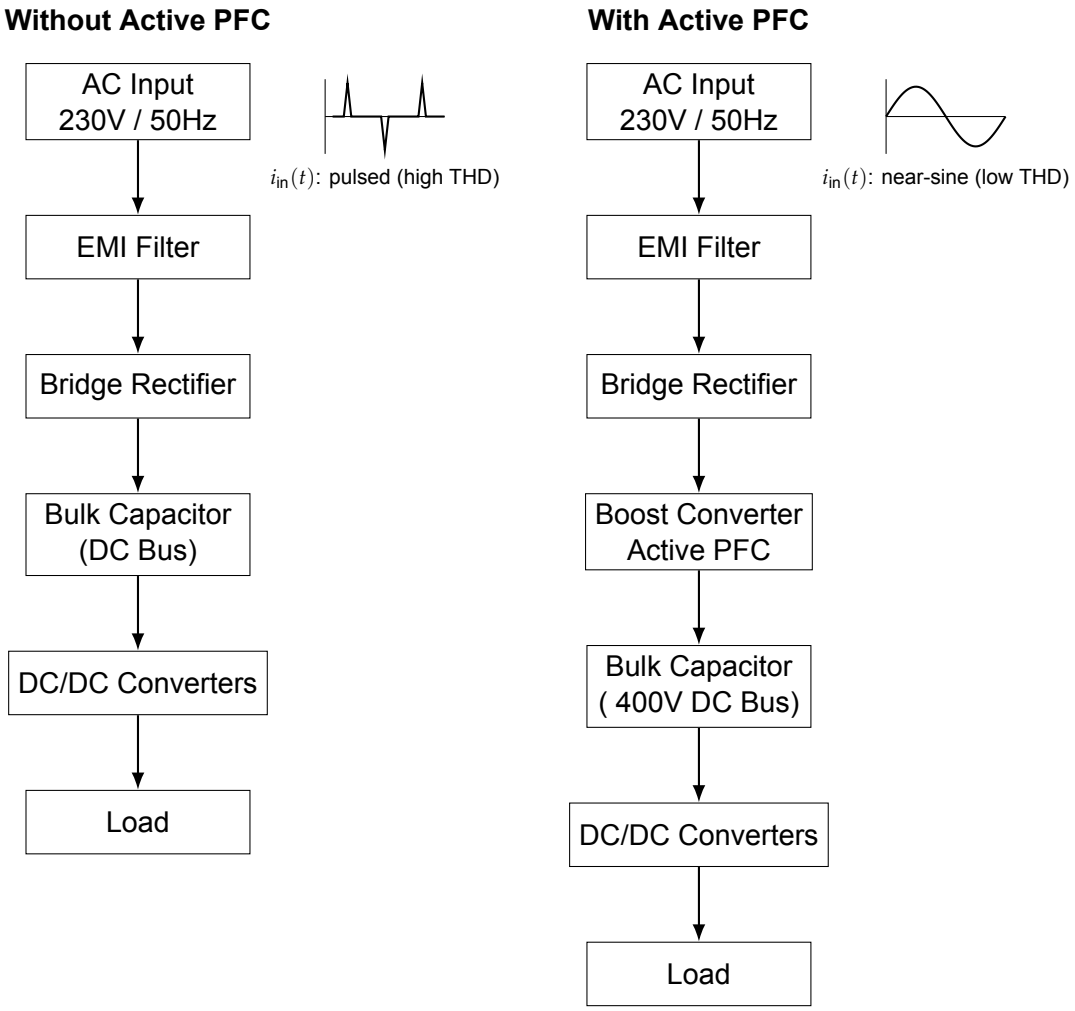

Figure 2.1: Comparison between PSU without Active PFC (left) and with Active PFC (right). Without the PFC, the current is drawn in pulses near the voltage peaks, producing a high harmonic distortion. Active PFC shapes the input current to follow the mains voltage waveform, reducing THD and improving power factor.

operation circumstances are shown in Figure 2.2. When a load was applied, the pulsed-current waveform of the supply without the PFC became even more noticeable. A pulsed current is likewise produced by the passive PFC supply, but it has a smoother shape and closely resembles a voltage waveform. The active PFC supply, on the other hand, produced the most linear current waveform; however, it was easier to observe the phase displacement angle φ between the voltage and current. This comparison provides a measured counterpart to the simplified topologies in Figure 2.1. The PFC design determines the baseline waveform shape, whereas the change from idle to load determines how strongly the shape is perturbed. These observations motivate later experiments, which test whether the resulting changes in the power factor and individual harmonic amplitudes are sufficiently structured to carry information.

Across the three rows, the idle-to-load transition primarily changes the magnitude and conduction profile of the input current, whereas the mains voltage waveform remains approximately

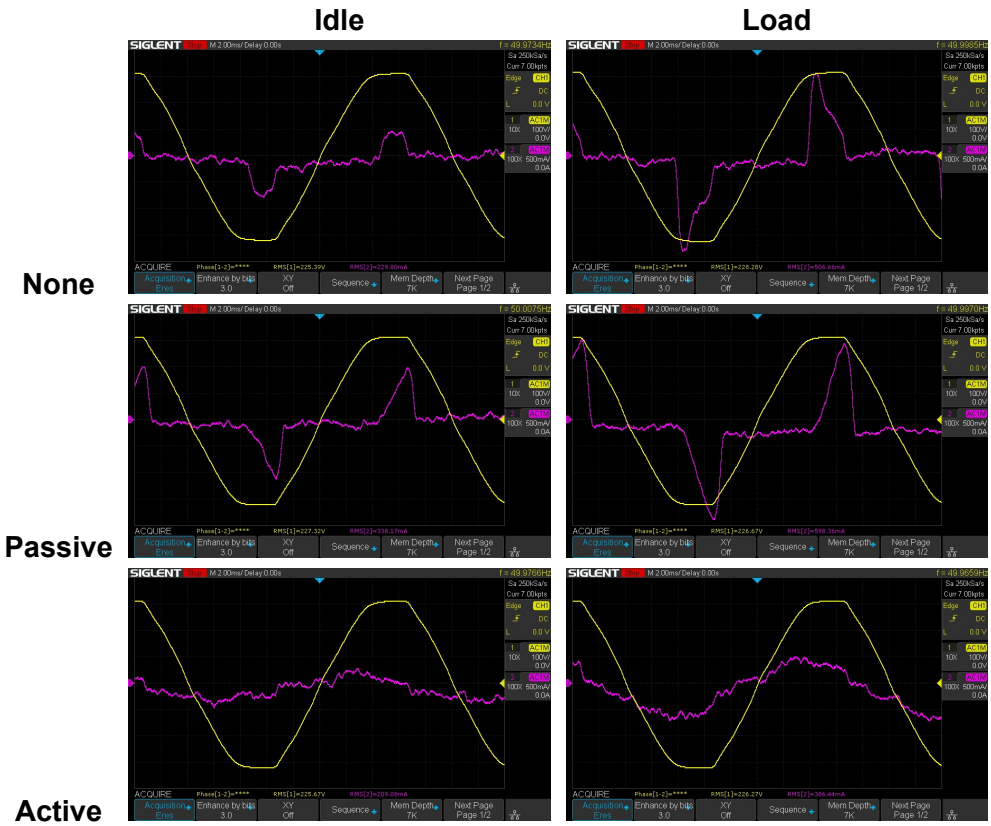

Figure 2.2: Input current waveforms for SMPS without PFC, with passive PFC and with active PFC, evaluated under both no-load (idle) and loaded operating conditions.

unchanged. For the non-PFC and passive PFC units, the loaded condition concentrates more current within particular portions of the voltage cycle, which is expected to produce stronger low-order odd harmonics. In the active PFC unit, the current remains distributed over a larger fraction of the cycle; therefore, workload changes appear as smaller amplitude and phase variations rather than as sharp current pulses.

These oscilloscope captures were used for qualitative comparison rather than as the basis of the decoder. Their purpose was to show why no single electrical metric is expected to behave identically across all PSU topologies. Therefore, the experiments in Chapter 4 derived both the absolute harmonic RMS amplitudes and the true power factors from the complete recorded waveforms. Absolute amplitudes preserve load-dependent scaling, whereas the power factor captures the combined effect of current-waveform distortion and phase displacement.

3. Related Work

Data exfiltration has been studied across network-based, physical access, and physical side-channel scenarios. A broad review of external exfiltration vectors and countermeasures was provided by Ullah et al. [13]. Conventional physical techniques, such as removable-media copying, hardware theft, or cold-boot attacks [14], depend on direct interaction with the target and are generally addressed through access control, peripheral restrictions, hardware protection, and operational monitoring. Network covert channels instead require connectivity and conceal data in legitimate-looking traffic, protocol fields, or packet timing [15]. Although they have potentially high bandwidths, they remain exposed to network filtering and traffic analysis.

However, these conventional mechanisms are not the focus of this study. The relevant scenario is a compromised air-gapped system in which normal network interfaces are unavailable, and the physical removal of data is restricted. In this setting, malware may encode information into the physical effects produced by existing hardware, including electromagnetic, acoustic, optical, thermal, magnetic, vibrational, and electrical emissions. Such channels generally trade bandwidth for stealth and depend strongly on the emitting component, receiver placement, environmental noise, and the attacker's ability to control the underlying physical process.

The remainder of this chapter therefore concentrates on representative air-gap covert channels and, in particular, on techniques that modulate the computational workload or exploit the electrical infrastructure. This focus allows the proposed method to be compared directly with the closest prior works, particularly PowerHammer [1], the NoDE channel introduced in *Your Noise, My Signal* [2], and *BREAKING TEMPEST* [16].

3.1. Air-gap Security

Table 3.1 presents a comparative summary of the representative air-gapped data exfiltration methods across several physical domains. Instead of emphasizing specific implementations, the comparison focuses on three fundamental factors that shape real-world applicability: the detectability of the technique, achievable data rate, and conditions under which it can operate effectively.

The first key observation is the fundamental trade-off between bandwidth and stealth. High-throughput optical methods, such as LED-it-GO[22] and BitJabber[27] can transmit data at rates ranging from hundreds to thousands of bits per second, but they require an unobstructed

Table 3.1: Comparison of air-gap data exfiltration research

Research / Technique	**Detectability**	**Speed**	**Feasibility Summary**
AirHopper[17] (EM via video cable)	Medium	100–480 bps	Requires malware + nearby FM receiver; works meters away from indoors.
GSMem[18] (EM in cellular bands)	Medium	∼1–2 bps	Needs CPU-modulated EM + nearby phone; low bandwidth, but stealthy.
PowerHammer[1] (power line)	Medium	10–1000 bps	Requires tapping building wiring; malware modulates CPU load.
BitWhisper (thermal)[19]	Low	1–8 bits/h	Needs two adjacent machines: very slow, but covert.
Fansmitter (acoustic fan noise)[20]	Medium	∼1 bps	Works without speakers and requires a microphone within meters.
DiskFiltration (HDD acoustics)[21]	Medium	10–200 bits/min	Use HDD actuator noise; requires a nearby microphone.
LED-it-GO (HDD LED optical)[22]	Low	100–4000 bps	Requires LOS camera; malware blinks HDD LED at high rate.
CTRL-ALT-LED (keyboard LEDs)[23]	Low	120–3000 bps	LOS required; uses keyboard LED blinking patterns.
MAGNETO / ODINI (magnetic CPU load)[24]	Low	0.1–2 bps	Requires phone magnetometer very close (cm) and bypasses Faraday cages.
Air-Fi (Wi-Fi bus emissions) [25]	Medium	∼16 bps	Uses DDR memory EM leakage; requires nearby Wi-Fi receiver.
ViBrAtIoNs[26]	Medium	0.64-5 bps	Uses FAN, speaker, or vibrator engine.
BitJabber[27]	Low	300 kbps (ideal)	Uses more intense high-frequency DRAM memory.
LANTENNA[28]	Medium	1-10 bps	Requires physical network packets into an Ethernet cable, high reliability dependent on distance.
TEMPEST-LoRa[29]	Medium	21600 bps	Uses HDMI cable electromagnetic emissions; requires a nearby LoRa receiver capable of decoding CSS signals.
SideComm[30]	Medium	1 bps	Uses processor for electromagnetic emissions; requires a nearby LoRa receiver.
Your Noise, My Signal[2]	Low	28.48 bps	Uses PFC-switched electrical noise.
Breaking-Tempest[16]	Medium	0.5bps	Uses PMBus interface to measure voltage of power supply and CPU load to detect voltage drops across power grid.

line of sight and are comparatively easy to identify when optical emissions are being monitored. At the other extreme, very stealthy channels such as BitWhisper[19] (thermal) and MAGNETO[24]/ODINI[31] (magnetic) provide only a very low bandwidth, often just a few bits per hour or fractions of a bit per second; however, they operate with minimal observable side effects.

The necessity of the attacker's close proximity was the second recurrent topic. While magnetic

channels typically require distances of only a few centimeters, electromagnetic and optical channels typically permit separation on the order of meters. Acoustic channels are in the center, and their practical range depends on the microphone qualities and background noise. Power-line channels, such as PowerHammer[1], are unique in that they take advantage of the electrical infrastructure already present in buildings rather than depending on free-space propagation, enabling reception without line-of-sight but requiring physical access to power wiring.

Hardware dependencies are the subject of the third observation. Numerous methods rely on specific hardware components (such as HDDs, fans, LEDs, video connections, and DDR buses), suggesting that changing or eliminating these components can reduce the attack surface. On the other hand, techniques based on power-line or CPU load-driven magnetic modulation rely on the basic energy consumption properties of digital computation. Although previous research covered a wide range of side channels, it mostly focused on thermal, electromagnetic, acoustic, magnetic, and optical emissions. In contrast, power quality characteristics such as power factor fluctuations and harmonic distortion as intentional signaling channels have been the subject of comparatively few systematic studies. Although the CPU load may be adjusted over power lines, as demonstrated by PowerHammer[1], the wider possibility of sculpting harmonics and controlling the power factor behavior as a function of load remains insufficiently investigated in data exfiltration.

This deficiency drives the investigation presented in the following chapter: whether intentional manipulation of the computational load can result in observable and discernible alterations in harmonic amplitudes and power factors at the grid boundary, and whether these alterations can serve as a secret communication channel in real-world settings.

This work explicitly considers the power factor change and harmonic composition as observable signaling aspects, in contrast to previous studies that have mainly altered the overall current magnitude.

3.1.1 Audio Waves

Vibrations

Guri's ViBrAtIoNs (AiR-ViBeR)[26] attack presented the first vibration-based covert channel that operates on both air-gapped and audio-gapped machines. Instead of relying on specialized vibration hardware, AiR-ViBeR leverages the inherent mechanical vibrations generated by the chassis fans of a computer. By programmatically adjusting fan speeds, malware can convert data into faint low-frequency vibrations that travel through the surface on which a computer rests.

A nearby smartphone, either already compromised by malware or simply executing JavaScript in a web browser, can capture these vibrations with its built-in accelerometer, which requires no special permissions and offers high sampling accuracy. Experiments demonstrate that this channel can transmit approximately 0.64–5 bit/s over distances of up to two meters when the devices share the same surface.

Air

Fansmitter[20] presented an innovative acoustic covert channel designed for exfiltration of data from air-gapped systems lacking conventional audio hardware. Fansmitter is different from regular acoustic attacks that use speakers or other sound output parts. Instead, it takes advantage of the fact that modern computers have cooling fans that generate noise when running normally. Malware can encode binary information in small changes to the machine's sound output by changing the rotational speed (RPM) of the CPU's heat sink fan or chassis's fan. These changes are typically invisible to those in close proximity, but they can be detected by a remote microphone, such as one found in a smartphone, a hidden recording device, or a compromised IoT sensor.

The authors explained how computer fans behave acoustically and demonstrated how different RPM values produce different sound frequencies and amplitudes. Fansmitter, the name of the study, uses this feature to create low-bandwidth digital modulation techniques that allow encoded data to be sent over the air. The attack uses only software-based fan speed management, which is supported by standard operating system APIs or driver mechanisms in most modern operating systems. The results of this study show that the channel can deliver a relevant bit rate in realistic scenarios, even when the malware must avoid noticeable fluctuations in fan noise. The researchers concluded that by deliberate control, the fan speed can produce distinct audio patterns for accurate decoding across a range of distances and in the presence of varying levels of ambient noise. This study emphasizes that acoustic data leakage cannot be prevented by audio gap defenses, which prohibit speakers and other audio devices in secure areas. Cooling fans are an ongoing and inevitable attack surface because they are essential parts that cannot be removed without endangering system performance. Systems without conventional audio hardware can still generate exploitable acoustic signals, as demonstrated by the Fansmitter technique. The receiving side of this channel is similarly discrete. Any standard microphone built into everyday consumer devices can record modulated fan noise without requiring direct access or specialized gear.

DiskFiltration[21] is a covert channel that uses the acoustic signals produced by hard disk drives to extract data from air-gapped without speakers. To encode information into the drive's operational noise, the malware deftly manipulates the actuator arm movements of the HDD to produce sounds at particular frequencies. Exfiltrated data can be detected, demodulated, and reconstructed by a nearby device, such as a smartphone with an integrated microphone. The authors validated the channel over distances of up to 2 m and attained a throughput of approximately 180 bits per minute in their experiments. This technique is particularly sneaky because it does not rely on any common audio hardware in the target system; while many secure setups purposefully disable or omit speakers, HDDs are almost universal. DiskFiltration demonstrates how presumptions about ”audio-gapped” machines can be broken by taking advantage of only the mechanical features of the drive. Additionally, this study offers a thorough explanation of how to generate and identify acoustic signals, encode data onto them, and reliably decode them in the presence of realistic noise. The authors suggested substituting solid-state drives (SSDs), which do not have moving parts that produce exploitable acoustic emissions, with HDDs as a simple defense.

The POWER-SUPPLaY[32] work demonstrates how to convert a computer's power supply unit into a loudspeaker. Malware does not use standard audio outputs. Instead, it changes the power supply's internal switching frequency, which causes its capacitors and transformers to vibrate and produce sound over a wide range of frequencies, up to 24 kHz. These sound waves can be changed to encode binary data, which a nearby device, such as a smartphone, with a microphone can then receive and decode. This method also works on systems with audio gaps because it does not need any sound card or speaker hardware. According to the test results, POWER-SUPPLaY can steal data from up to 5 m away at speeds of up to 50 bits per second when conditions are perfect. The proof-of-concept malware runs completely in the user space, so it does not need higher privileges or hardware changes. This makes it stealthy and easier to execute. POWER-SUPPLaY enlarges the threat landscape for air-gapped environments, demonstrating that even the internal components of the power supply, usually not considered as communication devices, can be abused for covert data leakage. The authors suggested ways to reduce this risk, such as filtering the power line, keeping an eye on strange switching frequencies with sensitive microphones, or intentionally adding sound noise to the PSU emissions.

3.1.2 Magnetism

Magnetic covert channels are a category of data exfiltration techniques that leverage low-frequency magnetic fields generated by the internal components of a computer. Unlike electromagnetic channels, which operate at higher frequencies and are usually attenuated by standard shielding, magnetic emissions in the extremely low-frequency (ELF) band can penetrate physical barriers, such as Faraday cages. Consequently, magnetic channels are increasingly regarded as one of the few techniques capable of bypassing both electromagnetic isolation and audio-gap defenses. In this context, MAGNETO[24] and ODINI[31] attacks represent two of the most thoroughly characterized and experimentally validated instances of such attacks.

MAGNETO demonstrated that malware can intentionally generate modulated magnetic fields by manipulating the CPU activity. Because the power drawing is correlated with CPU activity, rapid changes in the CPU load can create changes in the strength of magnetic fields, which therefore create magnetic waves. A device nearby, such as a smartphone containing a magnetometer, is sensitive enough to detect these changes and decode them into binary data. MAGNETO can achieve a data rate of 1–2 bps, depending on the distance, typically a few centimeters from the source. Despite the short distance required, the attack requires no special privileges and depends on user space privileges, making it easier to execute. This demonstrates that even in a rigorously locked system, side channels can be created from internal components and decoded by nearby devices.

The ODINI attack, which extends MAGNETO, demonstrates that magnetic covert channels can remain feasible even when a target is placed inside a farady cage. This means that preventive measures are not sufficient. It refines the modulation techniques to improve the signal under heavy shielding and shows that low-frequency emissions can traverse the enclosures of the farady cage with limited attenuation. Although its bit rate is low, similar to MAGNETO, ODINI's ability to operate and break out of shielding forces the need to implement detective controls rather than just preventive ones.

Combining these two studies, they demonstrated that magnetic emissions can be a persistent risk and are often overlooked in data exfiltration fields. The magnetic fields can penetrate the shielding, making the detective controls more necessary to mitigate the risk of exfiltration.

3.1.3 Electromagnetic Waves

Because of their switching behavior, all electronic components of a computer, including CPUs, memory buses, display cables, and power circuits, naturally release electromagnetic radiation. Electromagnetic covert routes make use of this. These emissions cover a wide range of frequencies, from higher-frequency RF leakage from data buses and video links to low-frequency kilohertz harmonics brought on by power fluctuations. Malware can purposefully alter these emissions to encode data because software-driven activities control the structure of this activity. The signals can then be intercepted and decoded by a nearby receiver, such as a smartphone, radio, Wi-Fi interface, or software-defined radio device. Because it does not rely on network infrastructure, storage media, or direct physical access, this enables an exfiltration method that overcomes conventional isolation defenses. As a result, even in air-gapped systems, unintentional EM radiation can act as a conduit for the exfiltration of information.

EM-based attacks are particularly alarming because they can be executed without elevated privileges, extra hardware, or obvious changes in the behavior of a target. Prior studies, including GSMem[18], AirHopper[17], Air-Fi[25], BitJabber[27], and LCD TEMPEST[33], have shown that CPU utilization patterns, memory accesses, and display activities can all be used as RF transmitters, each with its own throughput and distance range. For instance, video cables, such as HDMI and VGA, can emit signals in the FM broadcast band, DRAM buses can produce RF signals in Wi-Fi frequency bands, and CPU instruction sequences can emit GSM-encapsulated packets that nearby GSM phones can detect and decode. Although these channels can vary in bandwidth and reliability, they all come from the same fundamental reality: electromagnetic emanations are intrinsic to electronics in digital systems and cannot be fully removed. Therefore, defending against EM covert channels requires detective or preventive measures, such as physical shielding, zoning or physical segregation, spectrum monitoring, and improved hardware and circuit design, rather than purely turning off radios or disconnecting network interfaces.

The GSMem[18] attack demonstrates how standard computers, when powered by carefully selected CPU and memory usage patterns, can unintentionally radiate electromagnetic signals in the GSM cellular bands. Malware can encode data by modulating the amplitude of the electromagnetic harmonics produced when the processor repeatedly executes instruction sequences. These emissions can then be detected by a nearby mobile phone whose baseband already uses these GSM frequencies, and low-bandwidth payloads can be decoded without the need for firmware modifications or increased privileges. This indicates that consumer electronics that are readily available and situated close to air-gapped systems can function as covert receivers.

Because the transmitter in GSMem is fully implemented in the software and operates without the need for administrative privileges, the attack is both deployable and stealthy. The achievable data rate is modest, usually only a few bits per second, yet sufficient for exfiltrating small secrets, such as encryption keys or passwords. From the receiver side, the attack can use built-in diagnostic features on a mobile phone that can be used by the software installed on the mobile phone. This study underscores the broader danger arising from unintentional electromagnetic emissions created by standard hardware. While shielding and strict physical access can reduce RF emissions or the likelihood of receiving them, and therefore shrink the attack surface, GSMem still demonstrates how minor lapses in the process or physical controls, such as permitting personal smartphones or hidden IoT devices in proximity to isolated machines, can permit these covert channels for data exfiltration. This emphasizes the importance of electromagnetic isolation security controls and security architectures that account for not only explicit antennas or wireless components but also the emission patterns generated by CPUs during their normal operation.

The AirHopper[17] attack demonstrated a covert channel that uses electromagnetic interference produced by a computer's video cable, similar to a VGA. Malware can produce RF emissions at the FM radio frequency by carefully manipulating the screen output. Nearby cellphones with integrated FM receivers can detect these signals. Thus, AirHopper allows data exfiltration from air-gapped devices and transforms a television cable into an unintentional radio transmitter.

This method's relatively large bandwidth compared with many other electromagnetic hidden channels, with a practical throughput frequently approaching several hundred bits per second, is one of its main advantages. Small secrets and larger data volumes, such as documents and keystroke records, can be transmitted because of this capability. Because the modulation does not noticeably change the on-screen content, malicious software can remain undetectable to consumers. Moreover, enhanced privileges beyond the ability to manipulate or alter the display output are not required for the attack.

AirHopper demonstrated how seemingly innocuous peripherals can act as potent electromagnetic radiation emitters. For example, display cables are often long, have little shielding, and contain many high-frequency transitions, which makes them ideal for unintentional radio frequency creation. This method shows that the display hardware can function as a covert RF transmission conduit, even if traditional wireless interfaces are turned off or physically removed. This insight has important implications for electromagnetic security and emphasizes the need

for administrative controls, shielding, and filtering to reduce video cable exposure in secure settings. DRAM modules are abused as a source of unwanted electromagnetic radiation in an Air-Fi[25] attack. In the 2.4 GHz band, which is also utilized by Wi-Fi networks, observable RF emissions are produced by DRAM activity, especially fast row switching and significant memory bus usage. Air-Fi manipulates this behavior by creating particular memory access patterns that encode data as AM-like fluctuations. The signal can then be received and demodulated by a nearby Wi-Fi-capable device, enabling data exfiltration from air-gapped workstations, even in cases where the transmitting system lacks a Wi-Fi interface.

One of the main advantages of Air-Fi is that it takes advantage of routine tasks carried out by the CPU and memory controller, making it more difficult to discern from genuine activity. Because DRAM is continually accessed by current computers, unless specialized RF analysis techniques are used, the concealed signal is obscured by regular noise. Furthermore, the attack does not require elevated privileges to initiate high-frequency memory operations and can be initiated from the user space. It can reach a throughput of 16 bit/s under appropriate circumstances. Air-Fi expands the range of electromagnetic side channels by proving that memory subsystems can function as useful RF transmitters. This is especially crucial, as DRAM is ubiquitous and necessary in computing devices, even those that are protected or severely constrained. The utilization of the 2.4 GHz band emphasizes how difficult it is to protect air-gapped systems because spectral leakage from fast-switching digital components cannot be stopped by simply disconnecting Wi-Fi devices. To keep prospective receivers at a safe distance from the device, Air-Fi advises countermeasures, such as adding shielding, ensuring physical separation, or randomizing DRAM access patterns.

Although conceptually comparable to Air-Fi, the BitJabber[27] work explores an electromagnetic covert channel produced by modulating activity on the memory bus, which can achieve even greater data speeds owing to more intensive access patterns. BitJabber takes advantage of the fact that when modern high-speed DRAM interfaces encounter bus contention or are purposefully overloaded with memory, they produce audible RF harmonics. A neighboring software-defined radio receiver can demodulate the modulated RF signal created by malware by inducing particular sequences of cache misses, memory bursts, or planned interference. According to experimental findings, BitJabber can achieve a significantly higher bandwidth than many current RF-based covert channels. Under ideal laboratory settings, it can reach 300 kilobits per second. This makes it one of the quickest electromagnetic channels to be discovered. However, odd CPU consumption patterns or performance-monitoring counters (PMCs) may

more readily detect such aggressive memory-access behavior. However, this study shows that with careful adjustment, high-throughput electromagnetic exfiltration is still possible.

Research on LCD TEMPEST[33] examines how electromagnetic emissions from display hardware, particularly LCD monitors, can be used to secretly transmit encoded data or recover on-screen content. LCD drive electronics generate distinctive EM leakage fingerprints that correspond to pixel changes and screen refresh activity because they run at relatively high frequencies. Depending on how the emissions are modulated, an attacker with the correct RF receiving equipment can capture these signals from a distance, allowing either active data exfiltration or passive eavesdropping. Modern flat-panel screens produce less radiation than older CRT displays but are not enough to completely eliminate security issues, according to a key result in the LCD TEMPEST work. Attackers can reconstitute words, pictures, or graphical pattern fragments from a distance of several meters under ideal low-noise radio frequency conditions. Additionally, some research has demonstrated that viruses can subtly alter pixel intensities to include secret information in the sent signal without causing the user to notice any changes. Therefore, even without specialized exfiltration hardware, the display essentially becomes an unwanted RF transmitter. A more general idea is demonstrated by LCD TEMPEST: visual display displays can act as electromagnetic side channels to get beyond air-gap safeguards. Monitors are essential user-facing components, making it challenging to protect them effectively without compromising their usability. As a result, this collection of studies emphasizes the significance of stricter electromagnetic hardening regulations, safe video signal cabling, and physical zoning procedures that regulate the location of possible receivers. LCD TEMPEST assaults show that visual output remains a significant EM leakage vector despite enhanced shielding and low-emission designs in contemporary electronics.

TEMPEST-LoRa[29], which uses inadvertent electromagnetic emissions from video interfaces to illustrate cross-technology covert communications, is another noteworthy development in electromagnetic covert channels. This study shows that electromagnetic signals suitable for LoRa (long-range) wireless receivers may be produced by controlling HDMI cables connected to a hacked machine. Malware can use Chirp Spread Spectrum (CSS) modulation, a physical-layer technique used by LoRa communication networks, to encode data in the radiated signal by carefully manipulating pixel transmission patterns across the HDMI channel. Consequently, the HDMI cable serves as an unintentional transmitting antenna, emitting signals that commercial LoRa receivers can decode. The capacity of TEMPEST-LoRa to link two seemingly unrelated domains–video display interfaces and low-power wide-area network (LPWAN) communication protocols–is its fundamental contribution. Because of the high sensitivity and wide

coverage of LoRa technology, the use of LoRa receivers allows signal detection at significantly longer ranges than many other electromagnetic hidden channels. Experiments conducted by the authors demonstrated that this method can provide secret data exfiltration from air-gapped computers over distances significantly greater than those typical of traditional attacks based on the TEMPEST standard. This study further emphasizes how inadvertent electromagnetic leakage from common computer connections can be used to create secret channels that get around typical network isolation.

A covert channel attack called LANTENNA[28] uses electromagnetic emissions from Ethernet cables. The main idea is that Ethernet twisted-pair cables can act as unintentional antennas and generate radio frequency signals when powered by carefully designed network traffic patterns. A nearby radio receiver can wirelessly record the modulated electromagnetic radiation that results from malware on an air-gapped system encoding data into patterns of network activity that excite the Ethernet connection. The authors showed that the malicious program can operate in user space, does not need higher rights, and can even execute within a virtual machine while remaining undetectable.

The intentional use of processor-induced electromagnetic side channels as wireless communication interfaces is another line of inquiry. Feng et al. presented SideComm[30], a system that turns a microprocessor's natural electromagnetic emissions into an inexpensive radio transmitter. This approach uses carefully crafted software instruction sequences that are run on the CPU to produce regulated RF signals instead of depending on specialized RF components. The system generates periodic electromagnetic pulses with spectral characteristics that can be detected by a nearby receiver by manipulating the instruction execution patterns of the processor. This makes it possible for low-resource embedded systems to send data wirelessly without the need for traditional radio circuits or antennas. Additionally, the authors demonstrated how controlled software execution can be used to tune these emissions to resemble LoRa chirp spread-spectrum signals. This technology generates frequency-swept waveforms similar to LoRa chirps by varying the timing of the nested loops, allowing packets that can be deciphered by LoRa-style receivers to be transmitted. According to experimental findings, this side-channel-based communication can be detected across a distance of more than 10 m, even in non-line-of-sight situations such as around corners or through walls. This study highlights the wider security dangers of CPU-generated electromagnetic emissions and their capacity to allow hidden wireless channels, even though the primary goal is to facilitate communication for limited IoT platforms.

Electric Current

To date, as this thesis is being written, three studies have been identified that use electric current as a modulation channel for data exfiltration, one of which was published after this dissertation was completed. One was the PowerHammer[1] study, which precisely demonstrated this approach. Two attack variants are introduced based on the point at which the power lines are tapped: line- and phase-level power hammering. In line-level attacks, the adversary installs a sensor near the power outlet feeding the target machine, allowing a comparatively high exfiltration throughput, experimentally reaching approximately 1,000 bits per second. When tapping at the building's main electrical service panel (phase level), the attacker can still decode the transmitted signals, but the achievable rate drops substantially (to approximately 10 bit/s) because of the noise from other loads on the power network. The observed bit rates indicate that PowerHammer can be used to leak both small, high-value secrets and, given sufficient time, larger datasets, depending on the attacker's physical and infrastructural access.

The second key work was *Your Noise, My Signal*[2], which introduced a NoDE-based covert channel. Unlike PowerHammer, which directly measures current fluctuations, this approach leverages high-frequency voltage ripples produced by power-factor correction circuits in desktop PSUs. The receiver does not need to sit in the target's direct current path; instead, it can monitor the voltage at another outlet on the same building power network and isolate the switching-noise band of the transmitter. This improves stealth and allows operation from a different room without a line of sight. The authors reported data rates of up to 28.48 bit/s over approximately 27.4 m and further demonstrated concurrent exfiltration from multiple machines by exploiting the near-orthogonality of their switching frequencies. Relative to this dissertation, the study is similar in that both exploit workload-induced changes in a system's electrical behavior, but it differs in the observable: NoDE focuses on high-frequency switching noise from active PFC, whereas this work examines lower-frequency power quality metrics, such as power factor and harmonic content.

A more recent advance is the bidirectional power-line covert channel *BREAKING TEMPEST*[16]. It extends previous power-line exfiltration research in two ways. First, it enables outbound leakage from the target to the attacker via CPU load modulation of the line current and inbound communication by modulating the supply voltage delivered to the target. Second, it shows that the channel remains usable even in the presence of TEMPEST-grade power filters, indicating that low-frequency modulation near the mains frequency can traverse filters primarily designed for higher-frequency conducted emissions. Their prototype used current sensing for

the outbound path and controlled voltage variation for the inbound path, achieving up to 0.5 bit/s outbound and 2 bit/s inbound. Compared with this dissertation, it is closer to the NoDE work in both the physical domain and threat model, since all rely on low-frequency effects observable in building power infrastructure or high-frequency noise in the case of NoDE. However, their goals differ. Instead of building a fully interactive raw power-line channel, this study examines whether standard power quality observables can serve as information-bearing quantities for covert signaling.

Light

A light-based covert channel is a type of side-channel communication in which a compromised system deliberately modulates optical signals to transmit data across an air gap without relying on standard communication interfaces. These signals can fall within the visible spectrum (e.g., monitor displays and status LEDs) or the invisible spectrum (e.g., infrared LEDs) and are produced by hardware components that are already present and required for normal system functionality. Malware encodes information by adjusting parameters such as light intensity, color, contrast, flicker rate, or timing in ways that evade human perception but can be recorded and interpreted by optical sensors, including cameras, photodiodes, and light detectors.

VisiSploit[34] proposed a new optical covert channel that exploits the limits of human visual perception to exfiltrate data from air-gapped machines by using a standard LCD monitor. Conventional optical side channels typically depend on clear visible light flashes or LED signaling, which are easy to notice and therefore not truly covert. By contrast, VisiSploit hides the encoded binary information in the screen output by subtly altering the contrast levels and flicker frequencies, which remain invisible to the human eye. Consequently, a compromised system can display concealed data on its screen without raising the user's suspicion.

The LED-it-GO[22] study describes a covert channel that leaks information from air-gapped computers by controlling internal indicator LEDs, most notably hard-drive activity lights. Unlike approaches that visibly blink displays or status LEDs in a manner that users can detect, LED-it-GO embeds data in high-frequency blinking sequences that take advantage of the rapid switching capability of LEDs, staying below the threshold of human perception. Nearby cameras or light sensors can capture and decode these sequences, thereby enabling the retrieval of sensitive data without direct physical access or network connectivity to the device.

The aIR-Jumper[35] research presented a covert channel that uses near-infrared (NIR) light and surveillance cameras to connect and infiltrate air-gapped networks. In contrast to channels

based on visible light, this method relies on IR radiation in the 800–900 nm band, which is invisible to humans but detectable by standard CMOS camera sensors, thereby ensuring a high degree of stealth in the detection process. The authors demonstrated that modern security cameras, which are widely deployed in high-security locations, can unintentionally function as transmitters and receivers in air-gap attack scenarios.

4. Power Factor and Harmonics As Covert Channel

This chapter presents an exploratory proof-of-concept evaluation of the power factor and harmonic distortion as covert-channel observables. The goal is to determine whether controlled workload changes can create measurable patterns at the electrical grid interface and whether a recorded transmission can be recovered under the evaluated conditions. The experiments were not intended to provide a statistical characterization of long-term channel reliability.

4.1. Experimental Setup and Instrumentation

In this section, we outline the methodology used to assess whether variations in the power factor and harmonic distortion can serve as covert channels. The experiments were conducted in a controlled environment in which the electrical behavior of the loads was observed while executing the managed workload patterns. The main objective of the measurement was to establish whether these load patterns could induce detectable changes owing to harmonic distortion and a power factor phase shift.

In the experiments, a switched-mode power supply was used for the target system. The voltage and current were measured using an oscilloscope, and their waveform data were saved for offline processing. The power consumption was changed by adjusting the CPU workload. With these changes, we measured and calculated the power factor, total harmonic distortion, and its individual parts with the amplitude.

In this study, three workload scenarios are explored. First, we collected background and normal operation data without altering the power drawn from the power supply. Next, we induced workloads into the CPU using different patterns, such as pulses, PWM-like, and ramps, without encoding any particular data, measured the data, and mapped it into a heatmap and time-series graphs. Finally, we used encoded On-Off Keying based modulation, with three payloads: One string ’Ola futuro mestre!’ - ASCII encoded, an AES-128 symmetric example key with 0xdeadbeef repeating patterns and one secret ECDSA key.

Data were collected using a digital oscilloscope with voltage and current probes at the terminals of the power supply under test and at the main office entrance. The collected data were then processed to generate visual representations and calculate power metrics such as active power, apparent power, power factor, displacement power factor, and harmonic content for each nth multiplier, either relative to the fundamental (dBc) or in absolute terms (Amps RMS).

4.1.1 Testbed Overview

The experimental setup comprised a computing device powered by the PSU under investigation and external measurement equipment that monitored the electrical signals at the PSU input. The input voltage and current were recorded concurrently to enable time-aligned analysis of the relevant electrical quantities.

A schematic of the measurement setup is shown in Figure 4.1. The voltage probe senses the mains voltage delivered to the PSU, whereas the current probe is inserted in series with the phase conductor feeding the device.

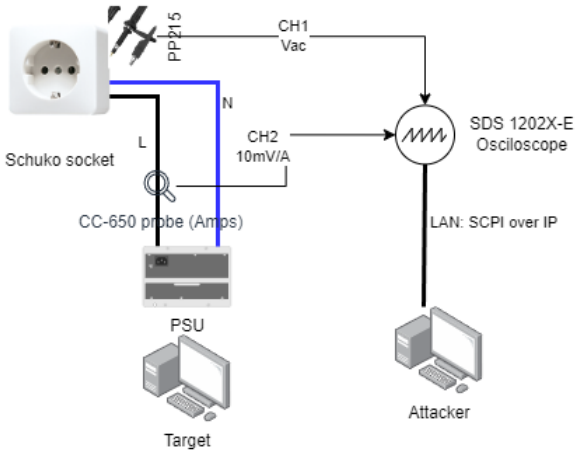

Figure 4.1: General Layout

During each experimental run, a predetermined workload pattern was executed on the target device. Changes in the computing load appeared as changes in the measured waveforms. By collecting this data and processing it offline, we can determine whether the patterns have induced detectable variations, and how and later decode it for a viable demonstration.

During the experiments, measurements were continuously conducted. over the payload size period. The raw waveform data were processed into a structured data h5 file containing representations of the waveforms, timestamps, and other metadata. This data file was then used to process the observables using a discrete Fourier transformation.

Figure 4.2 shows the second measurement site on the test bed. Multiple loads are connected to the office mains, and the current and voltage are recorded at the line entry; thus, any change in the load or power consumption is captured.

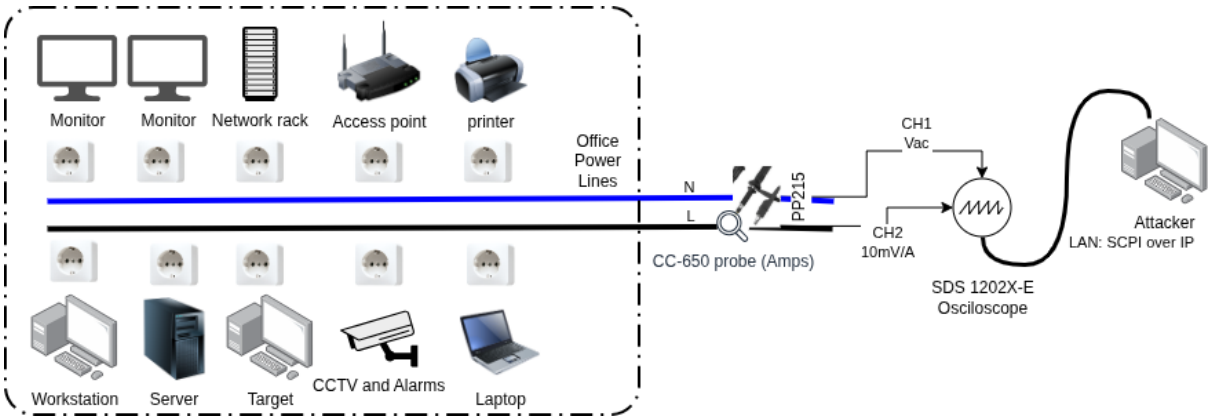

Figure 4.2: Office measurement location and loads

4.1.2 Instrumentation and Acquisition Parameters

Electrical measurements were performed using a digital storage oscilloscope configured to record the synchronized voltage and current waveforms. Two independent input channels were used to connect the oscilloscope to the measurement probes, enabling the concurrent monitoring of electrical quantities at the PSU input.

The primary measurement device employed in the tests was a SIGLENT SDS 1202X-E. The voltage was measured using a PP215 probe, and the current was measured using a Hantek CC-650. The voltage probe was connected in parallel to the PSU input terminals, and the current probe was installed in series with the phase conductor feeding the unit.

Waveform acquisition was performed periodically, and time-domain segments of both voltage and current were recorded. The oscilloscope was operated at a sampling frequency of 100kSa/s for 70s or 10kSa/s for 700s, and each recorded segment contained 7 000 000 samples. New waveform segments were captured continuously, except in the identified cases, enabling continuous tracking of the system's electrical behavior. The primary constraint on the acquisition rate was the oscilloscope's capability to capture data and transfer it to the measurement computer, as well as its resolution and memory capacity. Consequently, longer acquisition periods were traded off with a lower resolution.

The acquired waveforms were transferred to a host PC via the remote control interface of the oscilloscope and stored for subsequent offline analysis. Dedicated software was employed to calculate electrical quantities from the recorded signals, including the RMS voltage and current, instantaneous power, power factor, displacement power factor, and total harmonic distortion.

Harmonic analysis was conducted by applying a discrete Fourier transform to the measured waveforms. From the resulting spectrum, the magnitude of each harmonic component was

extracted and normalized with respect to the fundamental component. Harmonic magnitudes were then expressed in decibels relative to the fundamental (dBc) as

$$H_{dBc} = 20 \log_{10} \left(\frac{H_n}{H_1} \right)$$

where H_n denotes the amplitude of the n-th harmonic and H_1 is the amplitude of the fundamental.

These harmonic values were logged and arranged into time-indexed matrices, enabling heatmap visualization of their temporal evolution.

4.1.3 Devices Under Test

We conducted experiments using three different power supplies. One without any power factor correction, another with passive power factor correction, and finally with an active one. The objective of using different power-factor correction circuits is to understand how these circuits can affect the observation and viability of the proposed covert channel.

For each PSU, the same target was powered while maintaining the same workload and conditions. The target consisted of a desktop computer with an AMD Ryzen 5 3600 processor with six cores, running Windows 11 Education 25H2. The complete list of settings is presented in Table 4.1.

Component	**Target**	**Attacker**
CPU	AMD Ryzen 5 3600 6-Core Processor	AMD Ryzen 5 5500
Operating System	Windows 11 Education 25H2	Windows 11 Pro 25H2
RAM	8GB RAM	64GB RAM
Motherboard	A520M-A PRO (MS-7C96)	PRIME B550M-A
PSU	UNYKAch ATX 500W (no PFC) / 1Life PS:JET 600W ATX (Passive PFC) / MSI MAG A650BN 650W ATX (Active PFC)	Nox Hummer X650W 80+Gold

Table 4.1: Components of target and attacker

Table 4.2 summarizes the main characteristics of the tested power supplies.

PSU ID	Architecture	Rated Power	Notes
PSU-A	SMPS without PFC	500W	UNYKAch ATX
PSU-B	Passive PFC	600W	Life PS:JET 600W ATX
PSU-C	Active PFC	650W	MSI MAG A650BN 650W ATX

Table 4.2: Summary of the power supply units used in the experimental evaluation.

4.1.4 Workload Modulation Methods

To study how computation relates to electrical properties, controlled workloads were applied to the device to produce predictable changes in CPU load and enable the measurement of power-supply behavior under different conditions. A Python script modulated the workload by adjusting the number of active CPU threads and processor duty cycles, creating deterministic load profiles with well-defined timing within each execution.

Three types of workload patterns were examined. The first experiment employed a gradual sweep of CPU utilization, in which the processor load was ramped from idle to full usage over time. This setup was used to characterize the effect of changes in the computational load on the power quality metrics, such as the power factor and harmonic distortion.

In the second experiment, transient workloads consisting of rectangular CPU load pulses were used. The system alternated between idle and full load while the pulse duration increased from approximately 1 ms to 500 ms, allowing observation of the power supply's transient response and related electrical indicators.

The third experiment implemented a covert signaling method that conveyed messages by modulating CPU activity using On–Off Keying. Logical symbols are encoded as low or high CPU utilizations. The transmitted message was the ASCII string *"Ola futuro mestre!"*, converted to a binary sequence and framed with a synchronization preamble and a cyclic redundancy check. In later tests, the CRC was replaced with the Reed–Solomon error correction using four parity bytes.

The principal OOK experiments used a symbol duration of $T_s = 2\,\mathrm{s}$, whereas an additional $T_s = 1\,\mathrm{s}$ configuration was evaluated for the AES-128 payload. These values are conservative engineering choices for demonstrating the channel rather than the result of an optimization sweep. An alternating-bit preamble was employed to support synchronization during detection. The transmitted frame used a 16-bit synchronization preamble, followed by a size field, payload,

and Reed–Solomon redundancy. The rationale for the selected symbol duration is discussed in section 4.4.4.

These workload profiles support an exploratory assessment of the system's electrical behavior and whether deliberate workload modulation can convey recoverable information in the evaluated captures.

4.1.5 Experimental Design and Scope

The experimental evaluation was designed as an exploratory proof-of-concept intended to determine whether workload-induced changes in power quality observables can carry recoverable information. It was not designed as a statistical characterization of channel performance across repeated independent trials.

Each combination of payload, symbol duration, PSU topology, and receiver placement was transmitted and acquired once. Therefore, the number of independent runs for each experimental configuration was $N = 1$. In this dissertation, a run denotes one complete workload execution, waveform acquisition, feature extraction process, and decoding evaluation. The different electrical features extracted from the same waveform capture, including individual harmonic amplitudes, power factor, active power, and apparent power, are alternative receiver observables and are not independent repetitions.

Consequently, the reported bit error rates, signal-to-noise ratios, and decoding outcomes are the descriptive values for the corresponding recorded transmission. No mean, median, variance, standard deviation, or confidence interval across independent runs was reported. In particular, the BER value represents the observed fraction of erroneous bits in one decoded frame and should not be interpreted as an estimate of the expected BER of the channel.

Therefore, this study prioritizes physical interpretation and end-to-end feasibility. The results provide evidence that the proposed mechanism can operate under the evaluated conditions, but they do not establish run-to-run repeatability, long-term reliability, or statistical significance between PSU topologies or receiver placements. Such claims would require multiple independent transmissions for each experimental configuration.

4.2. Measurement Scenarios and Receiver Locations

Figure 4.3 illustrates the office-level measurement scenario and the approximate physical distances between the target area, oscilloscope measurement point, and main electrical panel.

The drawing is not to scale and should be interpreted as a schematic representation of the receiver placement rather than as a precise architectural survey.

The target area included the workload-modulated system and other nearby electrical loads connected to the same office environment. The oscilloscope represents the office-level measurement point used to evaluate whether the conducted signal remains observable from the immediate PSU connection. The red path indicates the electrical line between the main panel and the measurement/target area.

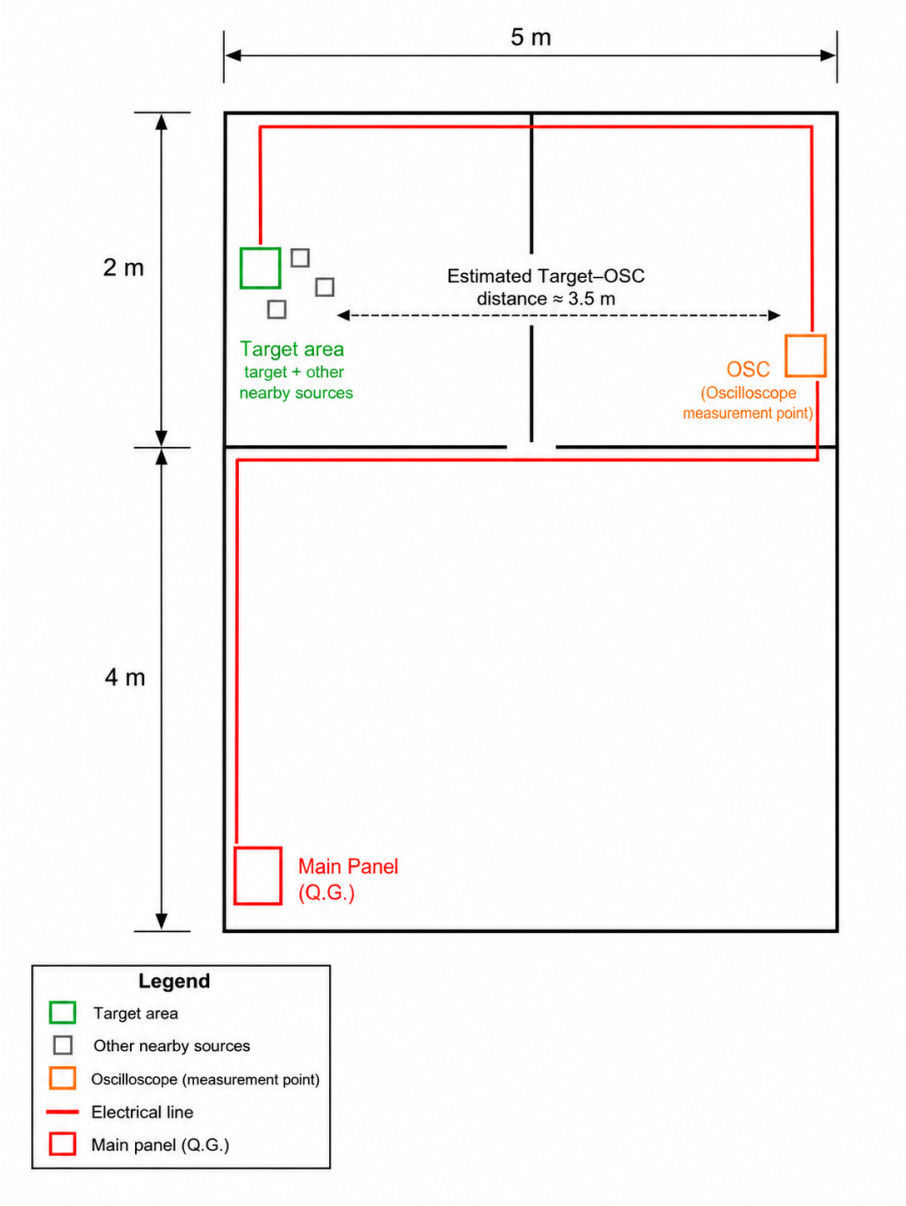

Figure 4.3: Representative office floor plan and electrical path for the office-level measurement scenario. Distances are approximate and the drawing is not to scale.

Table 4.3 summarizes the measurement setup, sampling rate, duration, transmission type, and transmitted payloads used in the experiments.

Table 4.3: Experimental configurations used in the exploratory proof-of-concept evaluation. Each listed PSU and receiver-placement configuration, where applicable, was acquired once.

Test	Sampling rate	Capture	Transmission	T_s	Runs	Description
1	100 kSa/s	70 s	Baseline	–	1	Background measurement with no payload transmitted.
2	100 kSa/s	70 s	CPU load sweep	–	1	CPU load sweep from 0 to 100% during the first 30 s.
3	100 kSa/s	70 s	CPU pulse ramp	–	1	PWM-like CPU load pulses of 100% during the first 30 s.
4	10 kSa/s	700 s	ASCII payload	2 s	1	16-bit preamble + “Ola futuro mestre!” + 4 Reed–Solomon parity bytes.
5	10 kSa/s	700 s	AES-128 payload	2 s	1	16-bit preamble repeated twice + `0xdeadbeef` $\times 4$ + 4 Reed–Solomon parity bytes.
6	10 kSa/s	700 s	AES-128 payload	1 s	1	Same frame as Test 5, evaluated with shorter symbol duration.
7	10 kSa/s	700 s	ECDSA payload	2 s	1	16-bit preamble + 32-byte ECDSA secret + 4 Reed–Solomon parity bytes.

For the OOK experiments, one waveform capture was obtained for each payload, symbol duration, PSU topology, and receiver placement combination. The “Runs” column therefore reports independent experimental executions, not the number of electrical features later extracted from each capture.

4.3. Symbol Encoding Model

The transmitter is modeled as a binary source that modulates the target system’s computational load over time using On-Off Keying: each bit maps to one of two CPU load levels held for a fixed symbol duration T_s (see Algorithm 1). The receiver does not observe these signals directly but infers them from their effect on power quality metrics, including harmonic amplitudes, power factor, active power, and apparent power.

Each transmission uses the frame

```
| preamble | size | payload | reed-solomon |
```

where the preamble is a known alternating bit pattern for synchronization, the payload holds the test message, and the RS field adds four Reed–Solomon parity bytes. This format supports the assessment of symbol-level detectability and full payload recovery.

Algorithm 1 On-Off Keying CPU based modulation

Require: Binary sequence $\mathbf{b} = [b_0, b_1, \ldots, b_{N-1}]$
Require: Symbol duration T_s
Require: CPU utilization levels u_0 and u_1
Require: Initial delay Δ_{start}
Ensure: CPU utilization waveform $u(t)$ encoding the symbol sequence

$t_{ref} \leftarrow$ MonotonicClock()
$t_0 \leftarrow t_{ref} + \Delta_{start}$
for $i = 0$ to $N - 1$ **do**
    $t_{start} \leftarrow t_0 + iT_s$
    $t_{end} \leftarrow t_0 + (i + 1)T_s$
    WaitUntil(t_{start})
    if $b_i = 0$ **then**
        SetCPULoad(u_0)
    else
        SetCPULoad(u_1)
    end if
    WaitUntil(t_{end})
end for
SetCPULoad(u_{idle})

The frame definition also allows the nominal communication rate to be separated from the payload goodput. For binary OOK, one bit is transmitted per symbol, so the raw rate is

$$R_{\text{raw}} = \frac{1}{T_s},$$

For a frame containing N_{frame} transmitted bits and N_{payload} payload bits, the frame duration and protocol payload goodput are

$$T_{\text{frame}} = N_{\text{frame}} T_s, \qquad G_{\text{payload}} = \frac{N_{\text{payload}}}{T_{\text{frame}}},$$

Table 4.4 applies these definitions to Tests 4 and 7, respectively. The frame lengths include the synchronization preamble, one 8-bit field, payload, and four Reed–Solomon parity bytes. Tests 5 and 6 used the repeated 32-bit preamble described in Table 4.3; the other payloads used a 16-bit preamble.

The table reports the protocol goodput during the transmitted frame. As each acquisition lasted 700 s and contained one payload frame, the corresponding capture-average payload rates were approximately 0.0257 B/s for the ASCII payload, 0.0229 B/s for either AES-128 test, and 0.0457 B/s for the ECDSA secret. The difference between the protocol goodput and capture-average

Table 4.4: Nominal rate and payload-goodput metrics for the OOK proof-of-concept transmissions.

Test	Payload	Payload (bytes)	T_s (s)	Frame (bits)	Frame time (s)	Raw rate (bit/s)	Payload goodput bit/s (B/s)
4	ASCII string	18	2	200	400	0.5	0.360 (0.0450)
5	AES-128 example key	16	2	200	400	0.5	0.320 (0.0400)
6	AES-128 example key	16	1	200	200	1.0	0.640 (0.0800)
7	ECDSA secret	32	2	312	624	0.5	0.410 (0.0513)

goodput reflects the acquisition interval outside the encoded frame, including the synchronization delay and the remaining idle portion of the capture.

For a decoded feature with an observed pre-correction bit error rate $\mathrm{BER}_{\mathrm{obs}}$, the observed number of erroneous bits per second over the compared bit interval is

$$E_{\mathrm{obs}} = \mathrm{BER}_{\mathrm{obs}} R_{\mathrm{raw}} = \frac{\mathrm{BER}_{\mathrm{obs}}}{T_s},$$

This is a descriptive conversion for the recorded frame and not an estimate of the long-term error process. For example, the best valid PF/harmonic result with $\mathrm{BER}_{\mathrm{obs}} = 0.0069$ at $T_s = 2\,\mathrm{s}$ corresponds to approximately 0.00345 erroneous bits per second before error correction or one observed error per approximately 290 s of compared symbols. A reported BER of zero means that no differing bits were observed in the decoded interval; it does not establish a zero expected BER for future transmissions.

Because the receiver measures the electrical quantities estimated over finite waveform windows rather than the instantaneous CPU command, symbol transitions appear as smoothed changes rather than as ideal steps. Their visibility depends on the PSU dynamics, chosen observable, and feature-estimation window. The main experiments used $T_s = 2\,\mathrm{s}$ configuration with an additional $T_s = 1\,\mathrm{s}$ configuration for the AES-128 payload. These durations are conservative proof-of-concept parameters that have not been claimed to be optimal.

4.4. Measured Quantities and Feature Extraction

The electrical waveforms recorded during the experiments were processed to obtain power quality indicators describing how the power supply load changes interacted with the electrical grid. These indicators were derived from synchronized voltage and current measurements in parallel and series configurations at the power supply input and office entrance.

The analysis aimed to determine whether changes in computational workload produced observable variations in external electrical parameters, focusing on two effects: variations in the

power factor and harmonic content of the input current.

The extracted metrics are the time series of the power factor and harmonic spectra from the Fourier analysis of the measurements. These were then used to assess whether controlled workload modulation could create detectable patterns suitable for the proposed covert channel.

4.4.1 Power and Power Factor Metrics

The power factor is the ratio of active power to apparent power in electrical networks and indicates how efficiently electrical energy is converted into useful work.

For measured voltage and current waveforms, $v(t)$ and $i(t)$, the active power P is

$$P = \frac{1}{T} \int_0^T v(t) i(t)\, dt,$$

where T denotes the observation interval. The apparent power S is

$$S = V_{\text{rms}} I_{\text{rms}},$$

and the power factor is

$$PF = \frac{P}{S},$$

With nonlinear loads, such as switched-mode power supplies, the current is typically non-sinusoidal; therefore, the measured power factor reflects both the phase displacement between the voltage and current and waveform distortion from harmonics.

In this study, the power factor is a key observable for detecting workload modulation. Changes in CPU utilization alter the instantaneous power draw of the system, affecting both the amplitude and shape of the input current. These changes modify the measured power factor and can create temporal patterns that reveal the underlying computational behavior.

4.4.2 Relative and Absolute Harmonic Metrics

With respect to harmonics, the current harmonic content demonstrates key electrical characteristics. This is owing to the nonlinearity of the current drawing, which contains harmonics at integer multiples of the main frequency.

To analyze these components, the current waveform was processed by a discrete Fourier transform, producing a spectrum that gives for each harmonic an amplitude relative to the fundamental.

For visualization, these amplitudes were normalized to the fundamental and expressed in decibels relative to the carrier (dBc),

$$H_{dBc} = 20 \log_{10} \left(\frac{H_n}{H_1} \right),$$

where H_n is the amplitude of the n-th harmonic and H_1 is the amplitude of the fundamental. These amplitudes were stored over time and in time-indexed matrices, on which they formed the basis of harmonic heatmaps.

These heatmaps reveal whether the modulation induces a correlated change across the harmonic n-order. A vertically aligned pattern in the harmonic content relative to time and n-order indicates patterns linked to deliberate workload modulation.

4.4.3 Windowing and Time Resolution

To limit gaps in acquisition due to transfer times of these frames into the measurement computer, the oscilloscope was configured to capture during the payload transmission in a single frame. This frame contained seven million points containing voltage, current, and waveform readings.

Temporal segmentation does not occur during acquisition; it only occurs during feature extraction. The hop Δ sets the time resolution of the features, whereas the window length T_w mainly controls the statistical reliability of each estimate: smaller hops yield denser time series, and longer windows improve averaging but smear temporal variations. To better reflect the experimental pipeline and replace the previous slot-based description, the method is best described as frame-based waveform capture followed by windowed post-processing.

This study looked at each harmonic's absolute RMS amplitude, in addition to providing harmonics in decibels relative to the fundamental. Because the dBc measure normalizes for each harmonic to the fundamental, this can be used to compare the spectral content independently of the overall current. However, this normalization can obscure certain modulation effects when the load scales the fundamental and higher-order components concurrently, and the ratio $\frac{H_n}{H_1}$ may remain almost constant despite the changes in absolute values. Consequently, the observables for symbol detection were the absolute harmonic RMS values. While harmonic RMS

represents absolute signal strength and can occasionally facilitate symbol separability, dBc is mostly employed in this dissertation to describe the relative spectral form.

4.4.4 Measurement Window and Symbol Duration

The measurement window and symbol duration play different roles. Window T_w specifies the interval over which each feature is estimated from the raw waveforms, whereas symbol duration T_s is the period for which the transmitter holds one of the two workload states used in OOK modulation.

These timescales are linked but distinct. If T_w is too long relative to T_s, symbol transitions are smoothed, and the two workload states become harder to distinguish. If T_w is too short, the power factor and harmonic estimates become more sensitive to noise and short transients. In addition, the input energy storage components and control loops of an SMPS prevent its grid-side electrical response from behaving as an instantaneous copy of the CPU command.

The principal value $T_s = 2\,\mathrm{s}$ was selected as the conservative engineering choice for this exploratory evaluation. On a 50 Hz supply, it spans approximately 100 mains cycles, allowing each workload state to persist over multiple electrical cycles and feature-estimation intervals. It was also compatible with the oscilloscope memory constraint that led to the 10 kSa/s, 700 s acquisition configuration. The additional $T_s = 1\,\mathrm{s}$ AES-128 test spans approximately 50 mains cycles and was included to assess whether a shorter, less conservative interval could still yield a recoverable frame in selected configurations.

No systematic sweep over the symbol duration was performed. Therefore, neither 1 s nor 2 s is presented as the optimal or minimum viable duration. Establishing the maximum symbol rate requires dedicated experiments over multiple T_s values and repeated independent transmissions.

At the decoding stage, the receiver does not rely on the exact prior knowledge of the recovered symbol timing. It processes the exported feature series, resamples it onto a uniform grid, and searches for candidate symbol durations and timing offsets before threshold detection and preamble matching.

4.4.5 Measurement and Estimation Uncertainty

The reported values are affected by acquisition, feature estimation, and decoding uncertainty. The oscilloscope has finite sampling resolution, memory, and transfer capability. The power

factor, THD, and harmonic amplitudes were not treated as direct instrument-grade readings with fully propagated uncertainty; they were estimated from finite waveform windows during offline processing. Their values therefore depend on the selected window length, hop size, resampling procedure, and spectral estimation method.

The decoding stage introduces additional sensitivity through moving-average smoothing, candidate symbol duration and timing-offset searches, threshold estimation, and preamble alignment. The BER reported for a feature is the error fraction observed after applying this processing chain to a recorded transmission.

Run-to-run variability is a separate source of uncertainty that has not yet been measured. Because each experimental configuration was acquired once, this study did not quantify the effect of repeated executions, changes in background loads, temperature, grid conditions, or acquisition timing. Statistical dispersion across runs and confidence intervals was, therefore, not reported. The results are descriptive proof-of-concept observations of the evaluated captures.

4.5. Baseline Characterization

We have explained the electrical behavior before looking into the purpose of modulation-based workload. The baseline can serve as a later reference condition.

In reality, loads connected to the electrical wiring change the electrical signals. Either by background loads, normal grid fluctuations, measurement noise, or interference by other devices nearby. A system that operates in idle or steady conditions without any deliberate workload changes can serve to set the baseline. Voltage and current waveform techniques for feature extraction are described in Section 4.4. For every power supply in this study, baseline data were gathered independently.

Idle Operation - No PFC PSU

Figure 4.4 shows the harmonic heatmap obtained when the computing platform was powered by an SMPS without power factor correction.

The low-order harmonics produced by the power supply without any power factor correction circuit significantly control the spectrum. This harmonic pattern does not change much over time, indicating that while the system is not in use, no structured temporal structure appears.

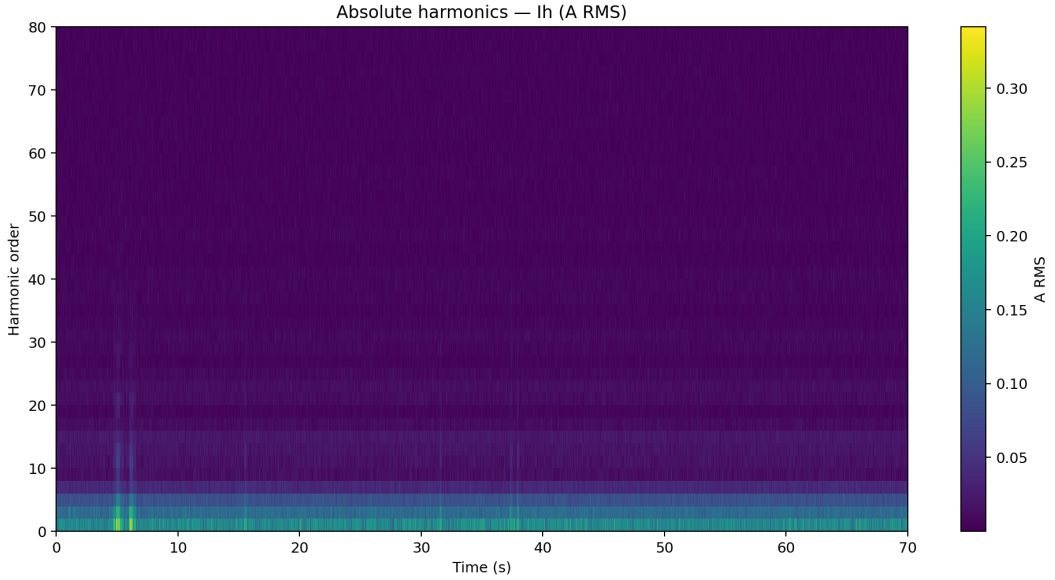

Figure 4.4: Harmonic heatmap of the input current during idle operation using an SMPS without power factor correction.

Idle Operation - Passive PFC PSU

The same experiment was repeated using a power supply with passive power factor correction. Figure 4.5 demonstrates the baseline for the passive PFC.

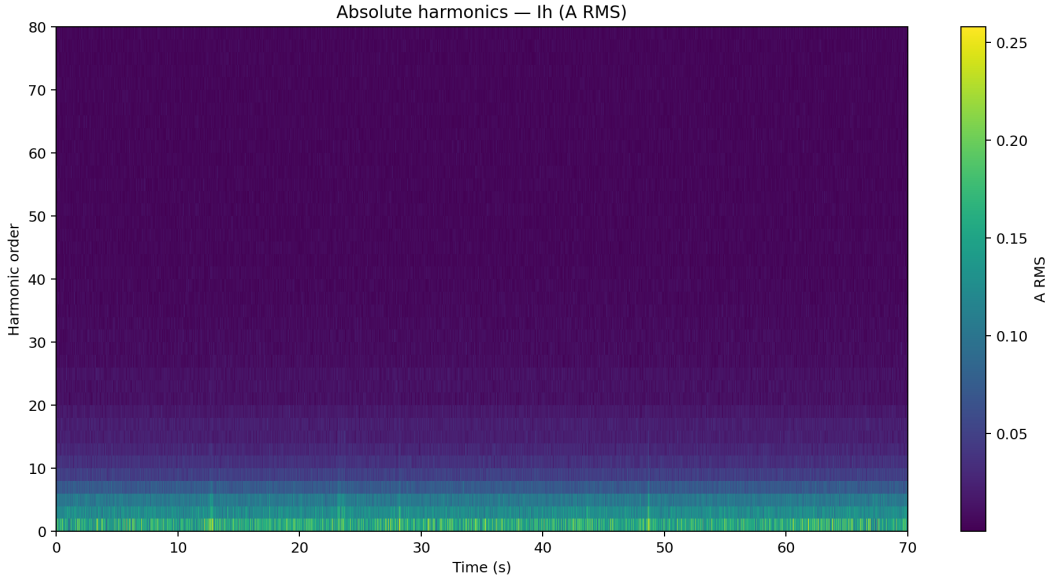

Figure 4.5: Harmonic heatmap during idle operation using a passive PFC PSU.

Compared with the earlier power supply, the harmonic amplitudes were lower, and the spectrum was more smoothly distributed. However, the harmonic content remained consistent over time, indicating the steady-state electrical behavior of the system.

Idle Operation - Active PFC PSU

Finally, the experiment was repeated using a power supply with an active power factor correction, as demonstrated in Figure 4.6.

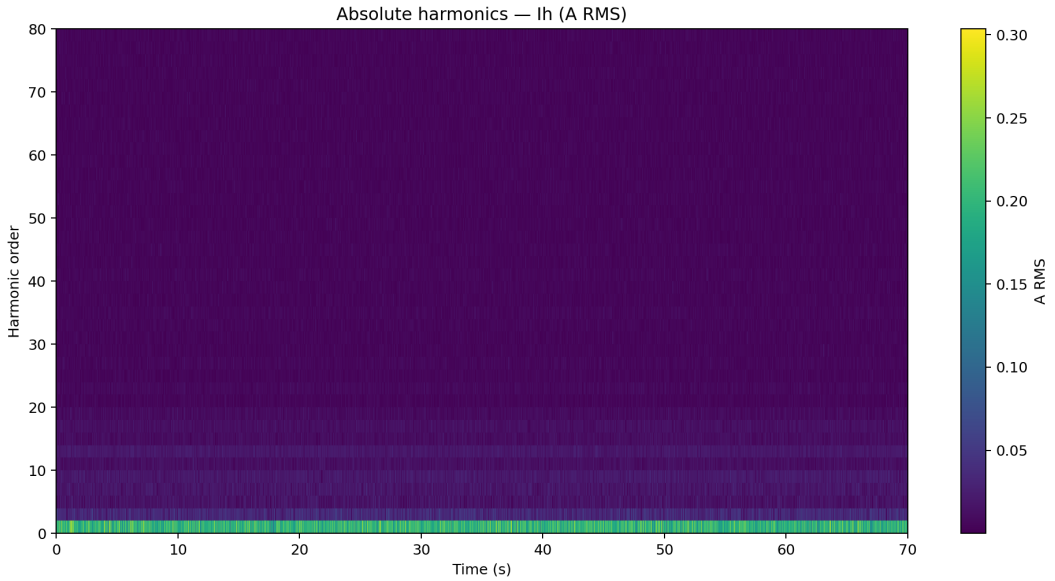

Figure 4.6: Harmonic heatmap during idle operation using an active PFC PSU.

As anticipated, the harmonic content of this circuit was significantly lower because of the current shaping provided by the active PFC circuit. The input current waveform became more sinusoidal, producing reduced harmonic amplitudes.

Baseline Comparison

The measurements revealed differences among the three power supplies. The relevant one, without the PFC, creates the most harmonic noise. In contrast, the active PFC power supply was the cleanest in harmonic content.

Despite these differences, each power supply maintained a stable harmonic content during idling. This serves as a baseline and comparison with further experiments with deliberate induced power drawn from these power supplies.

4.6. Load Sweep and Scaling Behaviour

Using the baseline characterization from the previous section, a series of controlled experiments were executed to determine how the load changes relate to observable electrical properties. The first purpose of this experiment was to see if changes in CPU usage caused changes in the power quality metrics.

To do this, a sweep load was applied, which slowly increased the CPU load over time from idle to fully used. As the system went through different load levels during the experiment, the voltage and current waveforms were constantly recorded. We used the feature extraction methods explained in Section 4.4 to convert these data into a time series of electrical indicators.

CPU Load Sweep Experiment

The load sweep test involved slowly adding more load to the CPU until all processor cores were fully used. This method causes the power consumption of the system to rise slowly, which in turn changes the measurable electrical properties.

The electrical measurements had to stabilize before moving on to the next workload level, so each load step was held for approximately 1 s. The entire sweep from the idle load to the peak load required approximately 30 s.

Therefore, the dataset we collected provides a continuous picture of how electrical parameters change as computational needs grow.

Power Factor Response

Figure 4.7 shows the measured power factor during the CPU load sweep experiment.

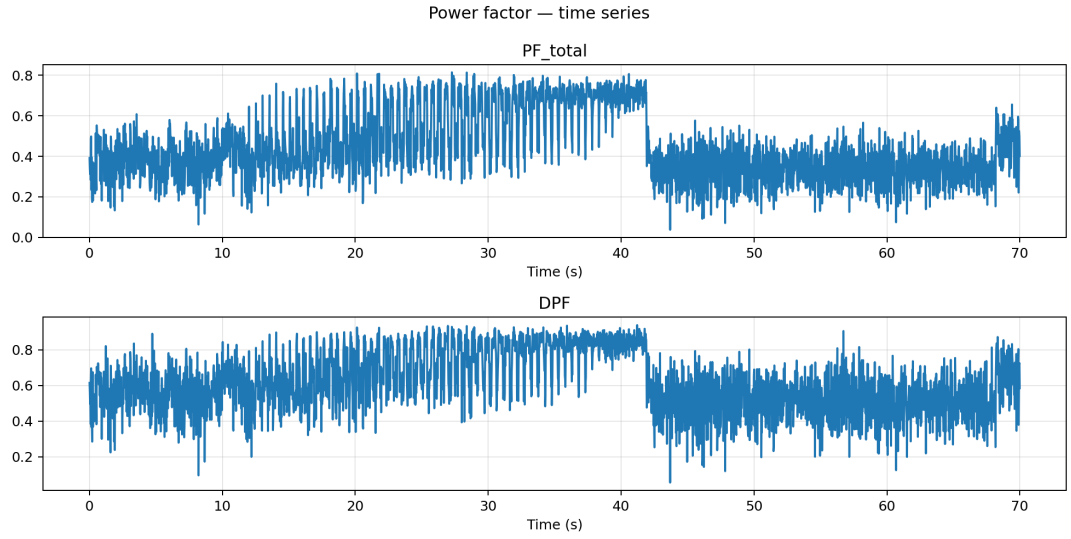

Figure 4.7: Power factor measured during the CPU load sweep experiment (Active PFC PSU).

The electrical behavior changed noticeably as the load increased. Changes in the CPU usage used for this experiment caused changes in the measured power factor. These changes were caused by changes in the waveform of the input current drawn by the power supply. Because

switched-mode power supplies draw current in a nonlinear manner, especially those without PFC, changes in the power demand change both the shape and size of the current waveform.

Harmonic Evolution

In addition to changes in the power factor, the harmonic composition of the input current also changes as the system load increases. Figures 4.8, 4.9 and 4.10 show the evolution of the harmonic spectrum over time during the load sweep as a function of the PFC design.

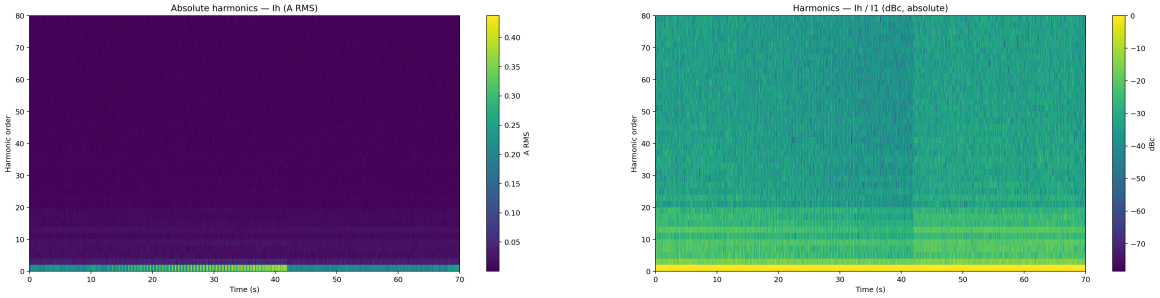

(a) Absolute in A_{rms}.

(b) Relative in dBc.

Figure 4.8: Absolute and relative harmonic content heatmap of the input current during the CPU load sweep experiment (Active PFC PSU).

Regarding the Active PFC, Figure 4.8 shows that the load sweep introduces very few harmonics into the spectra, as expected. The active PFC circuit attempts to follow the voltage waveform and reduce the harmonic distortion.

The heatmap in Figure 4.9a shows that the amplitudes of the harmonic components change over time for the passive power-factor correction circuit. The fact that a passive PFC does not track the voltage waveform indicates that harmonic distortion is created owing to nonlinearity.

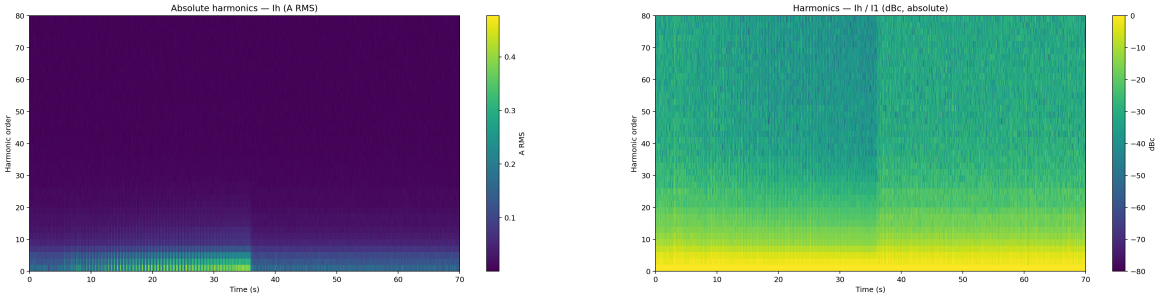

(a) Absolute in A_{rms}.

(b) Relative in dBc.

Figure 4.9: Absolute and relative harmonic content heatmap of the input current during the CPU load sweep experiment (Passive PFC PSU).

These results reveal that the load directly reshapes the harmonic content of the input-current waveform. This supports the idea that a controlled load can produce measurable patterns in the spectra.

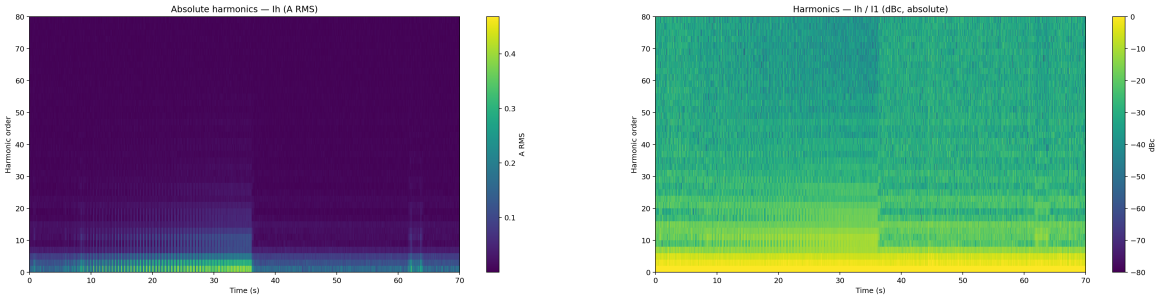

(a) Absolute in A_{rms}.

(b) Relative in dBc.

Figure 4.10: Absolute and relative harmonic content heatmap of the input current during the CPU load sweep experiment (Non-PFC PSU).

A comparison between Figure 4.10a and 4.8b shows that the power supply unit with active power factor correction concentrates most of the harmonic content in the fundamental current component, meaning at the bottom of the heatmap. Without a PFC circuit, Figure 4.10a clearly shows that it has a stronger signal than the other power supplies with a PFC.

This provides a physical foundation for the covert channel method described in the following sections.

4.7. Transient Load Modulation (Pulse Ramp)

In the previous section, we showed that changes in the load can affect the power factor and harmonic content. This is dependent on power consumption, which gives the RMS amplitude and power factor design, where no PFC means more harmonic content, whereas an active power factor means almost no harmonic content.

Load sweeps capture only gradual changes over time. In the next experiment, we tried a PWM-like load to mimic the encoded data. The load duration changed while the idle time also changed, where the pulse window duration was gradually increased to 100% duty.

Pulse Ramp Experiment

The pulse ramp experiment is characterized by moments of full load and idle usage. In this setup, we caused sudden changes in the computational activity, unlike the gradual load sweep that increases from 0% usage to 100%.

From this, a Python script was created to produce 0% to 100% load pulses by quickly switching between low and high CPU usage for a specific time. Each load pulse was a short spike of the full CPU load, followed by a return to the idle state.

The pulse durations were increased incrementally during the experiment, starting at approximately 1 ms and increasing up to 500 ms. This ramp allows us to examine how the system reacts over a range of timescales.

The total time for the experiment was approximately 30 s, during which the measurement setup described in Section 4.1 recorded the voltage and current waveforms continuously.

Power Factor Temporal Response

Figure 4.11 shows the power factor measured during the pulse-ramp experiment.

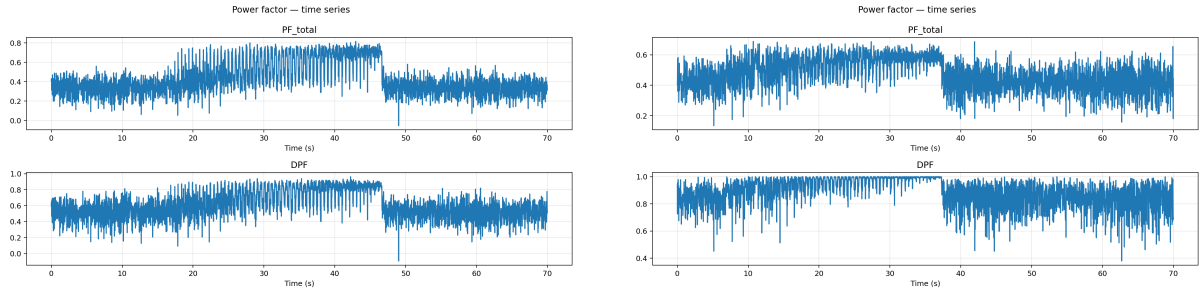

(a) Active PFC PSU (b) Non-PFC PSU

Figure 4.11: Power factor evolution during transient CPU load modulation.

Along with the sweep load results, the same effects are visually clear for the pulsed load experiments. The time series in Figure 4.11 shows that the power factor increases over time, as expected. This indicates that sudden increases in load lead to corresponding fluctuations in the power factor.

The envelope of the measured power factor increased as the pulse duty cycle increased, although the individual estimates remained highly variable.

Harmonic Response

Figure 4.12 presents the harmonic heatmap of the input current during the pulse-ramp experiment.

Compared with the baseline, the harmonic spectrum showed a regular time pattern. Sharp vertical bands were particularly visible in the non-PFC panels in several harmonic orders, aligned with the transitions between idle and load changes.

Thus, changes in workload cause shifts in multiple harmonic orders. Stronger signals from low-order harmonic to fading signal to higher-order harmonic. In effect, load modulation produces

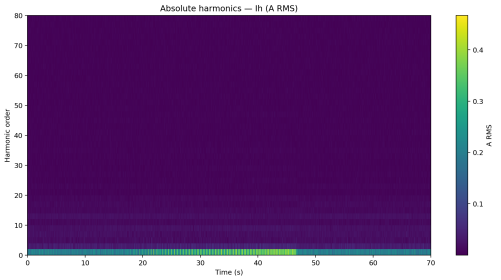

(a) Active PFC PSU in A_{rms}.

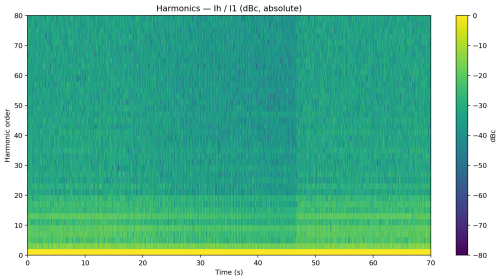

(b) Active PFC PSU in dBc.

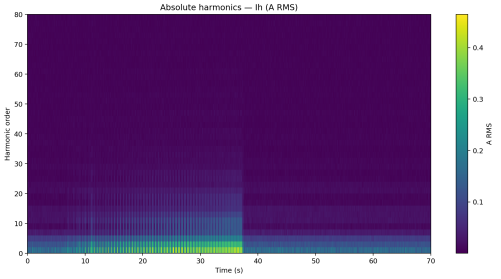

(c) Non-PFC PSU in A_{rms}.

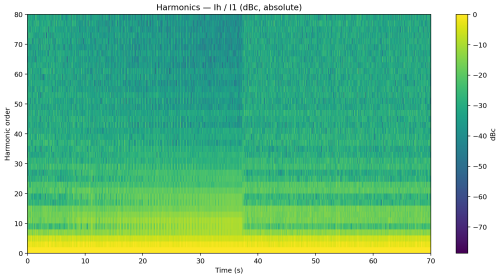

(d) Non-PFC PSU in dBc.

Figure 4.12: Harmonic response in absolute and relative.

clear spectral signatures, making encoding data into this harmonic and power factor variation feasible.

Implications for Covert Signaling

Previous experiments demonstrated that changing the power supply load and drawing more or less current from the mains, especially in a nonlinear way, can create signaling into the harmonic distortion and power factor metrics. These signal levels are measurable and can therefore be used to create a one-way channel for exfiltrating information from an air-gapped system.

Because we can only control the presence or absence of the signal, a modulation based on the signal amplitude must be used. Therefore, in the next section, we discuss on–off keying modulation and how we use it to encode digital symbols and later decode them.

4.8. On-Off Keying

On-Off Keying is a basic modulation scheme in which a signal is either present or absent, and this binary state is used to encode information.

The experiments used a transmission structure consisting of a preamble and payload (containing a 128-bit AES key, an ECDSA secret key, or UTF-8/ASCII data), followed by four parity bytes generated by a Reed-Solomon error-correcting code.

Each symbol had a duration of 2 s, with the exception of an additional test case for the AES example key, where a symbol duration of 1 s was used.

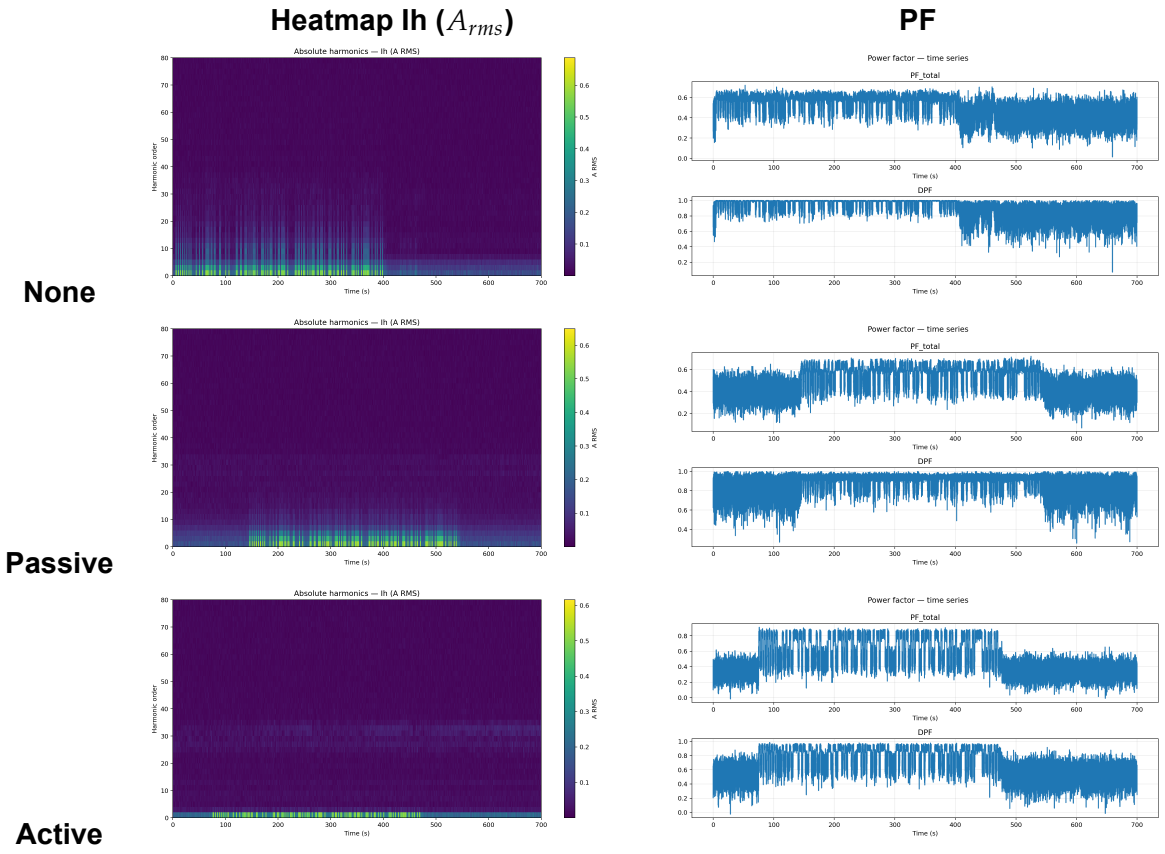

Figure 4.13: Heatmap and timeseries for each PSU design measured directly to the PSU.

Figure 4.13 shows the workload-correlated temporal bands, most clearly in the no-PFC and passive PFC measurements. The power factor series also exhibits the two workload states used for OOK.

4.8.1 Symbol Detection Model

Symbol detection simultaneously uses one electrical feature. Depending on the experiment, this feature was either a single relative harmonic amplitude (e.g., the 3rd, 5th, 7th, or 11th current harmonic in dBc), absolute in I_n or another scalar quantity such as the power factor, apparent power, or active power. A Python-written decoder was used, which processed an exported CSV time series and treated each feature separately.

Let $z(t)$ be the chosen feature over time. The decoder first cleans the series by removing invalid samples, sorting them by time, and merging duplicate timestamps if they are present. The data were then resampled onto a uniform time grid via linear interpolation, and a short moving-average filter was applied to suppress local fluctuations prior to symbol segmentation.

Because the exact symbol timing is unknown, the decoder scans a range of candidate symbol durations around the nominal duration T_s and a set of candidate timing offsets within each symbol period. For each duration–offset pair, the resampled sequence was divided into consecutive symbol intervals, and each symbol was represented by the mean of the samples in its corresponding interval. We denote this mean for symbol m by $\bar{z}_m$, yielding a one-dimensional sequence of symbol mean values.

A binary decision is obtained by thresholding the symbol mean. If no manual threshold is specified, it can be inferred from the empirical distribution of the symbol mean. Specifically, letting $q_{0.2}$ and $q_{0.8}$ denote the 20th and 80th percentiles of the sequence of symbol means, the threshold is set to

$$\tau = \frac{q_{0.2} + q_{0.8}}{2}.$$

Each symbol mean is then mapped to a bit via

$$\hat{b}_m = \begin{cases} 1, & \bar{z}_m \geq \tau \\ 0, & \bar{z}_m < \tau \end{cases}$$

with an optional inversion to account for feature polarity.

For every tested combination of symbol duration and offset, the resulting bitstream was searched for the known preamble. The decoder compares the preamble with all bit-aligned positions in the stream and assigns a score equal to the number of matching bits. The selected decoding is the one with the highest preamble score; ties are broken first by choosing the symbol duration closest to the nominal value and then by the earliest preamble position.

Once the best preamble alignment is determined, subsequent bits are parsed as packets. By default, the first byte encodes the payload size, followed by the payload data and a Reed–Solomon redundancy field, yielding a packet of the form `| preamble | size | data | reed-solomon |`. The decoder can also verify the Reed–Solomon code, and for this work alone, the expected payload is provided, and the bit error rate is computed.

Overall, the detection scheme uses single-feature threshold decoding with a joint search over symbol duration and offset, followed by preamble synchronization and packet parsing. It was used to assess the communication performance of each observable individually under different PSU types, payloads, and receiver positions considered in this study.

4.9. Receiver Placement and Electrical-Path Effects

In this section, the results are shown in a heatmap for harmonic features in dBc, Irms, and a time series for the true power factor, or ”PF_total” as shown in the figures, and displacement power factor. The DPF is determined from the phase angle between the current and voltage waveforms. In contrast, the true power factor accounts for the DPF and overall harmonic distortion that contribute to the power factor.

The tests were conducted as previously stated: an ASCII string, AES-128 example key, and ECDSA secret. The measurements were performed at two receiver placements: near the PSU and at the office-level measurement point, where the observed current included the target contribution mixed with other loads connected to the same electrical infrastructure.

All decoding statements in this section refer to the outcomes observed in the single recorded run for each configuration. A “successful decode” means that the payload was recovered and Reed–Solomon verification was valid for that feature/configuration combination; it does not imply a measured success probability across repeated transmissions.

For each payload, the comparison figure is arranged in two rows. The upper row contains the near-PSU relative harmonics, power-factor time series, and absolute harmonics, and the lower row contains the corresponding office-level measurements. This compact arrangement keeps the receiver locations directly comparable while leaving space for interpretation in the surrounding text.

4.9.1 Meaning of Distance in a Conducted Power-Line Channel

In this work, distance does not refer to the free-space separation between the transmitter and receiver. The physical layout used for the office-level measurements is shown in Figure 4.3. The geometric target of the OSC distance shown in the figure is an approximate physical separation, whereas the relevant distance for the conducted channel is the receiver placement along the shared electrical infrastructure. The proposed channel is not radiated through air, as in acoustic, optical, or electromagnetic covert channels. Instead, the signal is a conducted perturbation of the line current caused by workload-dependent changes in the target PSU input current.

Consequently, the relevant notion of distance is receiver placement along the shared electrical infrastructure. A receiver placed near the PSU observes a signal with minimal contribution from unrelated loads. A receiver placed upstream, such as at the office-level electrical distribution

point, observes an aggregate current composed of the target contribution and the contribution of other devices connected to the same electrical network. Therefore, the degradation is not expected to be purely proportional to the physical cable length. It depends on the electrical topology, conductor impedance, filtering, branching, and background load environment.

The measured signal at a receiver can be interpreted as

$$i_{\text{obs}}(t) = h_{\text{path}}(t) * i_{\text{target}}(t) + \sum_{k=1}^{K} i_{\text{load},k}(t) + \eta(t),$$

where $i_{\text{target}}(t)$ is the workload-modulated current drawn by the target system, $h_{\text{path}}(t)$ represents the transfer effect of the electrical path between the target and the receiver, $\sum_{k=1}^{K} i_{\text{load},k}(t)$ represents other loads connected to the same infrastructure, and $\eta(t)$ represents the measurement noise and external disturbances.

Thus, the experiments in this section did not attempt to establish a maximum metric range in meters. Instead, they evaluated whether the covert signal remained observable when the receiver was moved from a near-source measurement point to a more realistic upstream measurement point in the building wiring.

4.9.2 No PFC Results

A power supply without any power factor correction mechanisms produced the strongest harmonic signal. This is in line with expectations because it is not a linear load and has a worse harmonic performance and a larger variation in the total power factor, which accounts for phase delay and harmonic distortion. Figures 4.14a and 4.14d show that harmonic patterns are preserved from different measurement locations with some attenuation for the office location, figures 4.16a and 4.16d, owing to other loads connected to the office electrical grid.

The recorded non-PFC transmissions produced the largest set of successful feature decodes among the evaluated PSU categories. For the ASCII payload, both apparent and active power decoded correctly at the PSU and office locations, as shown in Table A.1, whereas the 11th harmonic decoded successfully only at the office location. In the AES-128 transmission (Table A.2), the recorded office capture yielded successful decodes through the 5th and 7th harmonics and through the apparent and active powers; several of these feature outcomes were better than those obtained from the corresponding near-PSU capture. For the ECDSA payload (Table A.3), the apparent and active power decoded correctly at both locations, whereas the 7th harmonic

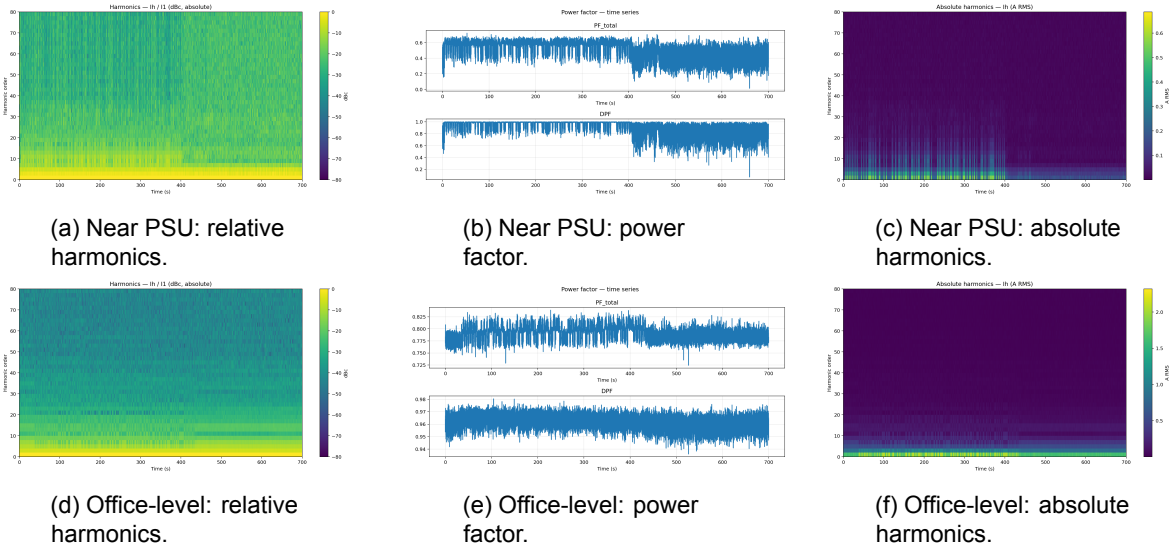

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.14: Comparison of measurement at PSU vs office measurements for the no-PFC PSU - ASCII

succeeded only at the PSU and the 11th harmonic only at the office. These observations identify the non-PFC setup as the strongest case in the recorded proof-of-concept dataset without establishing its expected success rate across repeated runs.

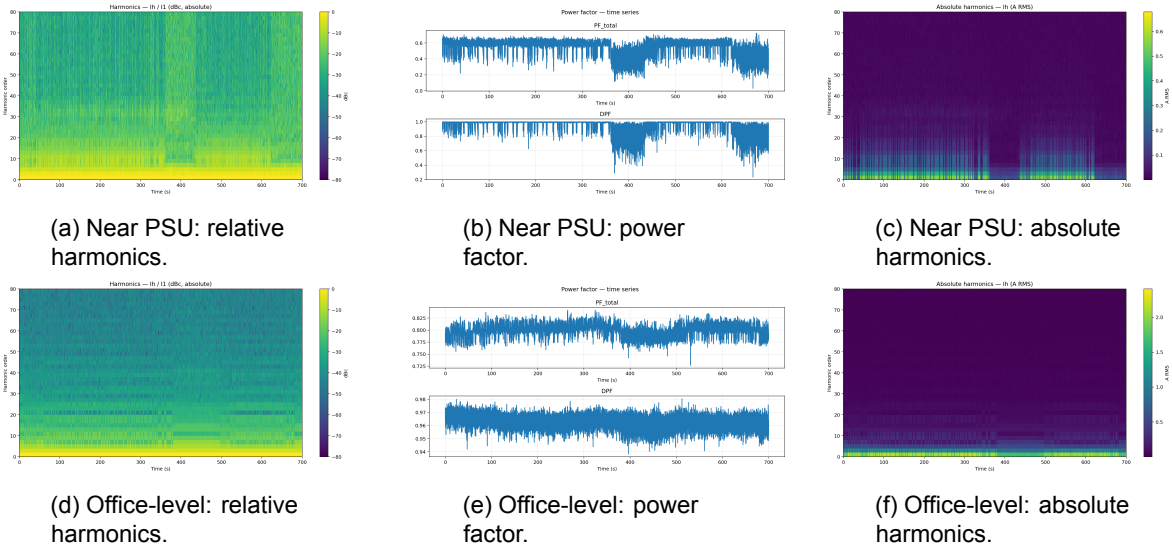

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.15: Comparison of measurements at the PSU and office locations for the no-PFC PSU - AES key

A key finding was that office measurements were not always worse than those at the PSU. In several instances, the office demonstrated a lower BER and more successful decoding, particularly for the AES-128 payload. This indicates that, with a strongly nonlinear load, the propagation path does more than just attenuates the signal. Depending on the chosen features and the electrical environment, the signal can remain distinguishable or even appear cleaner at a

more distant point. Thus, the non-PFC PSU clearly demonstrates that power line signaling can persist beyond the immediate source connection.

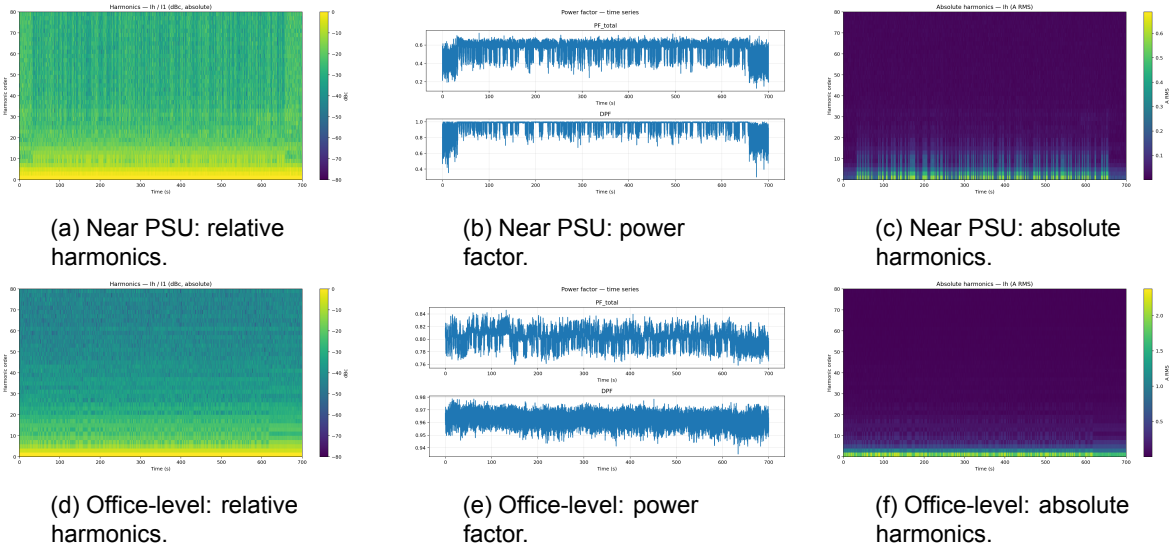

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.16: Comparison of measurements at the PSU and office locations for the no-PFC PSU - ECDSA

4.9.3 Passive PFC Results

The passive PFC PSU represents a middle ground between the strongly distorted non-PFC configuration and the tightly controlled active PFC setup. Because passive correction mitigates but does not remove distortion, the covert-channel features remain detectable, though typically with lower contrast and reduced consistency across payloads and receiver locations, as illustrated in figures 4.17d, 4.19d, and 4.18d. The communication results in Tables A.4, A.5, and A.6 indicate that some harmonic observables still convey useful information, but the channel is more selective and less uniformly dependable than in the non-PFC scenario.

For the ASCII payload, as shown in Table A.4, the 5th harmonic and active power decoded successfully in the near-PSU capture, whereas the 11th harmonic, apparent power, and active power decoded successfully in the office-level capture.

For the AES-128 payload, Table A.5 reports several successful feature decodes in the near-PSU capture, including the 7th dBc harmonic at $T_s = 1\,\mathrm{s}$, several RMS harmonic amplitudes, and the power factor at both symbol durations. The 5th dBc harmonic at $T_s = 2\,\mathrm{s}$ is marked as RS-invalid and is therefore not counted as a successful decode with valid error-control verification. For the ECDSA payload, Table A.6 shows that the 11th relative harmonic, 3rd through 11th RMS harmonics, and active power decoded correctly in the near-PSU capture, whereas only the 5th RMS harmonic decoded successfully in the office-level capture.

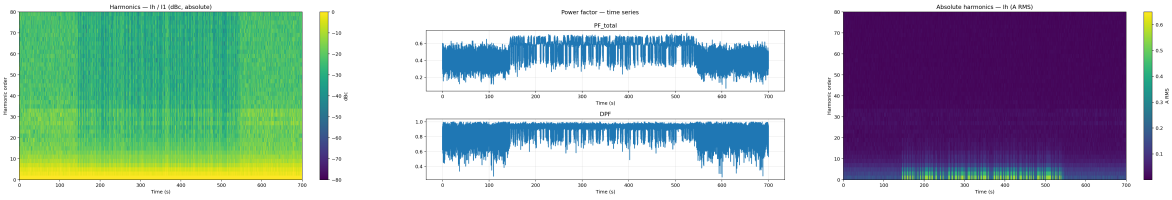

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

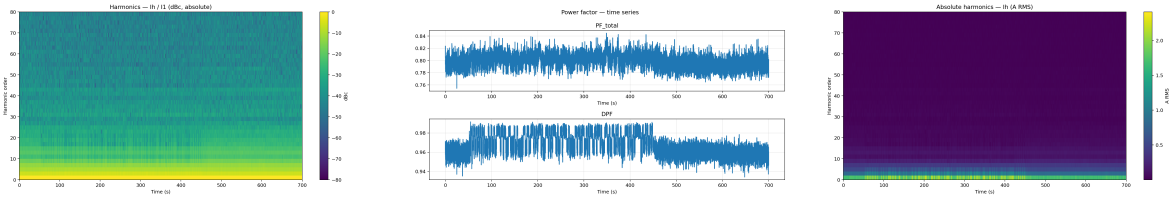

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.17: Comparison of Passive PFC PSU measurements at PSU and office locations - ASCII

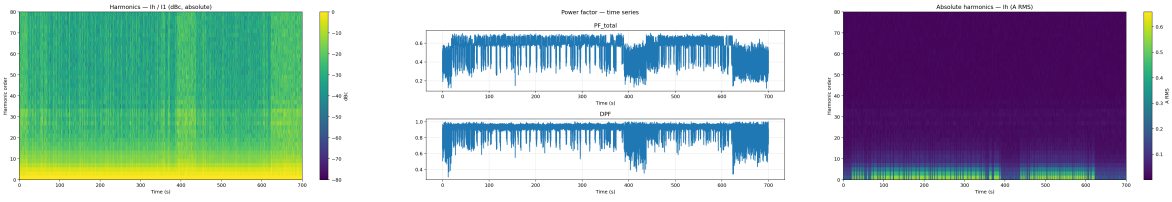

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

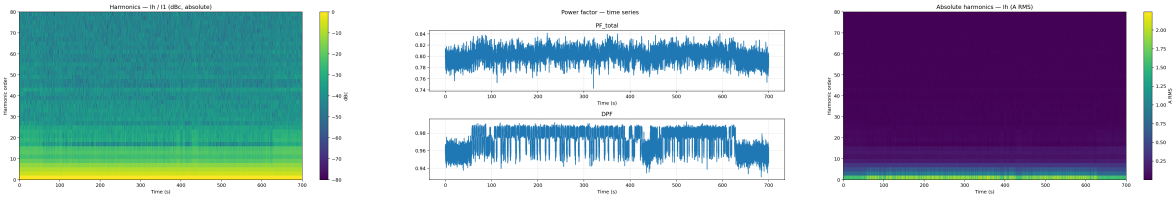

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.18: Comparison of Passive PFC PSU measurements at PSU and office locations - AES key

One particular interesting finding regarding passive PFC is that even though the signal carrier in harmonics is not strong at the office location, the DPF feature is mostly clear, as shown in figures 4.17e, 4.18e and 4.19e. Taken together, the recorded passive PFC transmissions demonstrate that covert signaling remained possible in the selected feature/configuration combinations, but fewer features succeeded at the office-level point than in the non-PFC case. The active power, apparent power, and selected harmonic components produce successful decodes depending on the payload and receiver position. Because each configuration was acquired once, these observations describe feature selectivity in the evaluated captures rather

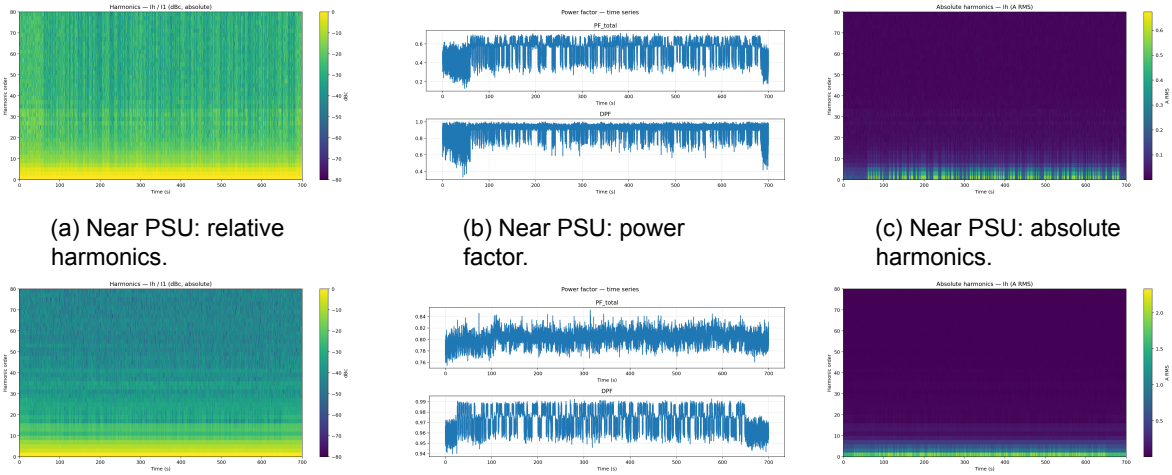

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.19: Comparison of Passive PFC PSU measurements at PSU and office locations - ECDSA

than run-to-run robustness or stability.

4.9.4 Active PFC Results

The recorded active-PFC measurements showed the largest reduction in harmonic observability between the near-PSU and office-level positions, as illustrated in Figures 4.20a, 4.20d, 4.21a, 4.21d, 4.22a, and 4.22d. This behavior is consistent with the intended electrical function of the active PFC: shaping the input current towards the voltage waveform reduces low-order harmonic distortion and suppresses several of the spectral features used by the proposed receiver. Tables A.7, A.8, and A.9 show successful decoding for a smaller set of feature/configuration combinations than for the other PSU classes in the recorded dataset.

For the ASCII payload, the active power was the only feature with a valid successful decode in near-PSU capture. At the office-level location, active power produced a valid decode, whereas the apparent-power candidate failed the Reed-Solomon verification. For the AES-128 payload, Table A.8 reports successful near-PSU decodes through the 5th and 11th dBc harmonics at $T_s = 1\,\mathrm{s}$, the 3rd RMS harmonic for both symbol durations, the power factor for both symbol durations, and power-based features. The 5th harmonic result at $T_s = 2\,\mathrm{s}$ did not produce a valid successful decode. At the office-level receiver, the performance degrades sharply; only the 3rd dBc harmonic at $T_s = 2\,\mathrm{s}$, power factor at $T_s = 2\,\mathrm{s}$, and active power at $T_s = 2\,\mathrm{s}$ remain successful, with most harmonic features failing.. For the ECDSA payload, the constraints tighten further: at the PSU, the 5th harmonic, power factor, apparent power, and active power work;

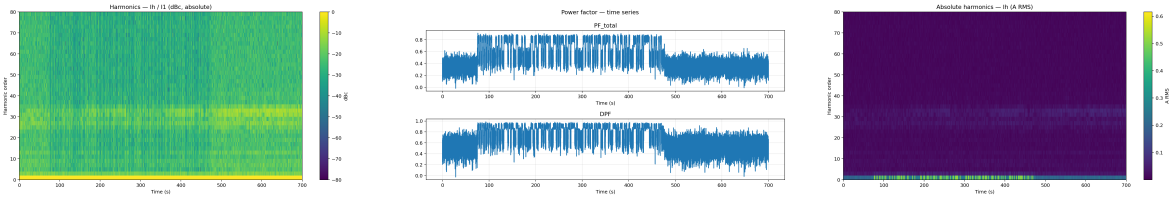

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

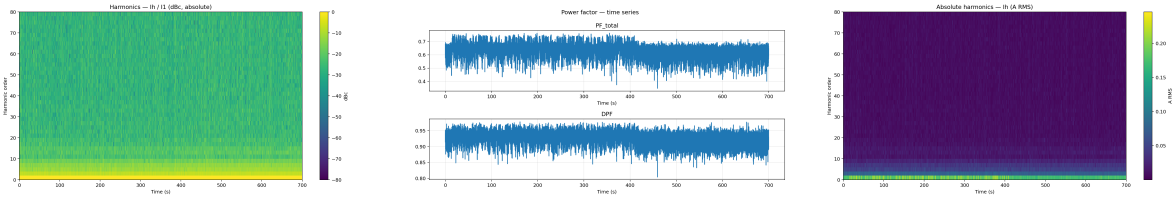

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.20: Comparison of Active PFC PSU measurements at PSU and office locations - ASCII

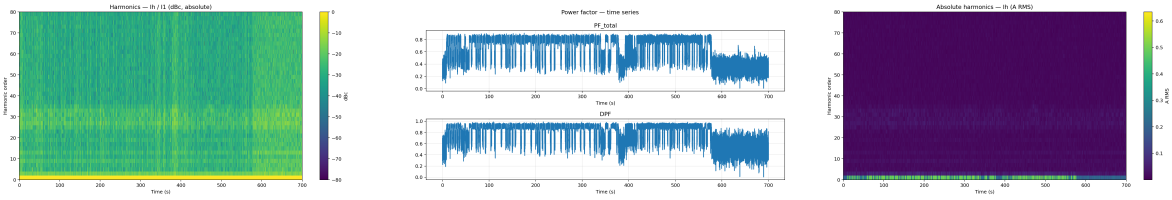

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

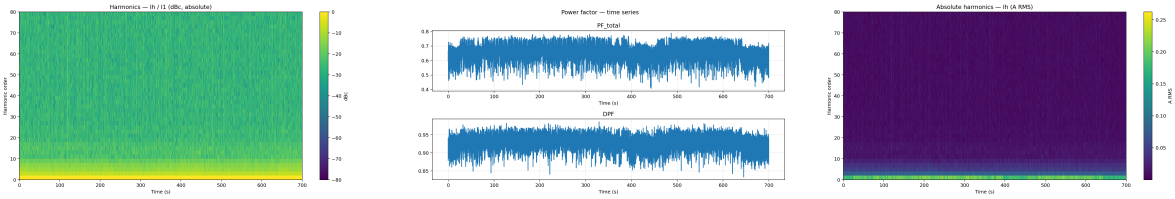

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.21: Comparison of Active PFC PSU measurements at PSU and office locations - AES key

however, at the office, only the apparent power remains successful, whereas all harmonics and the power factor fail.

Within the recorded proof-of-concept dataset, the active PFC was associated with the weakest harmonic features and smallest number of successful harmonic decodes at the office-level receiver. The successful outcomes were concentrated mainly in power-domain metrics, particularly active power, and apparent power in some captures. These observations are consistent with current waveform shaping suppressing the proposed harmonic signaling mechanism; however, the single-run design does not establish a general success probability or statistical

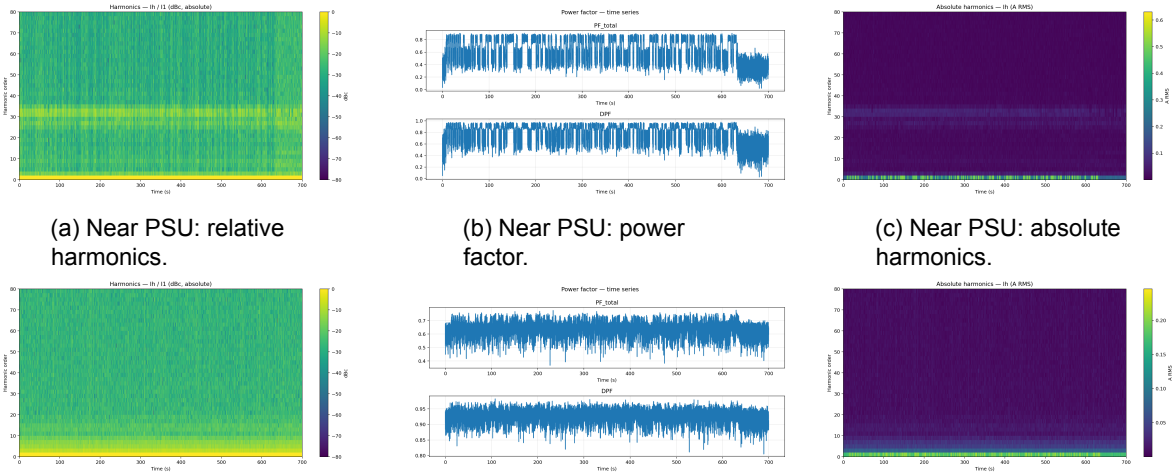

(a) Near PSU: relative harmonics.

(b) Near PSU: power factor.

(c) Near PSU: absolute harmonics.

(d) Office-level: relative harmonics.

(e) Office-level: power factor.

(f) Office-level: absolute harmonics.

Figure 4.22: Comparison of Active PFC PSU measurements at PSU and office locations - ECDSA

difference between the PSU classes.

4.10. Summary of Receiver Placement and PSU Effects

Table 4.5 provides a descriptive summary of the observables central to this dissertation: the power factor and harmonic amplitudes. Active and apparent power are intentionally excluded, even when they produce successful decodes, because the purpose of this table is to evaluate PF and harmonic distortion as information-bearing features. The interpretation immediately following the table separates the physical observability comparison from the communication rate metrics reported in Table 4.4.

Table 4.5: Descriptive PF and harmonic decoding outcomes for the recorded proof-of-concept transmissions.

PSU	Receiver	RMS	dBc	PF	Harm. SNR	A/K/E	Best BER A/K/E
No PFC	Near PSU	9/16 20.6 dB	3/16 9.0 dB	0/4 10.7 dB	**14.8 dB**	N/Y/Y	– / 0 / 0
No PFC	Office-level	9/16 16.5 dB	6/16 15.3 dB	0/4 10.2 dB	**15.9 dB**	Y/Y/Y	0 / 0 / 0
Passive PFC	Near PSU	14/16 17.2 dB	3/16 9.2 dB	2/4 11.3 dB	13.2 dB	Y/Y/Y	0.0069 / 0 / 0
Passive PFC	Office-level	7/16 13.2 dB	5/16 9.6 dB	0/4 9.2 dB	11.4 dB	Y/Y/Y	0 / 0 / 0
Active PFC	Near PSU	2/16 9.1 dB	4/16 8.0 dB	3/4 13.9 dB	8.5 dB	N/Y/Y	– / 0 / 0
Active PFC	Office-level	0/16 8.1 dB	1/16 8.2 dB	1/4 8.1 dB	8.1 dB	N/Y/N	– / 0 / –

A/K/E denotes ASCII, AES-128 key, and ECDSA key payloads, respectively.
RMS and dBc cells report successful feature/configuration decodes with valid Reed–Solomon verification, followed by the arithmetic mean SNR for that feature family in the recorded dataset.
PF cells report successful PF feature/configuration decodes with valid Reed–Solomon verification, followed by the arithmetic mean PF SNR.
Harm. SNR is the arithmetic mean across the evaluated RMS and dBc harmonic features only; PF is not included.
RS-invalid entries are excluded from the successful-decode counts.
The counts summarize alternative features and decoding conditions applied to the recorded captures; they do not represent repeated independent trials.
RMS and dBc denominators are 16: four harmonic orders evaluated over four decoding conditions. The PF denominator is 4: one PF feature is evaluated over the same four decoding conditions.

In the recorded datasets, the no-PFC configuration produced the highest mean harmonic SNR for both receiver placements. At the office-level receiver, all three payloads were recovered using at least one harmonic feature. passive-PFC captures also produced successful decodes, particularly through absolute RMS harmonic amplitudes; however, the successful feature set depended more strongly on the payload and receiver placement. The active-PFC office-level capture produced the smallest successful feature count: one dBc-harmonic decode and one PF decode across the evaluated conditions.

The table also shows that the feature count and SNR should not be interpreted as independent measures of repeated-channel reliability. The fractions count feature/configuration combinations evaluated on the recorded captures, not successes from independent trials. Near the

source, the no-PFC configuration produced the highest harmonic SNR, with successful RMS-harmonic decodes but no PF decode, whereas the passive-PFC configuration produced the largest successful feature count, mainly through RMS harmonics. At the office-level point, the no-PFC row combines the strongest recorded harmonic SNR with the recovery of all three payloads; the passive-PFC results were more feature-selective but also recovered all three payloads. The active-PFC captures produced fewer harmonic decodes, with a greater contribution from the PF near the PSU, and the office-level row contained the smallest PF/harmonic feature count and recovery only of the AES-128 payload. These comparisons support the physical interpretation that active current shaping reduces the low-order harmonic observability under the tested conditions.

Using Equation 4.3 and the rates in Table 4.4, the only non-zero best valid BER listed in Table 4.5 is 0.0069 for the passive PFC near-PSU ASCII result. At 0.5 bit/s, this is approximately 0.00345 observed erroneous bits per second before Reed–Solomon correction. Best BER entries of zero correspond to zero observed errors in the compared interval, while “–” indicates that no valid PF/harmonic payload recovery was available for that payload in the summarized row. As in this section, these values describe the selected recorded decodes and do not estimate future error rates or packet-success probabilities.

5. Discussion

This chapter discusses the exploratory proof-of-concept results from Chapter 4 with respect to the research questions defined in Chapter 1. Rather than revisiting each measurement individually, it interprets the observed workload-induced power quality changes, end-to-end payload recovery, and the effect of PSU topology and receiver placement. Because each experimental configuration was acquired once, the discussion distinguishes the evidence of feasibility from the statistical repeatability or expected channel reliability.

5.1. Discussion with Respect to the Research Questions

5.1.1 RQ1: Can workload modulation produce measurable and temporally structured changes in power factor and harmonic distortion?

The recorded load sweep and pulse modulation experiments provide affirmative evidence for the measurable component of RQ1. Changes in CPU activity were accompanied by observable variations in the active and apparent power, power factor, and harmonic content of the input current. During idle captures, each PSU presented a comparatively stable harmonic profile without a deliberate temporal structure. When controlled workload patterns were introduced, the extracted features changed in temporal alignment with the imposed CPU activity.

The observed effect depended on the PSU topology. The no-PFC capture contained the strongest low-order harmonic content, the passive-PFC capture produced an intermediate response, and the active-PFC capture contained the weakest harmonic signature, consistent with current shaping. These observations support the physical premise that workload modulation can affect measurable power quality quantities.

However, the experiment did not establish statistical repeatability across independent runs. Each workload, PSU, receiver placement, payload, and symbol-duration configuration was acquired once. The recurrence of symbol-related changes within one captured waveform is not equivalent to reproducing the same BER, SNR, or decode outcome during a new execution. Therefore, RQ1 can be answered affirmatively with respect to measurability and temporal structure under the evaluated conditions, whereas run-to-run repeatability remains unquantified.

5.1.2 RQ2: Can these measurable patterns support end-to-end data recovery, and what decoding performance is observed under the evaluated conditions?

OOK experiments provide proof-of-concept evidence that the measured changes can carry recoverable data. Complete recovery was observed for an ASCII string, an AES-128 example key, and an ECDSA secret in selected combinations of PSU topology, receiver placement, symbol duration, and electrical features. This demonstrates the end-to-end operation of the corresponding recorded transmissions.

The channel had a low nominal raw rate. With $T_s = 2\,\mathrm{s}$, binary OOK produces $0.5\,\mathrm{bit/s}$; the $T_s = 1\,\mathrm{s}$ AES-128 test produces $1\,\mathrm{bit/s}$. The preamble, size field, and Reed–Solomon parity reduce the payload goodput below these raw rates. Moreover, the 1 s and 2 s durations were conservative proof-of-concept parameters rather than values selected through a systematic rate-optimization experiment.

Capacity and throughput implications

The dominant practical constraint is the time required for a workload change to become distinguishable in a power factor, power, or harmonic estimate computed over a finite window. The receiver does not observe the CPU command directly; it observes a grid-side electrical response shaped by the PSU input stage, selected feature-estimation window, and surrounding electrical environment. Therefore, shorter symbols may be increasingly smoothed or obscured by transient and background effects.

In the recorded captures, the set of successful feature decodes depended strongly on the PSU topology. The no-PFC configuration produced the highest harmonic SNR at both receiver locations and allowed all three payloads to be recovered through harmonic features at the office-level point. The passive PFC captures produced a larger number of successful RMS-harmonic decodes near the PSU, but a more feature-selective set at the office location. The active-PFC office-level capture produced only one successful dBc-harmonic condition and one successful PF condition in the summary, restricted to PF and harmonic observables.

Receiver placement also changed the observed feature set. The office-level receiver measured the target contribution, together with other connected loads and path effects. Some no-PFC office-level features produced equal or better observed BER than their near-PSU counterparts,

whereas active-PFC office-level capture retained substantially fewer successful harmonic features. These are descriptive comparisons between individual captures and should not be interpreted as expected performance distributions.

The practical implication is that the channel is unsuitable for bulk transfer but can carry compact, high-value data under favorable conditions. The evaluated rates are compatible with secrets such as credentials, cryptographic keys, short configuration values, commands, or beacons. A claim about the expected packet-success probability, sustained goodput, or maximum rate would require repeated transmissions and a dedicated sweep over symbol durations.

5.1.3 RQ3: How can the power supply topology and receiver placement affect the observability and practical viability of the proposed covert channel?

The recorded data indicate that the PSU topology is a central determinant of channel observability. As summarized in Table 4.5, the mean harmonic SNR at the office-level receiver was 15.9 dB for the no-PFC configuration, 11.4 dB for passive PFC, and 8.1 dB for active PFC. At the near-PSU point, the corresponding mean harmonic SNR values were 14.8 dB, 13.2 dB, and 8.5 dB. These values are arithmetic summaries of features in the recorded datasets and not means across experimental repetitions.

Successful feature counts followed the same broad pattern. At the office-level receiver, the no-PFC captures recovered all three payloads using harmonic features, whereas the active-PFC summary contained no successful RMS-harmonic conditions, one successful dBc-harmonic condition, and one successful PF condition. The passive PFC occupied an intermediate position and remained dependent on the chosen feature and payload.

Receiver placement should not be interpreted as a geometric distance alone. The near-PSU receiver observes a comparatively direct contribution from the target, whereas the office-level receiver observes an aggregate current containing the target, background loads, electrical path effects, and measurement noise. The experiments therefore demonstrated that a signal was recoverable at both evaluated observation points in the selected configurations, but they did not establish a maximum distance or a general attenuation law.

Overall, the single-run proof-of-concept results are consistent with active PFC suppressing the low-order harmonic observables used by the channel and with upstream measurement reducing the available feature set in the most strongly corrected configuration. Statistical confirmation of these differences would require repeated independent acquisitions under controlled and varying background load conditions.

5.2. Comparison with Existing Power-Line Covert Channels

The proposed method differs from previous power-line covert channels. PowerHammer [1] modulates the CPU load and measures the resulting line-current variations, with the performance depending on the receiver's access to the relevant electrical path. In contrast, this study focuses on the derived power quality quantities, particularly low-order current harmonics and power factor, as indirect manifestations of workload-dependent electrical behavior.

The objective was not to maximize the data rate but to test whether power-quality metrics can serve as information-bearing observables. The receiver operates on quantities that can be derived from voltage and current waveforms and that are also associated with power analysis and monitoring systems. This shifts the emphasis from raw current modulation to power quality variation as a covert side channel.

Compared with the broader air-gap methods surveyed in Chapter 3, the proposed channel offers substantially lower throughput than high-bandwidth optical or electromagnetic methods. However, it does not require line of sight, audio peripherals, or a nearby free-space receiver. Instead, its operation depends on access to a shared electrical path and the observability of the target's contribution among other connected loads.

A further distinction is the role of the PSU topology. In the evaluated captures, no-PFC, passive-PFC, and active-PFC units produced markedly different harmonic signatures and successful feature sets. Therefore, for this class of channels, the power-conversion hardware is not a minor implementation detail but part of the communication path. The measurements are also consistent with the threat model in Section 1.4: an adversary would benefit from an absent, defective, or deliberately degraded PFC. The present study did not evaluate hardware tampering or supply chain manipulation; therefore, this remains a threat-model implication rather than an experimentally validated attack step.

From a defensive perspective, this comparison suggests two complementary measures. First, the effective active PFC and cleaner input-current shaping reduced the harmonic observability in the evaluated captures. Second, power quality monitoring may help identify unusual periodic changes in the active power, apparent power, power factor, or low-order harmonic amplitudes. Therefore, the proof of concept broadens the electrical-infrastructure attack surface to include both raw conducted signals and derived power-quality observables.

5.3. Defensive Implications

The electrical infrastructure should be considered as part of the attack surface of air-gapped systems. The recorded transmissions demonstrate that the derived power-quality metrics can contain recoverable information under the evaluated conditions, even though the nominal data rate is low.

The first implication concerns the architectural aspect. In the recorded active-PFC captures, harmonic amplitudes were weaker, and fewer office-level harmonic conditions produced valid payload recovery. Effective PFC and cleaner input-current shaping are, therefore, plausible mitigation factors for this specific mechanism, although the single-run design does not quantify the expected reduction in attack success.

The second implication of this study was monitoring. Periodic symbol-like variations in the active power, apparent power, power factor, or low-order harmonic amplitudes may indicate workload-driven signaling. Power-quality monitoring systems could therefore complement existing emanation-security controls, provided that detection thresholds account for legitimate workload changes and normal background-load variability.

The third implication concerns hardware integrity. A defective, degraded, or deliberately modified PFC stage can increase harmonic leakage and improve signal observability. PSU selection, procurement controls, and periodic validation are consequently relevant in high-assurance environments, although deliberate PFC manipulation was not tested in this study.

5.4. Limitations

The central limitation is the absence of repeated independent runs. Each combination of payload, symbol duration, PSU topology, and receiver placement was transmitted and acquired once. Consequently, the reported BER, SNR, and packet-recovery outcomes describe individual captures and do not characterize the run-to-run variance, expected BER, confidence intervals, long-term reliability, or statistical significance between configurations.

The different features evaluated from one waveform, for example, several harmonic orders, power factor, active power, and apparent power, were not independent repetitions. Similarly, the occurrence of many symbols within one frame does not replace independent experimental executions. Therefore, comparisons between PSU topologies and receiver placements should be interpreted as exploratory evidence and hypothesis-generating observations.

The experiments were also performed with a limited number of PSU types, receiver placements, payloads, and background load conditions. Different buildings, wiring topologies, grounding conditions, electrical panels, filters, UPS systems, temperatures, and nearby nonlinear loads can change the observed channel.

The two receiver locations represent practical observation points near the PSU and at an upstream office level point. They do not establish a maximum range in meters. For this conducted channel, the relevant factor is the electrical path observability, including branching, impedance, filtering, and the amount of unrelated load current mixed with the target contribution.

The modulation and decoding scheme was intentionally simplified. The transmitter uses binary OOK, and the receiver evaluates one feature at a time using threshold-based decoding. More advanced modulation, adaptive thresholds, multi-feature classifiers, or alternative error-control schemes may improve the observed performance, but such improvements require separate validation.

Finally, this study prioritizes interpretability and physical feasibility over rate optimization. The selected $T_s = 1\,\mathrm{s}$ and $T_s = 2\,\mathrm{s}$ values were conservative engineering choices, and no systematic symbol-duration sweep was performed. Therefore, the results show that power-quality observables carry recoverable information in selected recorded transmissions, but they do not represent an optimized or statistically characterized covert channel.

6. Conclusion

This dissertation investigated whether deliberate computational-workload modulation can produce measurable changes in electrical power quality observables, and whether those changes can carry data from an air-gapped system. The study combined a theoretical analysis of the power factor and harmonic generation in switched-mode power supplies with an exploratory proof-of-concept evaluation covering three PSU topologies, two receiver placements, two symbol durations, and three representative payloads.

The recorded experiments provide proof-of-concept evidence that workload changes can produce temporally structured variations in low-order harmonic amplitudes, power factor, active power, and apparent power. Complete payload recovery was observed in selected combinations of PSU topology, receiver placement, symbol duration, and electrical features. Binary OOK produced nominal raw rates of 0.5 bit/s for $T_s = 2$ s and 1 bit/s for the additional $T_s = 1$ s AES-128 configuration before accounting for the preamble, size field, and Reed–Solomon redundancy. Because each experimental configuration was transmitted and acquired once ($N = 1$), these outcomes demonstrated feasibility under the evaluated conditions; they did not establish an expected BER, packet success probability, run-to-run repeatability, or long-term channel reliability.

The PSU topology was a central determinant of the observed channel. In the recorded office-level datasets, the mean harmonic SNR was 15.9 dB for the no-PFC configuration, 11.4 dB for the passive PFC, and 8.1 dB for the active PFC. The no-PFC office-level captures allowed all three payloads to be recovered through at least one harmonic feature, whereas the active-PFC office-level summary contained only one successful relative harmonic condition and one successful power-factor condition. These values summarize features within individual captures rather than repeated trials, but they are consistent with active current shaping, which suppresses the low-order harmonic observables used by the proposed receiver.

The main contribution of this dissertation is the demonstration that derived low-frequency power-quality quantities can serve as information-bearing observables. Unlike PowerHammer, which focuses primarily on raw line-current modulation, and NoDE, which exploits high-frequency switching noise, this work evaluates the true power factor and individual low-order current-harmonic amplitudes across different PFC designs and receiver placements. Therefore, the

results extend the electrical infrastructure covert-channel taxonomy from raw conducted emissions to routinely derived power-quality metrics.

Another observation from the recorded dataset is that the absolute harmonic RMS amplitudes produced more successful feature decodes than harmonics normalized to the fundamental in dBc. Relative harmonics remain useful for characterizing the spectral shape, but normalization can conceal modulation when the fundamental and higher-order components scale together. Under the evaluated conditions, the absolute amplitudes more directly preserved the workload-dependent change used for symbol detection. This is a descriptive result from the available captures and should be verified through repeated experiments.

From a defensive perspective, the findings support treating the electrical supply as part of the attack surface of air-gapped systems. Effective active PFC and cleaner input-current shaping reduced harmonic observability in the recorded captures, whereas periodic symbol-like changes in power factor, power, or low-order harmonic amplitudes may provide indicators for monitoring. PSU procurement, hardware-integrity controls, and validation of PFC operations are also relevant in high-assurance environments, particularly when considering deliberate degradation or supply chain tampering.

The principal limitation is the absence of repeated independent runs. Each combination of payload, symbol duration, PSU topology, and receiver placement was acquired once, and the multiple electrical features extracted from one waveform were alternative analyses of the same transmission rather than independent repetitions. Therefore, the study does not report variance, confidence intervals, statistical significance, or expected performance distributions. Further limitations include the small hardware set, two receiver locations, one office electrical environment, limited background-load conditions, binary OOK, single-feature threshold decoding, and the absence of a systematic symbol-duration sweep or maximum-range study.

Future work should first repeat every configuration across independent transmissions and report the BER and packet-recovery distributions with confidence intervals. It should then evaluate additional PSU and computing platforms, varied and controlled background loads, different wiring topologies, filters and UPS systems, and a systematic sweep of symbol duration. Multifeature detection, adaptive thresholds, improved error control, and alternative modulation schemes may increase the achievable goodput. Finally, defensive research should determine whether workload-induced power quality patterns can be distinguished from legitimate computational activity with an acceptable false-positive rate.

In conclusion, the recorded proof-of-concept experiments show that deliberate workload modulation can create power quality variations capable of carrying recoverable data under selected conditions. The channel is low-rate and strongly dependent on the PSU topology, selected observables, receiver placement, and electrical environment. Although the present study does not establish statistical repeatability, it demonstrates that physical isolation does not remove the electrical supply as a possible communication path, and that power-quality observables deserve consideration in the security design and monitoring of high-assurance systems.

References

[1] M. Guri, B. Zadov, D. Bykhovsky, and Y. Elovici, “PowerHammer: Exfiltrating data from air-gapped computers through power lines,” 4 2018. [Online]. Available: https://arxiv.org/abs/1804.04014 [Cited on pages 5, 19, 20, 21, 30, and 69.]

[2] Z. Shao, M. A. Islam, and S. Ren, “Your noise, my signal: Exploiting switching noise for stealthy data exfiltration from desktop computers,” 1 2020. [Online]. Available: https://arxiv.org/abs/2001.06729 [Cited on pages 5, 19, 20, and 30.]

[3] MITRE ATT&CK. (2020) Proxy: Domain fronting. [accessible - 29/06/2026]. [Online]. Available: https://attack.mitre.org/techniques/T1090/004/ [Cited on page 9.]

[4] Defense Security Cooperation Agency. (2025) Secret internet protocol router network (siprnet). [accessible - 29/06/2026]. [Online]. Available: https://apps.dtic.mil/sti/html/tr/ADA512554/index.html [Cited on page 9.]

[5] U.S. Department of Defense. (2019) Intelligence communications system gets tech refresh. [accessible - 29/06/2026]. [Online]. Available: https://www.war.gov/News/News-Stories/Article/Article/1954347/intelligence-communications-system-gets-tech-refresh/ [Cited on page 9.]

[6] Microsoft. (2000) Microsoft security advisor: Nato cronos. [accessible - 29/06/2026]. [Online]. Available: https://web.archive.org/web/20000229230759/http://www.microsoft.com/security/resources/NATOCaseStudy.asp [Cited on page 9.]

[7] J. Scahill, J. Begley, and G. Greenwald. (2015, 2) The great sim heist: How spies stole the keys to the world’s cell phones. The Intercept. [accessible - 29/06/2026]. [Online]. Available: https://theintercept.com/2015/02/19/great-sim-heist/ [Cited on page 9.]

[8] Juniper Networks. (2015, 12) Juniper discovers unauthorized code in its firewall os. [accessible - 29/06/2026]. [Online]. Available: https://supportportal.juniper.net/s/article/2015-12-Out-of-Cycle-Security-Bulletin-ScreenOS-Multiple-Security-issues-with-ScreenOS-CVE-2015-7755-CVE-2015-7756 [Cited on page 9.]

[9] Kaspersky Lab, “Equation group: Questions and answers,” Kaspersky Lab, Tech. Rep., 2015, [accessible - 29/06/2026]. [Online]. Available: https://securelist.com/equation-group-questions-and-answers/68750/ [Cited on page 9.]

[10] D. Genkin, I. Pipman, and E. Tromer, “Get your hands off my laptop: physical side-channel key-extraction attacks on pcs,” *Journal of Cryptographic Engineering*, vol. 5, pp. 95–112, 2015. [Online]. Available: https://doi.org/10.1007/s13389-015-0100-7 [Cited on page 10.]

[11] D. Genkin, A. Shamir, and E. Tromer, “RSA key extraction via low-bandwidth acoustic cryptanalysis,” Cryptology ePrint Archive, Paper 2013/857, 2013. [Online]. Available: https://eprint.iacr.org/2013/857 [Cited on page 10.]

[12] OMIE, “Day-ahead and intraday electricity market rules: Non-binding translation of the electricity market rules,” OMI, Polo Español S.A. (OMIE), Tech. Rep., 2025, [accessible - 29/06/2026]. [Online]. Available: https://www.omie.es/sites/default/files/2025-03/r_28022025_en.pdf [Cited on page 13.]

[13] F. Ullah, M. Edwards, R. Ramdhany, R. Chitchyan, M. A. Babar, and A. Rashid, “Data exfiltration: A review of external attack vectors and countermeasures,” *Journal of Network and Computer Applications*, vol. 101, pp. 18–54, 1 2018. [Cited on page 19.]

[14] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten, “Lest we remember,” *Communications of the ACM*, vol. 52, pp. 91–98, 5 2009. [Cited on page 19.]

[15] H. Seong, I. Kim, Y. Jeon, M.-K. Oh, S. Lee, and D. Choi, “Practical covert wireless uni-directional communication in ieee 802.11 environment,” *IEEE Internet of Things Journal*, pp. 1–1, 2022. [Cited on page 19.]

[16] T. D. Balsdon, A. Grisel-Davy, and S. Fischmeister, “Breaking tempest: Low-frequency bidirectional covert channel on power lines,” in *Proceedings of the 12th International Conference on Information Systems Security and Privacy*. SCITEPRESS - Science and Technology Publications, 2026, pp. 422–433. [Online]. Available: https://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0014331000004061 [Cited on pages 19, 20, and 30.]

[17] M. Guri, G. Kedma, A. Kachlon, and Y. Elovici, “Airhopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies,” in *2014 9th International Conference on Malicious and Unwanted Software: The Americas (MALWARE)*. IEEE, 10 2014, pp. 58–67. [Cited on pages 20, 25, and 26.]

[18] M. Guri, A. Kachlon, O. Hasson, G. Kedma, Y. Mirsky, and Y. Elovici, "GSMem: Data exfiltration from air-gapped computers over gsm frequencies," in *24th USENIX Security Symposium (USENIX Security 15)*. Washington, D.C.: USENIX Association, 8 2015, pp. 849–864. [Online]. Available: https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/guri [Cited on pages 20 and 25.]

[19] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *2015 IEEE 28th Computer Security Foundations Symposium*. IEEE, 7 2015, pp. 276–289. [Cited on page 20.]

[20] M. Guri, Y. Solewicz, and Y. Elovici, "Fansmitter: Acoustic data exfiltration from air-gapped computers via fans noise," *Computers and Security*, vol. 91, 4 2020. [Cited on pages 20 and 22.]

[21] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, "Diskfiltration: Data exfiltration from speakerless air-gapped computers via covert hard drive noise," 2016. [Cited on pages 20 and 23.]

[22] M. Guri, B. Zadov, and Y. Elovici, "LED-it-GO: Leaking (a lot of) data from air-gapped computers via the (small) hard drive led," in *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA 2017)*, ser. Lecture Notes in Computer Science, vol. 10327. Springer, 2017, pp. 161–184. [Online]. Available: https://arxiv.org/abs/1702.06715 [Cited on pages 19, 20, and 31.]

[23] M. Guri, B. Zadov, D. Bykhovsky, and Y. Elovici, "CTRL-ALT-LED: Leaking data from air-gapped computers via keyboard leds," 7 2019. [Online]. Available: https://arxiv.org/abs/1907.05851 [Cited on page 20.]

[24] M. Guri, "Magneto: Covert channel between air-gapped systems and nearby smartphones via cpu-generated magnetic fields," *Future Generation Computer Systems*, vol. 115, pp. 115–125, 2 2021. [Cited on pages 20 and 24.]

[25] ——, "Air-fi: Leaking data from air-gapped computers using wi-fi frequencies," *IEEE Transactions on Dependable and Secure Computing*, vol. 20, pp. 2547–2564, 5 2023. [Cited on pages 20, 25, and 27.]

[26] ——, "Exfiltrating data from air-gapped computers via vibrations," *Future Generation Computer Systems*, vol. 122, pp. 69–81, 9 2021. [Cited on pages 20 and 21.]

[27] Z. Zhan, Z. Zhang, and X. Koutsoukos, “Bitjabber: The world’s fastest electromagnetic covert channel,” in *2020 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. IEEE, 12 2020, pp. 35–45. [Cited on pages 19, 20, 25, and 27.]

[28] M. Guri, “LANTENNA: Exfiltrating data from air-gapped networks via ethernet cables,” 9 2021. [Online]. Available: https://arxiv.org/abs/2110.00104 [Cited on pages 20 and 29.]

[29] X. Sun, Y. Zheng, W. Xi, Z. Chen, Z. Chen, H. Hao, Z. Jiang, and S. Zhong, “TEMPEST-LoRa: Cross-technology covert communication,” 6 2025. [Online]. Available: https://arxiv.org/abs/2506.21069 [Cited on pages 20 and 28.]

[30] J. Feng, T. Jacques, O. Abari, and N. Sehatbakhsh, “Everything has its bad side and good side: Turning processors to low overhead radios using side-channels,” in *IPSN 2023 - Proceedings of the 2023 22nd International Conference on Information Processing in Sensor Networks*. Association for Computing Machinery, Inc, 5 2023, pp. 288–301. [Cited on pages 20 and 29.]

[31] M. Guri, B. Zadov, and Y. Elovici, “Odini: Escaping sensitive data from faraday-caged, air-gapped computers via magnetic fields,” *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1190–1203, 2020. [Cited on pages 20 and 24.]

[32] M. Guri, “POWER-SUPPLaY: Leaking data from air-gapped systems by turning the power supplies into speakers,” 5 2020. [Online]. Available: https://arxiv.org/abs/2005.00395 [Cited on page 23.]

[33] M. Guri and M. Monitz, “Lcd tempest air-gap attack reloaded,” in *2018 IEEE International Conference on the Science of Electrical Engineering in Israel (ICSEE)*. IEEE, 12 2018, pp. 1–5. [Cited on pages 25 and 28.]

[34] M. Guri, O. Hasson, G. Kedma, and Y. Elovici, “An optical covert-channel to leak data through an air-gap,” in *2016 14th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 2016, pp. 642–649. [Online]. Available: https://arxiv.org/abs/1607.03946 [Cited on page 31.]

[35] M. Guri and D. Bykhovsky, “air-jumper: Covert air-gap exfiltration/infiltration via security cameras and infrared (ir),” *Computers and Security*, vol. 82, pp. 15–29, 5 2019. [Cited on page 31.]

A. Decode results

This appendix reports descriptive decoding outcomes from the exploratory proof-of-concept experiments. Each combination of payload, symbol duration, PSU topology, and receiver placement corresponds to one recorded transmission, so the number of independent runs per configuration is $N = 1$. Different harmonics, power factor, active power, and apparent power extracted from the same waveform are alternative receiver observables; they are not independent experimental repetitions.

The BER values are the fractions of differing bits reported by the decoder for the corresponding recorded frame; where present, the value in parentheses is the observed number of differing bits. These values describe individual captures and must not be interpreted as estimates of the expected channel BER or as evidence of run-to-run reliability. In the "Outcome" column, "Yes" denotes payload recovery with valid Reed–Solomon verification, "No" denotes that a valid payload was not recovered, and "RS invalid" denotes a candidate recovery for which Reed–Solomon verification failed. RS-invalid entries are excluded from the successful-decode counts reported in Chapter 4.

A.1. No Power Factor Correction Power Supply

Table A.1: Descriptive decoding outcomes for Test 4 (ASCII string payload) – no-PFC PSU ($N = 1$ per receiver placement).

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	2	6.77	0.5278 (76)	No
At PSU	5th dBc	2	8.30	1	No
At PSU	7th dBc	2	9.11	0.1319 (19)	No
At PSU	11th dBc	2	9.33	1	No
At PSU	3rd RMS	2	20.00	0.0208 (3)	No
At PSU	5th RMS	2	19.34	0.0278 (4)	No
At PSU	7th RMS	2	22.32	0.1528 (22)	No
At PSU	11th RMS	2	19.63	0.0417 (6)	No
At PSU	PF	2	10.00	1	No
At PSU	Apparent power	2	21.42	0	Yes
At PSU	Active power	2	18.97	0.007 (1)	Yes

Continued on next page

Table A.1 – continued from previous page

Location	Feature	$\mathbf{T_s}$ (s)	SNR (dB)	Observed BER	Outcome
Office-level	3rd dBc	2	8.8392	1	No
Office-level	5th dBc	2	18.1479	0.042 (6)	No
Office-level	7th dBc	2	19.9559	0.049 (7)	No
Office-level	11th dBc	2	14.1316	0	Yes
Office-level	3rd RMS	2	13.08	0	Yes
Office-level	5th RMS	2	13.31	1	No
Office-level	7th RMS	2	17.01	0.1528 (22)	No
Office-level	11th RMS	2	19.95	0	Yes
Office-level	PF	2	10.4453	0.0208 (3)	No
Office-level	Apparent power	2	17.1188	0	Yes
Office-level	Active power	2	16.7393	0	Yes

Table A.2: Descriptive decoding outcomes for Tests 5 and 6 (AES-128 payload, $T_s = 2$ s and $T_s = 1$ s) – no-PFC PSU ($N = 1$ per configuration).

Location	Feature	$\mathbf{T_s}$ (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	1	5.83	1	No
At PSU	3rd dBc	2	5.83	0.5096 (53)	No
At PSU	5th dBc	1	9.54	1	No
At PSU	5th dBc	2	9.54	0.5547 (71)	No
At PSU	7th dBc	1	11.18	0.3906 (50)	No
At PSU	7th dBc	2	11.18	0	Yes
At PSU	11th dBc	1	11.35	0.3828 (49)	No
At PSU	11th dBc	2	11.35	0.0156 (2)	Yes
At PSU	3rd RMS	1	20.96	0.0078 (1)	Yes
At PSU	3rd RMS	2	20.96	0	Yes
At PSU	5th RMS	1	20.26	0.1250 (16)	No
At PSU	5th RMS	2	20.26	0	Yes
At PSU	7th RMS	1	21.72	0.0078 (1)	Yes
At PSU	7th RMS	2	21.72	0.0313 (4)	No
At PSU	11th RMS	1	19.84	0.1953 (25)	No
At PSU	11th RMS	2	19.84	0	Yes

Continued on next page

Table A.2 – continued from previous page

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	PF	1	11.51	0.3906 (50)	No
At PSU	PF	2	11.51	0.3671 (47)	No
At PSU	Apparent power	1	22.28	0.4062 (52)	No
At PSU	Apparent power	2	22.28	0	Yes
At PSU	Active power	1	19.47	0.3984 (51)	No
At PSU	Active power	2	19.47	0	Yes
Office-level	3rd dBc	1	8.72	0.4766 (61)	No
Office-level	3rd dBc	2	8.72	1	No
Office-level	5th dBc	1	17.26	0	Yes
Office-level	5th dBc	2	17.26	0	Yes
Office-level	7th dBc	1	19.30	0.0156 (2)	Yes
Office-level	7th dBc	2	19.30	0	Yes
Office-level	11th dBc	1	16.04	0.0234 (3)	No
Office-level	11th dBc	2	16.04	0.0938 (12)	No
Office-level	3rd RMS	1	12.80	0	Yes
Office-level	3rd RMS	2	12.80	0	Yes
Office-level	5th RMS	1	12.56	0	Yes
Office-level	5th RMS	2	12.56	0.5391 (69)	No
Office-level	7th RMS	1	19.25	0.2186 (28)	No
Office-level	7th RMS	2	19.25	0	Yes
Office-level	11th RMS	1	20.08	0.0781 (1)	Yes
Office-level	11th RMS	2	20.08	0.0469 (6)	No
Office-level	PF	1	10.07	0	RS invalid
Office-level	PF	2	10.07	0.0469 (6)	No
Office-level	Apparent power	1	15.23	0	Yes
Office-level	Apparent power	2	15.23	0	Yes
Office-level	Active power	1	15.07	0	Yes
Office-level	Active power	2	15.07	0	Yes

Table A.3: Descriptive decoding outcomes for Test 7 (ECDSA secp256k1 secret-key payload) – no-PFC PSU ($N = 1$ per receiver placement).

Location	Feature	$\mathbf{T_s}$ (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	2	6.28	1	No
At PSU	5th dBc	2	8.90	0.5039 (129)	No
At PSU	7th dBc	2	10.14	0.004 (1)	Yes
At PSU	11th dBc	2	9.75	0.070 (18)	No
At PSU	3rd RMS	2	21.47	0	Yes
At PSU	5th RMS	2	20.37	0	Yes
At PSU	7th RMS	2	21.18	0.0078 (2)	Yes
At PSU	11th RMS	2	19.76	0	Yes
At PSU	PF	2	9.95	0.4531 (116)	No
At PSU	Apparent power	2	22.93	0	Yes
At PSU	Active power	2	20.03	0	Yes
Office-level	3rd dBc	2	8.47	0.5398 (95)	No
Office-level	5th dBc	2	19.49	0.0508 (13)	No
Office-level	7th dBc	2	19.23	0.1797 (46)	No
Office-level	11th dBc	2	14.33	0	Yes
Office-level	3rd RMS	2	17.11	0	Yes
Office-level	5th RMS	2	17.12	0.1758 (45)	No
Office-level	7th RMS	2	18.55	0.1406 (36)	No
Office-level	11th RMS	2	18.91	0	Yes
Office-level	PF	2	10.10	0.0469 (12)	No
Office-level	Apparent power	2	16.50	0.0078 (2)	Yes
Office-level	Active power	2	15.63	0	Yes

A.2. Passive Power Factor Correction Power Supply

Table A.4: Descriptive decoding outcomes for Test 4 (ASCII string payload) – passive-PFC PSU ($N = 1$ per receiver placement).

Location	Feature	$\mathbf{T_s}$ (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	2	7.10	0.465 (67)	No
At PSU	5th dBc	2	9.24	0.007 (1)	Yes
At PSU	7th dBc	2	9.68	0.097 (14)	No

Continued on next page

Table A.4 – continued from previous page

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	11th dBc	2	9.024	0.0694 (10)	No
At PSU	3rd RMS	2	20.38	0.1944 (28)	No
At PSU	5th RMS	2	18.89	0.0208 (3)	Yes
At PSU	7th RMS	2	14.66	0.0903 (13)	No
At PSU	11th RMS	2	12.63	0.0069 (1)	Yes
At PSU	PF	2	10.21	0.0347 (5)	No
At PSU	Active power	2	19.25	0.007 (1)	Yes
Office-level	3rd dBc	2	9.76	0.0208 (3)	No
Office-level	5th dBc	2	8.43	0.475 (57)	No
Office-level	7th dBc	2	10.13	0.0347 (5)	No
Office-level	11th dBc	2	10.90	0	Yes
Office-level	3rd RMS	2	17.58	0.0138 (2)	Yes
Office-level	5th RMS	2	16.45	0	Yes
Office-level	7th RMS	2	8.23	1	No
Office-level	11th RMS	2	8.78	1	No
Office-level	PF	2	8.62	0.4722 (68)	No
Office-level	Apparent power	2	16.68	0.007 (1)	Yes
Office-level	Active power	2	15.45	0.007 (1)	Yes

Table A.5: Descriptive decoding outcomes for Tests 5 and 6 (AES-128 payload, $T_s = 2$ s and $T_s = 1$ s) – passive-PFC PSU ($N = 1$ per configuration).

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	1	6.50	1	No
At PSU	3rd dBc	2	6.50	1	No
At PSU	5th dBc	1	10.37	0.0313 (4)	No
At PSU	5th dBc	2	10.37	0	RS invalid
At PSU	7th dBc	1	11.16	0.008 (1)	Yes
At PSU	7th dBc	2	11.16	0.0781 (10)	No
At PSU	11th dBc	1	9.76	0.0625 (8)	No
At PSU	11th dBc	2	9.76	0.0625 (8)	No
At PSU	3rd RMS	1	21.46	0	Yes

Continued on next page

Table A.5 – continued from previous page

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd RMS	2	21.46	0	Yes
At PSU	5th RMS	1	19.91	0	Yes
At PSU	5th RMS	2	19.91	0	Yes
At PSU	7th RMS	1	15.56	0	Yes
At PSU	7th RMS	2	15.56	0	Yes
At PSU	11th RMS	1	12.75	0.0078 (1)	Yes
At PSU	11th RMS	2	12.75	0	Yes
At PSU	PF	1	12.02	0	Yes
At PSU	PF	2	12.02	0.0078 (1)	Yes
At PSU	Apparent power	1	21.62	0	Yes
At PSU	Apparent power	2	21.62	0	Yes
At PSU	Active power	1	19.80	0	Yes
At PSU	Active power	2	19.80	0	Yes
Office-level	3rd dBc	1	10.55	0	Yes
Office-level	3rd dBc	2	10.55	0.008 (1)	No
Office-level	5th dBc	1	8.73	1	No
Office-level	5th dBc	2	8.73	0.1875 (24)	No
Office-level	7th dBc	1	10.79	0.008 (1)	Yes
Office-level	7th dBc	2	10.79	0	Yes
Office-level	11th dBc	1	8.52	0.0156 (2)	Yes
Office-level	11th dBc	2	8.52	0.5 (64)	No
Office-level	3rd RMS	1	19.30	0	Yes
Office-level	3rd RMS	2	19.30	0	Yes
Office-level	5th RMS	1	17.44	0.0078 (1)	Yes
Office-level	5th RMS	2	17.44	0	Yes
Office-level	7th RMS	1	8.39	1	No
Office-level	7th RMS	2	8.39	0.3984 (51)	No
Office-level	11th RMS	1	8.44	0.4922 (63)	No
Office-level	11th RMS	2	8.44	0.4297 (55)	No
Office-level	PF	1	8.73	0.3906 (59)	No
Office-level	PF	2	8.73	0.4772 (42)	No

Continued on next page

Table A.5 – continued from previous page

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
Office-level	Apparent power	1	17.10	0.0078 (1)	Yes
Office-level	Apparent power	2	17.10	0	Yes
Office-level	Active power	1	15.68	0.0078 (1)	Yes
Office-level	Active power	2	15.68	0	Yes

Table A.6: Descriptive decoding outcomes for Test 7 (ECDSA secp256k1 secret-key payload) – passive-PFC PSU ($N = 1$ per receiver placement).

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	2	7.20	0.5 (124)	No
At PSU	5th dBc	2	9.82	0.0117 (3)	No
At PSU	7th dBc	2	10.64	0.0156 (4)	No
At PSU	11th dBc	2	9.36	0.0039 (1)	Yes
At PSU	3rd RMS	2	21.86	0	Yes
At PSU	5th RMS	2	20.08	0	Yes
At PSU	7th RMS	2	15.11	0.0039 (1)	Yes
At PSU	11th RMS	2	12.44	0.0039 (1)	Yes
At PSU	PF	2	10.88	0.0234 (6)	No
At PSU	Apparent power	2	21.96	0.0742 (19)	No
At PSU	Active power	2	20.00	0	Yes
Office-level	3rd dBc	2	7.20	0.0156 (4)	No
Office-level	5th dBc	2	9.82	0.4 (48)	No
Office-level	7th dBc	2	10.64	0.0156 (4)	No
Office-level	11th dBc	2	9.36	0.0664 (17)	No
Office-level	3rd RMS	2	18.77	0.2188 (56)	No
Office-level	5th RMS	2	17.72	0	Yes
Office-level	7th RMS	2	8.57	0.4861 (70)	No
Office-level	11th RMS	2	8.50	0.2070 (53)	No
Office-level	PF	2	10.87	0.5903 (85)	No
Office-level	Apparent power	2	21.96	0.2344 (60)	No
Office-level	Active power	2	20.00	0.2305 (59)	No

A.3. Active Power Factor Correction Power Supply

Table A.7: Descriptive decoding outcomes for Test 4 (ASCII string payload) – active-PFC PSU ($N = 1$ per receiver placement).

Location	Feature	$\mathbf{T_s}$ (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	2	7.25	0.2708 (39)	No
At PSU	5th dBc	2	8.72	0.0278 (4)	No
At PSU	7th dBc	2	8.46	0.4375 (63)	No
At PSU	11th dBc	2	8.52	0.4931 (71)	No
At PSU	3rd RMS	2	9.03	0.0278 (4)	No
At PSU	5th RMS	2	9.01	0.4931 (71)	No
At PSU	7th RMS	2	8.85	0.4097 (59)	No
At PSU	11th RMS	2	9.09	1	No
At PSU	PF	2	13.20	0.0625 (9)	No
At PSU	Apparent power	2	17.10	0.0278 (4)	No
At PSU	Active power	2	18.56	0.007 (1)	Yes
Office-level	3rd dBc	2	8.215	0.368 (53)	No
Office-level	5th dBc	2	8.1112	0.4861 (70)	No
Office-level	7th dBc	2	7.888	0.490 (51)	No
Office-level	11th dBc	2	8.5633	1	No
Office-level	3rd RMS	2	8.07	1	No
Office-level	5th RMS	2	7.93	0.4861 (40)	No
Office-level	7th RMS	2	7.93	0.5167 (62)	No
Office-level	11th RMS	2	8.50	1	No
Office-level	PF	2	7.47	0.4375	No
Office-level	Apparent power	2	9.34	0	RS invalid
Office-level	Active power	2	11.51	0	Yes

Table A.8: Descriptive decoding outcomes for Tests 5 and 6 (AES-128 payload, $T_s = 2$ s and $T_s = 1$ s) – active-PFC PSU ($N = 1$ per configuration).

Location	Feature	$\mathbf{T_s}$ (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	1	6.82	0.2891 (37)	No
At PSU	3rd dBc	2	6.82	0.6406 (82)	No
At PSU	5th dBc	1	8.23	0.0078 (1)	Yes
At PSU	5th dBc	2	8.23	0.4141 (53)	No

Continued on next page

Table A.8 – continued from previous page

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	7th dBc	1	7.92	0.3438 (44)	No
At PSU	7th dBc	2	7.92	0.3438 (4)	No
At PSU	11th dBc	1	8.19	0.0078 (1)	Yes
At PSU	11th dBc	2	8.19	0.0703 (9)	No
At PSU	3rd RMS	1	9.11	0.0078 (1)	Yes
At PSU	3rd RMS	2	9.11	0	Yes
At PSU	5th RMS	1	9.15	1	No
At PSU	5th RMS	2	9.15	0.5078 (65)	No
At PSU	7th RMS	1	9.05	0.5078 (65)	No
At PSU	7th RMS	2	9.05	0.5268 (59)	No
At PSU	11th RMS	1	9.16	0.4609 (59)	No
At PSU	11th RMS	2	9.16	0.4766 (61)	No
At PSU	PF	1	14.60	0	Yes
At PSU	PF	2	14.60	0	Yes
At PSU	Apparent power	1	17.23	0	Yes
At PSU	Apparent power	2	17.23	0	Yes
At PSU	Active power	1	18.86	0	Yes
At PSU	Active power	2	18.86	0.0078 (1)	Yes
Office-level	3rd dBc	1	8.24	0.0234 (3)	No
Office-level	3rd dBc	2	8.24	0.0078 (1)	Yes
Office-level	5th dBc	1	7.91	0.0156 (2)	No
Office-level	5th dBc	2	7.91	1	No
Office-level	7th dBc	1	7.78	0.0156 (2)	No
Office-level	7th dBc	2	7.78	0.0859 (11)	No
Office-level	11th dBc	1	8.69	1	No
Office-level	11th dBc	2	8.69	1	No
Office-level	3rd RMS	1	8.01	0.6696 (75)	No
Office-level	3rd RMS	2	8.01	0.3281 (42)	No
Office-level	5th RMS	1	7.92	1	No
Office-level	5th RMS	2	7.92	0.6875 (44)	No
Office-level	7th RMS	1	7.96	0.3984 (51)	No

Continued on next page

Table A.8 – continued from previous page

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
Office-level	7th RMS	2	7.96	0.4688 (60)	No
Office-level	11th RMS	1	8.47	1	No
Office-level	11th RMS	2	8.47	0.3359 (43)	No
Office-level	PF	1	8.46	0.0313 (4)	No
Office-level	PF	2	8.46	0	Yes
Office-level	Apparent power	1	7.88	0.0313 (4)	No
Office-level	Apparent power	2	7.88	0.0156 (2)	No
Office-level	Active power	1	11.15	0.0078 (1)	No
Office-level	Active power	2	11.15	0.0078 (1)	Yes

Table A.9: Descriptive decoding outcomes for Test 7 (ECDSA secp256k1 secret-key payload) – active-PFC PSU ($N = 1$ per receiver placement).

Location	Feature	T_s (s)	SNR (dB)	Observed BER	Outcome
At PSU	3rd dBc	2	7.09	0.0898 (23)	No
At PSU	5th dBc	2	8.57	0	Yes
At PSU	7th dBc	2	8.15	0.0313 (8)	No
At PSU	11th dBc	2	8.28	0.0781 (20)	No
At PSU	3rd RMS	2	9.03	0.0156 (4)	No
At PSU	5th RMS	2	9.12	0.4758 (118)	No
At PSU	7th RMS	2	9.03	0.3988 (67)	No
At PSU	11th RMS	2	9.03	0.4875 (78)	No
At PSU	PF	2	13.37	0.0078 (2)	Yes
At PSU	Apparent power	2	17.58	0	Yes
At PSU	Active power	2	18.62	0	Yes
Office-level	3rd dBc	2	8.26	0.1094 (28)	No
Office-level	5th dBc	2	8.20	0.0156 (4)	No
Office-level	7th dBc	2	7.81	0.4688 (120)	No
Office-level	11th dBc	2	8.62	0.4844 (124)	No
Office-level	3rd RMS	2	8.10	0.4866 (109)	No
Office-level	5th RMS	2	8.01	1	No
Office-level	7th RMS	2	8.00	1	No

Continued on next page

Table A.9 – continued from previous page

Location	**Feature**	$\mathbf{T_s}$ (s)	**SNR (dB)**	**Observed BER**	**Outcome**
Office-level	11th RMS	2	8.60	1	No
Office-level	PF	2	7.94	0.082	No
Office-level	Apparent power	2	8.44	0.0039 (1)	Yes
Office-level	Active power	2	11.11	0.0195 (5)	No