UNIVERSITY OF MANNHEIM
School of Business Informatics and Mathematics

The Internet of the Things: The Security of the Wireless

Inauguraldissertation zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften der Universität Mannheim

vorgelegt von

Christian Müller, M. Sc.
aus Bochum

Mannheim, 2026

Dekan:	Prof. Dr. Claus Hertling, Universität Mannheim
Referent:	Prof. Dr. Frederik Armknecht, Universität Mannheim
Korreferent:	Prof. Dr. Mathy Vanhoef, Katholieke Universiteit Leuven, Belgien

Tag der mündlichen Prüfung: 22. September 2026

Abstract

The Internet of Things (IoT) with its increased use and proliferation into numerous everyday aspects connects various types of devices from various domains to the Internet, e. g., wearables like smartwatches or home automation gadgets like smart heating and lighting devices. While this allows for new insights and uncovering potential (resource) savings or optimization, data needs to be collected from within private spheres and often shared with service providers. Many devices communicate wirelessly, i. e., an IoT device is connected to some gateway using a wireless communication protocol like Zigbee, and the gateway which is connected to the Internet, e. g., via Wireless Local Area Network or a wired connection, sends received data to the service provider. Unfortunately, it is not always obvious what kind of information seemingly innocuous data, e. g., relative humidity or temperature collected in a smart home, may carry and reveal if analyzed accordingly. Furthermore, wireless communication protocols intended for IoT devices may suffer from security vulnerabilities.

In this thesis, we focus on the privacy and security of wirelessly transmitted data from IoT devices. First, we investigate the possibility to extract private information from seemingly innocuous room climate data like relative humidity and temperature. We demonstrate how presence and number of people, i. e., occupancy detection and estimation, and certain activities within a room can be inferred from collected relative humidity and temperature data by utilizing standard machine learning techniques. This example underpins the need to also safeguard such data from unauthorized access. Then, we shift our attention towards the popular lightweight communication standard Zigbee (based on IEEE 802.15.4) which includes encrypted (and authenticated) communication. To integrate new devices into an existing Zigbee network, the *touchlink commissioning* procedure is offered. Our analysis discovers design flaws leading to exploitable security vulnerabilities, e. g., extracting sensitive key material Over-The-Air (OTA) or performing device takeovers over a distance of up to 190 m. As we believe that some of the vulnerabilities' root causes stem from underlying misconceptions regarding wireless communication, we consequently propose a novel Wireless Communication Model (WCM) and Wireless Attacker Model (WAM) specifically tailored to wireless communication. This model incorporates three main properties, i. e., communication channels, signals, and locality, thus, proving the security of Distance Bounding protocols and proposing a new security definition, dubbed IND-CSA (Indistinguishability under Chosen Signal Attack). Furthermore, it is more precise as it respects that information is modulated onto a carrier wave before actually being transmitted. Inspired by our WCM and WAM, we revisit friendly jamming and lower the bound regarding hardware requirements for an attacker to successfully recover friendly-jammed information. More specifically, we consider Binary Phase-Shift Keying-modulated communication as defined in the IEEE 802.15.4 specification and present a one-antenna attacker with a maximal Bit Recovery Rate of 0.92 (average 0.70) in a real-world setting with jammed OTA transmissions, thus showing the feasibility of real-world attacks.

Zusammenfassung

Das Internet der Dinge (*englisch:* Internet of Things, IoT) mit seiner zunehmenden Nutzung und Verbreitung in zahlreiche alltägliche Bereiche verbindet verschiedenste Arten von Geräten aus unterschiedlichen Domänen mit dem Internet, z. B. Wearables wie Smartwatches oder Hausautomatisierungsgeräte wie intelligente Heizungs- und Beleuchtungssysteme. Dies ermöglicht zwar neue Erkenntnisse sowie das Aufdecken potenzieller (Ressourcen-)Einsparungen oder Optimierungen, erfordert jedoch die Erhebung von Daten aus privaten Bereichen, die häufig mit Dienstanbietern geteilt werden. Viele Geräte kommunizieren drahtlos, d. h. ein IoT-Gerät ist über ein drahtloses Kommunikationsprotokoll wie Zigbee mit einem Gateway verbunden, und das Gateway, welches mit dem Internet verbunden ist, z. B. über WLAN oder kabelgebunden, übermittelt die empfangenen Daten an den Dienstanbieter. Leider ist nicht immer offensichtlich, welche Art von Informationen scheinbar harmlose Daten, z. B. in einem Smart Home erfasste relative Luftfeuchtigkeit oder Temperatur, enthalten und bei entsprechender Analyse preisgeben können. Darüber hinaus können drahtlose Kommunikationsprotokolle, die für IoT-Geräte vorgesehen sind, Sicherheitslücken aufweisen.

In dieser Arbeit konzentrieren wir uns auf die Privatsphäre und Sicherheit drahtlos übertragener Daten von IoT-Geräten. Zunächst untersuchen wir die Möglichkeit, private Informationen aus vermeintlich harmlosen Raumklimadaten wie relativer Luftfeuchtigkeit und Temperatur zu extrahieren. Wir zeigen, wie aus erfassten Daten zur relativen Luftfeuchtigkeit und Temperatur mithilfe gängiger Machine-Learning-Techniken die Anwesenheit und Anzahl von Personen sowie bestimmte Aktivitäten innerhalb eines Raumes abgeleitet werden können. Dieses Beispiel unterstreicht die Notwendigkeit, auch solche, vermeintlich harmlose Daten vor unbefugtem Zugriff zu schützen. Anschließend richten wir unseren Fokus auf Zigbee (basierend auf IEEE 802.15.4), ein verbreiteter leichtgewichtiger Kommunikationsstandard, der verschlüsselte (und authentifizierte) Kommunikation spezifiziert. Um neue Geräte in ein bestehendes Zigbee-Netzwerk zu integrieren, gibt es das *Touchlink Commissioning*-Verfahren. Unsere Analyse identifiziert Designfehler, die zu ausnutzbaren Sicherheitslücken führen, z. B. das drahtlose Extrahieren sensiblen Schlüsselmaterials oder die Übernahme von Geräten über eine Distanz von bis zu 190 m. Wir gehen davon aus, dass einige der Schwachstellen auf zugrundeliegenden Fehlannahmen hinsichtlich drahtloser Kommunikation beruhen, und schlagen folglich ein speziell auf drahtlose Kommunikation zugeschnittenes Kommunikations- und Angreifermodell (*englisch:* Wireless Communication Model, WCM; Wireless Attacker Model, WAM) vor. Dieses Modell beinhaltet drei Haupteigenschaften — Kommunikationskanäle, Signale, Lokalität — und ermöglicht Sicherheitsbeweise für Distance-Bounding-Protokolle sowie die Einführung einer neuen Sicherheitsdefinition namens IND-CSA (Indistinguishability under Chosen Signal Attack). Darüber hinaus berücksichtigt es, dass Informationen vor der eigentlichen Übertragung auf eine Trägerwelle moduliert werden. Inspiriert von WCM und WAM betrachten wir Friendly Jamming und senken die minimalen Hardwareanforderungen für einen Angreifer, um Informationen, die durch einen Friendly Jammer gestört wurden, erfolgreich wiederherzustellen. Konkret betrachten wir mit binärer Phasenumtastung modulierte Kommunikation, wie in IEEE 802.15.4 spezifiziert, und präsentieren einen Ein-Antennen-Angreifer mit einer maximalen Bit-Rekonstruktionsrate von 0,92 (durchschnittlich 0,70) in einer realen Umgebung mit durch einen Friendly Jammer gestörten drahtlosen Übertragungen, wodurch die realistische Durchführbarkeit solcher Angriffe gezeigt wird.

Acknowledgments

Special thanks to my supervisor and advisor, Frederik Armknecht, whose professional and personal support helped me through the highs and lows of life, the universe, and everything. I am grateful to have had the opportunity to explore various research ideas and areas, for his critical thinking, for guiding and challenging me, and for his understanding and heartfelt support in personal matters and through difficult times.

Many thanks to Mathy Vanhoef, who found time in his schedule and agreed to be the second assessor of this thesis. I greatly appreciate the time and effort he dedicated to reviewing my thesis.

Some of this work would not have been possible without others. I want to thank Philipp Morgner and Zinaida Benenson for the fruitful collaboration on a DFG research project; and Vasily Mikhalev and Yves T. Staudenmaier, with whom I have burnt the midnight oil more than once and shared both the joys of writing papers and the arduous process of getting them accepted.

Furthermore, I wholeheartedly thank all my former and current colleagues for making my time at the university unforgettable for various reasons. In particular, I thank Matthias Hamann, Alex Moch, Jochen Schäfer, and Linda Scheu-Hachtel for turning late-night emergency server maintenance into fun and for the unparalleled BBB and karaoke evenings, which proved that PUR music can be moving. Thank you, Alex, for constantly finding good coffee and food. I am also grateful to Jasmin Zalonis for the nice talks and our shared suffering through tedious project work; to Philipp Schaber for his help, encouragement, and perspective on work and life in general; and to Marcel Mildenberger for his knack for brightening up an otherwise glum day. Thanks to Gabi Atkinson and Karin Teynor for their outstanding support in all organizational matters, and to Matthias Krause for always lending a sympathetic ear, for his support, and for his great hiking and wine recommendations. Likewise, I want to express my gratitude to Andreas Neuenkirch for chairing the committee for my thesis defense.

I want to thank Daniel Flachs, who shares my passion for typesetting and LaTeX issues, as well as the dubious pleasure of going through multiple review cycles for author proofs produced by publishers.

To Kati and Tobi: Thank you for always being there for me, day or night, and for the many great memories we have created.

To my family: Thank you for your unconditional love and support from the very beginning, whether you are still by my side today or supported me for as long as you were able to. I will forever carry you in my heart. In particular, I would like to express my gratitude to Sarah for being my safe haven, taking so many burdens off my shoulders, and allowing me to concentrate on the many small pieces that ultimately came together to form this thesis.

Last but not least, I sincerely thank J. R. R. Tolkien and Douglas Adams for their great and inspiring works, which have influenced my life probably more than I am willing to admit.

Contents

Abstract iii

Zusammenfassung v

Acknowledgments vii

List of Figures xv

List of Tables xvii

List of Acronyms xix

1 Introduction 1
- 1.1 Background and Motivation . . . 2
- 1.2 Need for Protection of Room Climate Data . . . 2
- 1.3 Vulnerabilities of Wireless Communication Protocols . . . 3
- 1.4 Theoretical Models Incorporating Real-World Effects . . . 4
- 1.5 Lowering Requirements for Recovering Jammed Bits . . . 5
- 1.6 Publications . . . 6
- 1.7 Organization of This Thesis . . . 8

2 Keep Room Climate Data Secret — Keep Room Climate Data Safe 9
- 2.1 Introduction . . . 10
  - 2.1.1 Research Questions . . . 10
  - 2.1.2 Experiments . . . 11
  - 2.1.3 Summary of Results . . . 11
  - 2.1.4 Organization of This Chapter . . . 12
- 2.2 Related Work . . . 12
  - 2.2.1 Occupancy Detection . . . 12
  - 2.2.2 Occupancy Estimation . . . 12
  - 2.2.3 Activity Recognition . . . 14
  - 2.2.4 This Work . . . 15
- 2.3 Threat Model . . . 15
  - 2.3.1 Attack Scenario and Goals . . . 15
  - 2.3.2 Attacker Model and Capabilities . . . 16
  - 2.3.3 Attack Procedure . . . 16
  - 2.3.4 Scope . . . 17
- 2.4 Experimental Design and Methods . . . 17
  - 2.4.1 Experimental Setup and Tasks . . . 17
  - 2.4.2 Sensor Data Collection . . . 18
    - 2.4.2.1 Hardware . . . 18

2.4.2.2 Software . . . 19
2.4.2.3 Collected Data . . . 19
2.4.3 Experimental Procedure . . . 20
2.4.4 Participants and Ethical Principles . . . 20
2.4.5 Classifier Design . . . 21
2.4.5.1 Data Segmentation . . . 21
2.4.5.2 Feature Extraction . . . 21
2.4.5.3 Feature Selection and Classification . . . 22
2.4.5.4 Evaluation Metrics . . . 22
2.5 Results . . . 22
2.5.1 Visual Inspection . . . 22
2.5.2 Occupancy Detection . . . 23
2.5.3 Occupancy Estimation . . . 24
2.5.4 Activity Recognition . . . 25
2.5.5 Multi-Sensor Classification . . . 26
2.6 Further Observations . . . 27
2.6.1 Length of Measurement Windows . . . 27
2.6.2 Selected Features . . . 29
2.6.3 Size and Layout of Rooms . . . 29
2.6.4 Position of Sensors . . . 30
2.7 Discussion . . . 30
2.7.1 Privacy Implications . . . 30
2.7.2 Policy Implications . . . 31
2.7.3 Location-Independent Classification . . . 32
2.8 Conclusion . . . 33

3 Speak Touchlink and Enter 35
3.1 Introduction . . . 36
3.1.1 Contribution . . . 36
3.1.2 Organization of This Chapter . . . 37
3.2 Zigbee . . . 37
3.2.1 System Model . . . 38
3.2.2 Security . . . 38
3.2.3 Commissioning . . . 38
3.2.3.1 EZ-Mode Commissioning . . . 39
3.2.3.2 Touchlink Commissioning . . . 39
3.3 Threat Model . . . 40
3.3.1 Security Goals and Attacker Model . . . 40
3.3.2 Threat Scenarios . . . 41
3.4 Security Analysis of the Touchlink Commissioning Procedure . . . 42
3.4.1 Penetration Testing Framework Z3sec . . . 42
3.4.2 Testbed . . . 43
3.4.3 Denial-of-Service Attacks . . . 44
3.4.3.1 Active Device Scan . . . 44
3.4.3.2 Identify Action Attack . . . 45
3.4.3.3 Reset to Factory-New Attack . . . 46
3.4.3.4 Permanent Disconnect Attack . . . 47

3.4.4 Attacks to Gain Control . . . 47
3.4.4.1 Hijack Attack . . . 48
3.4.4.2 Network Key Extraction . . . 48
3.4.5 Evaluation of Wireless Range . . . 49
3.4.5.1 Active Attacks . . . 50
3.4.5.2 Passive Attacks . . . 50
3.4.6 Recovery . . . 51
3.5 Disclosure and Response . . . 51
3.6 Discussion . . . 52
3.6.1 Usage of Touchlink . . . 52
3.6.2 Manufacturer-Specific Mitigation . . . 53
3.6.3 Renewal of the Touchlink Preconfigured Link Key . . . 53
3.6.4 EZ-Mode Commissioning . . . 53
3.6.5 Limitations of Our Work . . . 54
3.7 Related Work . . . 54
3.8 Conclusion . . . 55

4 A Cryptographic Model for Wireless Communication Is Never Late 57
4.1 Introduction . . . 58
4.1.1 Contribution . . . 58
4.1.2 Organization of This Chapter . . . 59
4.2 Background . . . 59
4.2.1 IEEE 802.15.4 . . . 59
4.2.2 Adversarial Jamming . . . 60
4.2.2.1 Reactive Jamming . . . 61
4.2.2.2 Selective Jamming . . . 61
4.2.3 Friendly Jamming . . . 61
4.2.4 Generating Jamming Signals . . . 62
4.3 Motivation . . . 62
4.3.1 The Man-in-the-Middle Model . . . 62
4.3.2 Communication Channels . . . 63
4.3.2.1 Discussion . . . 63
4.3.2.2 Lessons Learned . . . 63
4.3.3 Signals . . . 64
4.3.3.1 Discussion . . . 64
4.3.3.2 Lessons Learned . . . 64
4.3.4 Locality . . . 64
4.3.4.1 Discussion . . . 64
4.3.4.2 Lessons Learned . . . 65
4.4 Wireless Communication Model . . . 65
4.4.1 Communication Channels in the WCM . . . 65
4.4.2 Signals in the WCM . . . 67
4.4.2.1 Physical Signals . . . 67
4.4.2.2 Modulation and Demodulation . . . 68
4.4.2.3 Sending and Receiving Message Symbols . . . 69
4.4.2.4 Signal Collision . . . 70
4.4.3 Locality in the WCM . . . 71

4.5 Wireless Attacker Model . . . 73
4.5.1 Capabilities . . . 73
4.5.2 Modification of Signals Through Jamming . . . 73
4.6 Applications . . . 75
4.6.1 Frequency-Hopping Spread Spectrum . . . 75
4.6.2 Distance Bounding Protocols . . . 76
4.6.3 Friendly Jamming for Confidentiality . . . 78
4.6.3.1 Formalization of Friendly Jamming . . . 79
4.6.3.2 IND-CSA-2 Attacker . . . 80
4.7 Related Work . . . 81
4.8 Conclusion . . . 83

5 One Antenna to Recover Them All 85
5.1 Introduction . . . 86
5.1.1 Contribution . . . 87
5.1.2 Organization of This Chapter . . . 87
5.2 Background . . . 88
5.2.1 System and Attacker Model . . . 88
5.2.2 Sub-GHz IEEE 802.15.4 . . . 88
5.2.2.1 Binary Phase-Shift Keying . . . 89
5.2.2.2 Simultaneous Transmissions . . . 89
5.2.3 Evaluation Metrics . . . 89
5.3 Related Work . . . 90
5.3.1 Foundations and Implementations of Friendly Jamming . . . 90
5.3.2 Comparing Our Method with Others . . . 91
5.4 Attack Description . . . 93
5.4.1 Attack Scenario and Objectives . . . 93
5.4.2 Attack Procedure . . . 93
5.4.2.1 Our Attack in a Nutshell . . . 94
5.4.2.2 Illustration . . . 94
5.4.3 Detailed Attack Description . . . 94
5.5 Implementation . . . 99
5.5.1 Simulation Experiment . . . 100
5.5.1.1 Message Bit Jamming . . . 100
5.5.1.2 Chip Bit Jamming . . . 101
5.5.2 Real-World Experiment . . . 101
5.5.2.1 Test Setup . . . 101
5.5.2.2 GNU Radio Implementation . . . 103
5.5.2.3 Recovery Evaluation . . . 105
5.5.3 Attack Specificities . . . 107
5.6 Results . . . 107
5.6.1 Simulation Attack Results . . . 107
5.6.2 Real-World Attack Results . . . 108
5.7 Discussion . . . 109
5.7.1 Generalizability and Practicality of Our Attack . . . 110
5.7.2 Influence of Communication Distance . . . 111

5.7.3 Misalignments 112
5.7.3.1 Chip Misalignment 112
5.7.3.2 Phase and Frequency Misalignment 113
5.7.4 Differing Thresholds 113
5.7.5 Countermeasures 113
5.8 Conclusion 114

6 Conclusion **115**

Bibliography **119**

List of Figures

2.1 Floor plans of the three experiment spaces including sensor node locations . . 18
2.2 The four defined tasks performed by participants at Location A 19
2.3 Visualization of two examples of room climate measurements 23
2.4 Classification accuracy for occupancy detection and activity recognition . . . 26
2.5 Average accuracy over all sensors from each location for occupancy detection 29

3.1 Zigbee's security network models 38
3.2 Touchlink commissioning protocol 40
3.3 The four evaluated Zigbee-certified connected lighting systems 43
3.4 Acknowledgment spoofing 45
3.5 Outdoor testbed for attacking Zigbee-certified devices 49

4.1 The PPDU structure according to IEEE 802.15.4. 60
4.2 Linkage from $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}$ to $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}$ 72
4.3 A challenge-response protocol run between verifier V and prover P 79

5.1 Jamming plot from the simulation experiment 95
5.2 Waveform plot of three jammed messages from the real-world experiment . . 95
5.3 Detailed waveform plot of a jammed message from the real-world experiment 96
5.4 Flowchart illustrating the processing steps for recovering jammed bits 97
5.5 Chip Bit Jamming plot from the simulation experiment 101
5.6 Lab setup for the real-world experiment 102
5.7 The developed GNU Radio flowgraph comprising data sender and jammer . . . 104
5.8 The developed GNU Radio flowgraph for the attacker 106

List of Tables

2.1 Overview of previous experiments on occupancy detection, occupancy estimation, and activity recognition . . . 13
2.2 Measured sensor data of all three locations . . . 20
2.3 Demographic data of participants per location . . . 21
2.4 Classification accuracy for occupancy detection . . . 24
2.5 Occupancy detection and estimation accuracies for two- and three-class tasks at Locations A and B . . . 25
2.6 Occupancy detection and estimation accuracies for two- and three-class tasks at Locations A and B based on two sensors . . . 27
2.7 Activity Recognition accuracies for all three locations based on a single sensor and on two sensors . . . 28

3.1 Overview of attacks . . . 42
3.2 Maximum duration of the identify action attack for each evaluated lighting system . . . 46
3.3 Maximum ranges of an active attack in our evaluation for each evaluated lighting system . . . 50

5.1 Comparison of research studies on data recovery under adversarial and friendly jamming scenarios . . . 92
5.2 Overview of the scaling factors chosen in the respective attack step for simulation and practical evaluation . . . 108
5.3 Results of our simulation-based attack experiment for Message Bit Jamming . 109
5.4 Results of our simulation-based attack experiment for Chip Bit Jamming . . 110
5.5 Comparison of mean, minimum, and maximum BRR per antenna type and SJR level in the real-world experiment . . . 111
5.6 Comparison of mean, minimum, and maximum BRR per SJR level with communication distance of 10 m and the omnidirectional antenna . . . 112

List of Acronyms

AAJ Active Anti-Jamming 90

ACK MAC Ackowledgment 45

AES Advanced Encryption Standard 4, 38, 47

BER Bit Error Rate 89–91

BPSK Binary Phase-Shift Keying 59, 62, 85, 87–89, 93, 94, 101, 103, 105, 113, 117

BRR Bit Recovery Rate 6, 85, 87, 89, 91, 108–110, 112–114

CRC Cyclic Redundancy Check 60, 61, 75, 105

CSS Chirp Spread Spectrum 59, 103

DB Distance Bounding 5, 57–59, 64, 75, 76, 78, 82, 83, 116

DoS Denial-of-Service 5, 7, 42, 44

DPIA Data Protection Impact Assessment 32

DSSS Direct-Sequence Spread Spectrum 63, 88, 103, 105

FCS Frame Check Sequence 60, 105

FHSS Frequency-Hopping Spread Spectrum 5, 59, 63, 75, 82, 116

FSK Frequency-Shift Keying 64, 87

GDPR General Data Protection Regulation 31, 32

HVAC Heating, Ventilation, and Air-Conditioning 14, 15

IMD Implantable Medical Device 58, 62, 78, 79, 83, 86

IND-CPA Indistinguishability under Chosen Plaintext Attack 80

IND-CSA Indistinguishability under Chosen Signal Attack . . . . . iii, v, 5, 57, 80, 83, 86

IoT Internet of Things . . . . . iii, v, 1–4, 10, 32, 35, 36, 54, 55, 57–59, 90, 116

ISM Industrial, Scientific, and Medical 44

ITM Interactive Turing Machine 62, 65, 82

LR-WPAN Low-Rate Wireless Personal Area Network 59

MAC Media Access Control Layer 38, 39, 45, 54, 59, 60, 103

MFR MAC Footer 60

MHR MAC Header 60, 61

MitMM Man-in-the-Middle Model 5, 57–59, 62–66, 70, 71, 73–75, 81, 85

NB-IoT NarrowBand Internet of Things 4, 59

NDA Non-Disclosure Agreement 36, 39, 48, 53

OQPSK Offset QPSK 59, 60, 103

OTA Over-The-Air . . . . . iii, 6, 19, 54, 68, 85, 87, 88, 114

PAN Personal Area Network 36, 42, 44, 46–48, 50, 53

PDR Packet Delivery Ratio 90, 91

PHR PHY Header 60, 100, 103, 105

PHY Physical Layer 38, 59, 60, 65, 67, 87, 88, 90, 101, 103, 110

PIR Passive Infrared 12, 14, 15

PLR Packet Loss Rate 90, 91

PPDU PHY Protocol Data Unit 60, 61, 88, 100, 103, 105

PSDU PHY Service Data Unit 100, 103, 105

PSK Phase-Shift Keying 64, 68, 111

QPSK Quadrature Phase-Shift Keying 57, 59, 62, 68, 69, 75, 86, 87, 89

RFID Radio-Frequency Identification 32, 82

RSSI Received Signal Strength Indicator 64, 65

SDR Software-Defined Radio 43, 101, 103, 105

SHR Synchronization Header 60, 61, 103, 105

SJR Signal-to-Jamming Ratio 86, 90, 91, 100, 107–110, 112, 114

SNR Signal-to-Noise Ratio 90, 100

USRP Universal Software Radio Peripheral 43, 44, 49, 50, 101, 102

WAM Wireless Attacker Model . . . . . iii, v, 5, 57, 59, 73, 75, 85, 116

WCM Wireless Communication Model . . . iii, v, 5, 57, 59, 62, 65, 70, 73–75, 81, 85, 116

WLAN Wireless LAN 37, 41, 43, 54, 59, 63, 66

WPAN Wireless Personal Area Network 37, 63, 87, 88

ZLL Zigbee Light Link 36, 37, 39–41, 43, 44, 46, 47, 50–53, 59, 65

Am Anfang war das Am

(Knorkator)

1 Introduction

Abstract

This chapter provides a brief introduction to the Internet of Things (IoT). Subsequently, we highlight the contribution, contents, and organization of this thesis.

1.1 Background and Motivation

The term IoT most probably first appeared in a speech by Peter T. Lewis in 1985 [Sha16] who described it as "the integration of people, processes and technology with connectable devices and sensors to enable remote monitoring, status, manipulation and evaluation of trends of such devices." [Sha16] Independently, Kevin Ashton devised the term using it as title for a presentation he gave in 1999 [Ash09]. Clarifying his vision of the IoT roughly ten years later, he concludes that "[if] we had computers that knew everything there was to know about things [...] we would be able to track and count everything" [Ash09]. Furthermore, he advocates that such computers should be able to "see, hear and smell the world for themselves" [Ash09]. In essence, the two interpretations are closely related, they consider sensing the physical world and having *remote* monitoring capabilities among their use cases. Although Ashton does not explicitly state the requirement of having Internet communication capabilities, he implies it as he does not differentiate between *things* and *computers*, while he considers the latter an integral part of the Internet [Ash09].

Nowadays, this IoT vision has become reality. The world is ever more interconnected, e. g., the number of fixed broadband Internet subscriptions worldwide has slightly more than doubled from approximately 0.7 million in 2014 to approximately 1.5 million in 2023 [ITU23]. Connections of IoT devices worldwide are forecast to increase from 9.5 *billion* in 2023 to 19 billion by 2028 in the enterprise sector respectively from 9 billion in 2023 to 13 *billion* in the consumer sector [Eri23] with cellular IoT connections predicted to rise up to approximately 3.3 million (using 4G or 5G technology) by 2028 as well [GSM23]. This encompasses devices like smart door locks or smart light bulbs which are able to interact with their physical surroundings while being connected to the Internet. Consequently, the IoT has established itself as an integral part of everyday life, e. g., in the form of wearables with 506.6 million units shipped in 2023 (and forecast 612.5 million in 2028) [IDC24], or smart homes with approximately 1.361 billion users in 2023 (and forecast approximately 1.889 billion in 2028) [Sta25]. As the private home is held to high privacy standards and, e. g., in Europe, protected by European law from unwarranted intrusion and surveillance (e. g., see Article 7 of the Charter of Fundamental Rights of the European Union), having devices or sensors communicating and sharing information about the state of this private sphere might lead to leaking information that is commonly recommended to be kept private [ALAM19; CAF19; Cho18; WAF17].

In the following, we provide the motivation and contribution of this thesis including the underlying publications as well as how they relate to one another, and the organization of this thesis.

1.2 Need for Protection of Room Climate Data

Considering a smart home scenario with lightweight devices communicating wirelessly, one might raise the question whether it would be necessary to protect each and every kind of data originating from within this home. This contrasts with the traditional IT world, in which most of the communication over the Internet is protected by Transport Layer Security nowadays [Goo26], regardless of type of data. In the smart home or general IoT landscape, this is not as ubiquitous [JSJ+25; Pal20; Zsc19]. Although IoT devices are also becoming more powerful in general, they might lack update capabilities and they are usually desgined

for specific purposes. Furthermore, there might be situations where the data originating from an IoT device is considered to not contain anything private.

Motivation. Consider a heating installation that keeps track of the rooms' climate by measuring temperature and (relative) humidity in regular intervals with the goal to satisfy user-defined comfort levels. Recording these two physical quantities only yields data points in time with (supposedly) minimal information, i. e., the currently measured temperature and (relative) humidity at that specific point in time. Considering the fact that humans breathe and dissipate heat [PKS23, Chapter 13; Hal15, Chapter 74], it is conceivable that such data could reveal the mere presence of a person in a room. If this kind of detection can be performed reliably, then this would allow for creating behavioral profiles, e. g., at what time and for how long a person is in a specific room, when the person usually sleeps or leaves the house, and so on. While this could technically also be (partially) observed by someone in physical proximity, using the collected data for deducing whether this person is reading a book or working on a laptop, or how many persons are present might seem more intricate and maybe even too far-fetched. Of course, reading or working might be perceived as harmless activities. Nonetheless, if these two somewhat similar activities can be distinguished from one another, a natural follow-up question is what other *intimate* or potentially incriminating activities could be disclosed involuntarily by properly analyzing the data. This shows that, even if the data seems innocuous at first glance, depending on accompanying circumstances, the data might indeed reveal private information — especially if this data is transmitted without any data protection measures to remote or third parties.

Contribution. In Chapter 2, we present our papers [ABM+19; MMR+17] in which we show that, using only two basic physical quantities, i. e., temperature and relative humidity, an attacker is able to infer presence (including the number of people) and to discriminate between four different types of activities. More specifically, we use room climate data collected at three different locations and train machine learning classifiers using data from one or two sensor nodes, each equipped with a combined humdity-temperature sensor, placed in different positions in the respective location. The results show that single-sensor occupancy detection is reliably possible with an accuracy of up to 93.5 %, while we achieve an accuracy of up to 72.2 % for single-sensor occupancy estimation considering the cases of no, one, and two persons present. For single-sensor activity detection, i. e., identifying which of the four defined tasks — *reading*, *standing*, *walking*, and *working on a laptop* — is performed, our approach yields an accuracy of up to 56.8 %. Further analysis reveals that considering fewer tasks improves the accuracy, e. g., *standing* and *walking* can be discriminated with an accuracy of up to 96.3 % using a single sensor. When feeding the machine learning classifiers with data from a second sensor residing in the same location, in most cases, the accuracy increases.

1.3 Vulnerabilities of Wireless Communication Protocols

The results above demonstrate that protecting even such rudimentary physical quantities is essential for privacy protection of smart home occupants. Thus, they represent a strong point for using established protocols and encryption schemes as potential safeguarding measures, and they ultimately answer the question whether data like temperature and relative humidity

should be protected from unwarranted access. A partial solution to protect data originating from lightweight devices could be to use a standardized wireless communication protocol offering data protection measures like data encryption, such that communicating devices use established encryption mechanisms to achieve confidentiality.

Motivation. As there are many competing standards targeting similar types of devices for different communication ranges, e. g., NarrowBand Internet of Things (NB-IoT) [3GP22a], LoRa and LoRaWAN [LoR20; LoR22], WirelessHART [IEC16], IEEE 802.15.4 [IEE11], Zigbee [zig17], 6LoWPAN [RFC4944], or Bluetooth Low Energy [Blu24], choosing a secure protocol is crucial. However, not all developed protocols may have received the same scrutiny. That is, they might contain exploitable vulnerabilities that, on the one hand, might render the protection stemming from encryption useless, or, on the other hand, prevent the rightful owner from interacting with the device in the intended way. Such vulnerabilities could be manifold, e. g., a symmetric key used by all nodes in the network might be derived or distributed insecurely when a new device is added, or it might be assumed that a *strong* signal implies the sender's proximity. Thus, careful security analyses of communication standards are of utmost importance for ensuring data is protected from illegitimate parties.

A popular and public communication standard for IoT devices is *Zigbee* [zig17]. It offers definitions for a plethora of smart devices with specific sets of instructions or mechanisms for network management, e. g., adding and removing new devices. Communication is encrypted using symmetric encryption based on the Advanced Encryption Standard (AES). For smart light bulbs, Zigbee has an optional method called *touchlink commissioning*, which is exclusive to light bulbs, for adding new devices to an existing Zigbee network. It specifically targets easy integration. Consequently, a question is how well-thought-out the specification is for this use case and whether it contains exploitable weaknesses ultimately deviating from the targeted level of security.

Contribution. Chapter 3 presents our paper [MMB+17] in which we analyze Zigbee's touchlink commissioning procedure and uncover several vulnerabilities an attacker can exploit, e. g., to take over devices. As some of the discovered vulnerabilities are rooted in design choices and defined in the Zigbee specification [zig17], these cannot be fixed without breaking backwards compatibility. Using our developed framework Z3sec and popular Zigbee-certified light bulbs, we demonstrate passive key extraction from a distance of 130 m and active device takeover from 190 m among others.

1.4 Theoretical Models Incorporating Real-World Effects

An underlying cause of the identified problems within Zigbee is the assumption that proximity is directly determined by the signal strength of wirelessly received communication. One important question is why such assumptions are made in critical parts of a standard that ultimately lead to weakened security.

Motivation. Clearly, certain aspects of wireless communication and potential attacker capabilities were not properly considered during the design phase of the Zigbee specification. A root cause could be that existing communication and attacker models from the wired world impose unrealistic restrictions, especially for lightweight, wireless devices. Or in other words,

models that are applicable to *real-world* wireless communication are either not widespread or inadequate. As a result, this misconception may lead to incorrect conclusions, such as assuming that a party is closer than they actually are.

Considering available (attacker) models commonly used for multi-party scenarios, e. g., the Man-in-the-Middle Model (MitMM), it is obvious they fall short of respecting specifics unique to wireless communication. They inherently require to consider an unrealistically strong attacker one needs to protect against, e. g., when designing a wireless communication protocol. Thus, such a model might be considered inapplicable due to its imposed requirements or its lacking representation of real-world effects or of physical properties in wireless communication. As wireless communication continues to proliferate, the need for adequate formal models therefore becomes even more important.

Contribution. Consequently, in our papers [AM24; AM25], presented in an extended version in Chapter 4, we introduce a new communication model and a new attacker model tailored to the peculiarities of wireless communication, dubbed Wireless Communication Model (WCM) and Wireless Attacker Model (WAM), respectively. There, we capture unique aspects such as communication channels, signals, and locality to adequately represent distinguishing features of wireless communication. More specifically, we question the generally accepted abilities of an attacker regarding the control of communication in a wireless setting. That is, we highlight the shortcomings of the MitMM in the context of wireless communication, and, at the same time, propose a concrete solution, i. e., WCM and WAM, respecting the fundamental properties of wireless communication in the real-world. Additionally, we demonstrate the benefit of our suggested model by showcasing how to formalize and analyze established mechanisms such as Frequency-Hopping Spread Spectrum (FHSS), Distance Bounding (DB) protocols, or friendly jamming. We also develop the new security notion IND-CSA (Indistinguishability under Chosen Signal Attack).

1.5 Lowering Requirements for Recovering Jammed Bits

The discovery of vulnerabilities in a communication protocol triggers the immediate question whether these can be mitigated by changing only a device's implementation or whether both the respective protocol and the devices implementing and utilizing this communication protocol, need to be changed. Depending on the exact nature of a vulnerability, mitigation might introduce a breaking change, i. e., the protocol is changed in a way that makes it incompatible with previous iterations of the protocol. Besides the difficulties of distributing upgrades to all devices in the field, some devices might not even have the capability of being updated. From a security perspective, as a consequence, such devices should be decommissioned unless there might be other options to protect these devices from exploitation, e. g., *friendly jamming* where an additional guarding device could block certain messages by listening for and jamming these messages, thus, preventing such messages from reaching the vulnerable device. This form of protective jamming complements the better-known *adversarial jamming*.

Motivation. While adversarial jamming represents a Denial-of-Service (DoS) attack as the attacker prevents the reception of transmitted information, friendly jamming targets authorization as in the example above, or, which is more common, confidentiality as the jammer

protects the information being transmitted, and the attacker tries to recover the original information. The difficulty here is to choose jamming signals that are effective, as these must be aligned to the protected transmission in form *and* time, and their effect must be reversible for the legitimate receiver. Ultimately, the attacker should not be able to recover a meaningful portion of the received communication.

One observation from the formalization described above is that the effect of jamming depends on the underlying modulation scheme among others. Another is that, under certain circumstances, jammed signals could be recovered. For instance, by utilizing two antennas for receiving, the successful recovery of friendly-jammed wireless communication is possible [TMRC13]. Inspired by the widely accepted assumption that effective signal recovery by attackers requires at least as many antennas as are involved for sending and jamming, we challenge this conventional wisdom by showing that effective signal recovery can be achieved by an attacker using *only one antenna*.

Contribution. Chapter 5 presents an extended version of our paper [MMSA25], where we introduce a novel attack method that effectively neutralizes the advantages of (friendly) jamming. By exploiting properties of the underlying modulation to estimate signal parameters, our attack recovers the original bits from jammed signals. To this end, we conducted simulation and real-world experiments. Our results are very promising and highlight the potential of simpler hardware for robust bit recovery, lowering costs and technical barriers for the attacker. More precisely, in the simulation, we achieve Bit Recovery Rates (BRRs) that are close to 1, in our real-world Over-The-Air (OTA) tests with a 15 dB stronger jamming signal, the average BRR is 0.70 with a maximum BRR of 0.92. With a jamming signal that is 20 dB stronger, the average BRR decreases to 0.34 with a maximum BRR of 0.76, still allowing for significant signal recovery even with substantial jamming. Results of an additional follow-up experiment with a communication distance of 10 m underscore our attack's feasibility in the real-world with an average BRR of 0.91 and 0.89 with a 15 dB and a 20 dB stronger jammer, respectively. Clearly, these results show a need to re-evaluate wireless security frameworks, especially against single-antenna attacks. Our code and data are publicly available to enable reproducibility and transparency.

1.6 Publications

The work presented in this thesis is based on the following publications.

[MMSA25]: *One Antenna is Enough for Neutralizing Jamming in IEEE 802.15.4 BPSK*, written together with *Vasily Mikhalev, Yves T. Staudenmaier, and Frederik Armknecht*, and presented at *IWSEC 2025*. The author of this thesis designed, tested, and evaluated the theoretical and practical experiments, as well as contributed major parts of the publication as main author. Together with the co-authors, some of the practical experiments were carried out, and the attack procedure and the paper's content were iteratively refined.

[AM24; AM25]: *Towards a Cryptographic Model for Wireless Communication: Extended Version*, written together with *Frederik Armknecht*, and published in *Security and Cryptography* which is an extended version of *Towards a Cryptographic Model for Wireless Communication*, written together with *Frederik Armknecht*, and presented at *SECRYPT 2024*. These

publications are main-authored by the author of this thesis who contributed the main content including the technical background. Frederik Armknecht helped with the formalization, developing the security proof, and further content refinements.

[MMB+17]: *Insecure to the touch: attacking ZigBee 3.0 via touchlink commissioning*, written together with *Philipp Morgner, Stephan Mattejat, Zinaida Benenson, and Frederik Armknecht*, and presented at *WiSec '17*. With this publication, a Zigbee penetration testing framework was released. The framework is a result of Stephan Mattejat's Bachelor's Thesis [Mat16]. While the author of this thesis discovered the possibility of DoS attacks inherent to the relevant specification and a (so far theoretical) possibility for exploitation, the practical evaluation was conducted by Stephan Mattejat and Philipp Morgner. Additionally, the author of this thesis contributed crucial parts to the publication, and all co-authors provided refinements iteratively.

[ABM+19; MMR+17]: *Privacy implications of room climate data*, written together with *Frederik Armknecht, Zinaida Benenson, Philipp Morgner, and Christian Riess*, and published in the *Journal of Computer Security* which is an extended version of *Privacy Implications of Room Climate Data*, written together with *Philipp Morgner, Matthias Ring, Björn Eskofier, Christian Riess, Frederik Armknecht, and Zinaida Benenson*, and presented at *ESORICS 2017*. Together with Philipp Morgner and Christian Riess, the two-sensor classification and occupancy estimation approaches, reported exclusively in the extended version, were devised, refined, and integrated into the publication. The author of this thesis considerably improved the initial experimental design developed by Philipp Morgner, managed and conducted the data collection at two of the three locations, and contributed significant content to the publications. In cooperation with all involved co-authors, the structure and content of the paper were iteratively and collectively improved.

The author of this thesis additionally contributed to the following publications, however, they are not presented in this thesis.

[SMA22]: *If You Like Me, Please Don't "Like" Me: Inferring Vendor Bitcoin Addresses From Positive Reviews*, written together with *Jochen Schäfer, and Frederik Armknecht*, and presented at *PETS 2022*. In close collaboration with Jochen Schäfer, the author of this thesis contributed important parts to the general content, e. g., the background or the attack's theoretical foundation and presentation, as well as the development and refinement of the formal concepts. Together with Jochen Schäfer and Frederik Armknecht, the paper's structure and content were iteratively improved.

[MAM17]: *On Ciphers that Continuously Access the Non-Volatile Key*, written together with *Vasily Mikhalev, and Frederik Armknecht*, and presented at *FSE 2017*. Besides enhancing the overall presentation, the publication's content regarding different EEPROM access possibilities and their respective influence on the considered cipher implementations respectively their throughput was main-authored by the author of this thesis. All parts of the publication received general refinements and were collaboratively improved with Vasily Mikhalev and Frederik Armknecht.

Further Publications. Additional publications to which the author of this thesis contributed, but which are not further elaborated on are the following.

[MSMA25]: *On The Recovery Of Jammed Bits*, written together with *Yves T. Staudenmaier, Vasily Mikhalev, and Frederik Armknecht*, and presented at *38th Crypto Day*.

[SMA24]: *0 Sterne für die Sicherheit: Wie Kundenbewertungen die Bitcoin-Adressen von Darknet-Händlern verraten*, written together with *Jochen Schäfer, and Frederik Armknecht*, and presented at *Sicherheit 2024*.

[MABM16]: *On the security of the ZigBee light link touchlink commissioning procedure*, written together with *Frederik Armknecht, Zinaida Benenson, and Philipp Morgner*, and presented at *SmartBuildingSec'16* (co-located with *Sicherheit 2016*).

[MBMA15]: *Design Space of Smart Home Networks from a Security Perspective*, written together with *Philipp Morgner, Zinaida Benenson, and Frederik Armknecht*, and presented at *Fachgespräch Sensornetze 2015*.

1.7 Organization of This Thesis

The remainder of this thesis is organized as follows. Chapter 2 investigates potential privacy risks stemming from room climate measurement data. In Chapter 3, we present our security analysis of Zigbee's touchlink commissioning procedure which uncovers several exploitable vulnerabilities. Chapter 4 presents our cryptographic model for wireless communication incorporating necessary real-world effects. In Chapter 5, we show the feasibility of recovering jammed bits using only a single antenna. Chapter 6 concludes this thesis with a summary of our results.

Asleep or awake, working or eating, indoors or out of doors, in the bath or in bed—no escape.

Nineteen Eighty-Four (George Orwell)

2 Keep Room Climate Data Secret — Keep Room Climate Data Safe

Abstract

Smart heating applications promise to increase energy efficiency and comfort by collecting and processing room climate data. While it has been suspected that the sensed data may leak crucial personal information about the occupants, this belief has not been supported by evidence, until now.

In this chapter, we investigate privacy risks arising from the collection of room climate measurements. We assume that an attacker has access to the most basic measurements only: temperature and relative humidity. We train machine learning classifiers to predict the presence and number of room occupants and to distinguish between different types of activities. On data that was collected at three different locations, we show that occupancy can be detected from data measured by a single sensor with up to 93.5 % accuracy. One can even distinguish between the cases that no, one, or two persons are present with up to 72.2 % accuracy. Moreover, four specific actions — reading, standing, walking, and working on a laptop — can be distinguished with up to 56.8 % accuracy, which is likewise clearly better than guessing (25 %). Constraining the set of actions allows for even higher accuracies. For example, we distinguish standing from walking occupants with 96.3 % accuracy. In addition, we show that the accuracy can be increased in most cases if an attacker has access to measurements from two different sensors located in the same room.

Our results provide evidence that even the leakage of such 'inconspicuous' data as temperature and relative humidity can seriously violate privacy. Thus, even these basic types of data should be protected from illegitimate access whenever possible.

Declaration of Origin: This chapter is based on the paper *Privacy implications of room climate data* [ABM+19], written together with *Frederik Armknecht, Zinaida Benenson, Philipp Morgner, and Christian Riess* which is an extended journal version of *Privacy Implications of Room Climate Data* [MMR+17], written together with *Philipp Morgner, Matthias Ring, Björn Eskofier, Christian Riess, Frederik Armknecht, and Zinaida Benenson*, presented at *ESORICS 2017*.

2.1 Introduction

The vision of the IoT is to enhance work processes, energy efficiency, and living comfort by interconnecting actuators, mobile devices, and sensors. These networks of embedded technologies enable applications such as smart heating, home automation, and smart metering, among many others. Sensors are of crucial importance in these applications. Data gathered by sensors is used to represent the current state of the environment, e. g., in smart heating, sensors measure the room climate. Using this information and a user-defined configuration of the targeted room climate, the application regulates heating, ventilation, and air conditioning.

While the collection of room climate data is obviously essential to enable smart heating, at the same time, it may pose a risk of privacy violations. Consequently, it is commonly *believed* among security experts that leaking room climate data may result in privacy violations and hence that the data needs to be cryptographically protected [CCC16; Eur15; Sel14]. However, these claims have not been supported by scientific *evidence* so far. Thus, one could question whether additional effort for protecting the data would be justified in practice.

The current situation with room climate data is comparable to the area of smart metering [Har89; JJK11; KDK11; YLQ+12]. In 1989, Hart [Har89] was the first to draw attention to the fact that smart metering appliances can be exploited as surveillance devices. Since then, research has shown far-reaching privacy violations through fine-grained power consumption monitoring, ranging from occupancy and everyday activities detection [MSF+10] up to recognizing which program a TV was showing [GGJL12].

Various techniques have been proposed over the years to mitigate privacy risks of smart metering [CPW10; JKD12; RD11; REC15; YLQ+12]. This issue has become such a grave concern that the German Federal Office for Information Security (*German*: Bundesamt für Sicherheit in der Informationstechnik, BSI) published a protection profile for smart meters in 2014 [BSI14]. By considering privacy implications of smart heating, we hope to initiate consumer protection research and policy debate in this area, analogous to the developments in smart metering described above.

2.1.1 Research Questions

In this chapter, we investigate room climate data from the perspective of possible privacy violations. More precisely, we address the following research questions:

- *Occupancy detection:* Can an attacker determine the presence of a person in a room using only room climate data, i. e., temperature and relative humidity?
- *Occupancy estimation:* Can an attacker determine the number of persons present in a room using only this room climate data?
- *Activity recognition:* Can an attacker recognize the occupants' activities in the room using only the temperature and relative humidity data?

Our threat scenario targets buildings with multiple rooms that are similar in size, layout, furnishing, and sensor placement. These properties are typical for office buildings, dormitories, cruise ships, and hotels, among others. Assuming that an attacker is able to train a classifier that recognizes pre-defined activities, possible privacy violations are, e. g., tracking

presence and working practices of employees in offices, or the disclosure of lifestyle and intimate activities in private spaces. All these situations present intrusions into the occupants' privacy. In contrast to surveillance cameras and motion sensors, occupants are not expecting to be monitored. Also, legal restrictions regarding privacy might apply to surveillance cameras and motion sensors, but not to room climate sensors.

2.1.2 Experiments

To evaluate these threats, we present experiments that consider occupancy detection, occupancy estimation, and activity recognition based on the analysis of room climate data from a privacy perspective. We measured room climate data in three office-like rooms and distinguished between the four activities *reading*, *standing*, *walking*, and *working on a laptop*. The data was collected from sensors measuring temperature and relative humidity at regular intervals of a few seconds.

To reflect the most restrictive scenario, we analyzed how much information can be derived from the measurements of a single sensor. Although we assume that in smart heating applications, only one sensor per room is most likely to be installed, each room was equipped with several sensors in order to evaluate any impact of the position of the sensor in the room. Additionally, we combined and evaluated data from two sensors simultaneously to investigate whether increasing the available data affects classification noticeably.

In our procedure, occupants performed a pre-defined sequence of tasks in the experimental space. In sum, we collected almost 115 hours of room climate sensor data from a total of 36 participants. The collected room climate data was analyzed using an off-the-shelf machine learning classification algorithm.

2.1.3 Summary of Results

Evaluating our collected room climate data, the attacker detects the presence of a person with detection rates of up to 93.5 % depending on location and the sensor position, which is significantly higher than random guessing (50 %). Discriminating between one and two persons present in the room is not a similarly easy task for the attacker as classification results do not exceed 68 %, with some around 48 %, which is in the region of guessing (50 %). Consequently, differentiating between absence, presence of one, and of two persons yields results lying between those of the former two tasks with rates of up to 72.2 %. The attacker can distinguish between four activities (reading, standing, walking, and working on a laptop) with detection rates of up to 56.8 %, which is significantly better than guessing (25 %). We can also distinguish between three activities (sitting, standing, and walking) with detection rates of up to 81.0 %, as opposed to 33.3 % if guessing. Furthermore, we distinguish between standing and walking with detection rates of up to 96.3 %. Additionally, evaluating the performance of pairs of sensors predominantly shows an increase of detection rates. However, the exact difference fluctuates from almost none to approximately 10 %. Thus, we show that the fears of privacy violation by leaking room climate data are well-justified. Furthermore, we analyze the influence of the room size, sensor positions, and amount of the measured sensor data on the accuracy. In summary, we take the first steps towards verifying the common belief that room climate data leaks privacy-sensitive information.

2.1.4 Organization of This Chapter

The remainder of this chapter is organized as follows. In Section 2.2, we give an overview of related work. Section 2.3 presents the considered threat model. In Section 2.4, we introduce the experimental design and methods. The results of our experiments and complementary observations are presented in Sections 2.5 and 2.6, respectively, and discussed in Section 2.7. This chapter concludes with Section 2.8.

2.2 Related Work

Over the last decade, several experiments have been conducted to detect occupancy in sensor-equipped spaces and to recognize people's activities as summarized in Table 2.1. Activity recognition has been considered for basic activities, such as leaving and arriving at home, or sleeping [LSS+10], as well as for more detailed views, including toileting, showering, and eating [KNEK08].

Most of the previous research uses types of sensors that are different from temperature and relative humidity. For example, CO_2 represents a useful source for occupancy detection and estimation [WBRW14]. Additionally, sensors detecting motion based on Passive Infrared (PIR) [AFG14; CFG+16; DAL+10; FXL+17; HGAK11; HGF12; LHD+09; PNP17; SCFS16; ZLCD12], sound [EBPF13; GRR+17; HGAK11; PNP17], barometric pressure [MSC15], and door switches [CFG+16; EBV+13; EBV+15; FXL+17; MPC16; SCFS16; YLBO14] are utilized for occupancy estimation. For evaluation, different machine learning techniques are used, e. g., hidden Markov models [DAL+10; FXL+17; KNEK08; LHD+09; LSS+10; WBRW14], auto-regressive hidden Markov models [AFG14; HGF12], Artificial Neural Networks [DAL+10; EBPF13; LHD+09; MPC16], Naïve Bayes [ZWF18], and decision trees [CF16; FXL+17; GRR+17; HGAK11; YLBO14; ZWF18].

2.2.1 Occupancy Detection

Hailemariam et al. [HGAK11] present a system aiming at occupancy detection. They conclude that using *only* PIR-based sensors delivers the most reliable results for their cubicle office setting. Using the information collected by a Netatmo weather station with a five-minute granularity, Wörner et al. [WBRW14] provide evidence that the station's CO_2 sensor data leaks information about a room's occupancy. Detecting the presence of one or two persons in a two-person office is presented by Candanedo and Feldheim [CF16]. The authors used temperature, humidity, light, and CO_2 sensors, whereas they achieved the best results when combining data from only two of those, concluding that any combination including the lighting sensor, and likewise, the time of day help overall performance. A single-person office and a three-room dormitory (shared by two persons), each equipped with temperature, humidity, CO_2, volatile organic compound, PIR, and acoustics sensors, are used for occupancy detection by Pedersen et al. [PNP17]. As seen before, CO_2 combined with PIR is a suitable solution for occupancy detection of a room.

2.2.2 Occupancy Estimation

Detecting occupancy levels ranging from zero to three in an open-plan office setting is presented by Lam et al. [LHD+09]. To this end, the authors use several different sensor types,

Table 2.1: Overview of previous experiments on occupancy detection (D), occupancy estimation (E), which aims at determining the number of people in a room, and activity recognition (A) with a focus on selected sensors; Rel. Hum. refers to relative humidity, Temp. to temperature, Vent. to ventilation, and AML denotes acoustic, motion, and lighting sensors.

Work	Target	Rel. Hum.	Temp.	CO_2	Vent.	AML	Switches	Wearables
Van Kasteren et al., 2008 [KNEK08]	A	○	○	○	○	○	●	○
Lu et al., 2010 [LSS+10]	A	○	○	○	○	●	●	○
Mehr et al., 2016 [MPC16]	A	○	○	○	○	○	●	○
Sprint et al., 2016 [SCFS16]	A	○	●	○	○	●	●	○
Cicirelli et al., 2016 [CFG+16]	A	○	○	○	○	●	●	●
Fan et al., 2017 [FXL+17]	A	●	●	○	●	●	●	●
Hailemariam et al., 2011 [HGAK11]	D	○	○	○	○	●	○	○
Wörner et al., 2014 [WBRW14]	D	○	○	●	○	○	○	○
Candanedo and Feldheim, 2016 [CF16]	D	●	●	●	○	●	○	○
Pedersen et al., 2017 [PNP17]	D	●	●	●	○	●	○	○
Lam et al., 2009 [LHD+09]	E	○	○	●	○	●	○	○
Dong et al., 2010 [DAL+10]	E	●	●	●	○	●	○	○
Han et al., 2012 [HGF12]	E	●	●	●	●	●	○	○
Zhang et al., 2012 [ZLCD12]	E	●	●	●	○	●	○	○
Ekwevugbe et al., 2013 [EBPF13]	E	○	○	●	○	●	○	○
Ebadat et al., 2013 [EBV+13]	E	○	●	●	●	○	○	○
Ai et al., 2014 [AFG14]	E	●	●	●	●	●	●	○
Masood et al., 2015 [MSC15]	E	○	●	●	○	○	○	○
Ebadat et al., 2015 [EBV+15]	E	○	●	●	●	○	●	○
Yang et al., 2014 [YLBO14]	D/E	●	●	●	○	●	●	○
Ghaffarzadegan et al., 2017 [GRR+17]	D/E	○	○	○	○	●	○	○
Zimmermann et al., 2018 [ZWF18]	D/E	●	●	●	○	○	○	○
This work	D/E/A	●	●	○	○	○	○	○

i. e., CO, CO_2, total volatile organic compounds, small particulates, acoustics, illumination, motion, temperature, and relative humidity sensors, whereas the latter two are not contributing to their findings. The authors deem the CO_2 and acoustics information the most valuable. This work is reaffirmed and extended by Dong et al. [DAL+10] with occupancy levels ranging from zero to four. Zhang et al. [ZLCD12] use the same workplace and sensor node settings. Here, relative humidity and temperature contribute to their results. In an office setting, Han et al. [HGF12] present a system dedicated to detecting occupancy levels ranging from zero to six occupants while, at the same time, providing a measure of indoor air quality. They deploy sensor nodes comprising PIR, CO_2, air velocity, temperature, and relative humidity sensors. Continuing this work, the authors of [AFG14] re-used the previously developed system to increase its performance for occupancy estimation. Ekwevugbe et al. [EBPF13] deployed a sensor node network in an open-plan office. They evaluate the usefulness of using acoustics, (desktop PC case) temperature, CO_2, volatile organic compound, PIR, temperature, and relative humidity sensors. For their scenario, i. e., estimating occupancy levels ranging from zero to five, they focused on acoustics, (desktop PC case) temperature, CO_2, and PIR as the best combination. Ebadat et al. [EBV+13] present estimating occupancy levels from zero to four in an office setting, using Heating, Ventilation, and Air-Conditioning (HVAC) actuation, temperature, and CO_2 sensors. Their work is extended by Ebadat et al. [EBV+15]. They additionally use the information from deployed door switches. Yang et al. [YLBO14] evaluate the usefulness of various sensor types in an office building. For occupancy detection and estimation, they find CO_2, door status, and lighting levels the most valuable, thereby excluding the deployed temperature and relative humidity sensors. Masood et al. [MSC15] aim to estimate occupancy levels of a classroom for 35 people, with a maximum of 22 occupants being present simultaneously, utilizing CO_2, temperature, and barometric pressure sensors. Especially, the latter was deemed useful. However, they targeted five different occupancy levels, each spanning five values, i. e., 0–4, 5–9, and so forth. Another interesting approach is to derive (the level of) occupancy from acoustics sensors only in a residential and commercial setting as presented by Ghaffarzadegan et al. [GRR+17]. They use audio data from a living room with levels of occupancy from zero to three and a single-person office with occupancy ranging from zero to two. Using an indoor air quality system including CO_2, volatile organic compound, temperature, and relative humidity sensors to detect occupancy levels in four student dormitories is shown by Zimmermann et al. [ZWF18]. The occupancy levels ranged from zero to three, whereas they consider temperature and humidity far less valuable than CO_2.

2.2.3 Activity Recognition

Van Kasteren et al. [KNEK08] deployed 14 *state-change* nodes (or digital switches) that were installed in key positions in a residential apartment. For instance, one node was installed as a toilet flush sensor, which would notice and report toileting activity. Their goal was to automatically detect activities of daily living via a strategically deployed sensor/switch network. Similarly, Mehr et al. [MPC16] present a system to detect routine activities of daily living using data collected from 77 state-change sensors deployed in a residential apartment. The main goal of Lu et al. [LSS+10] is to achieve savings by autonomously regulating the heating of a residential HVAC system. By using information from deployed motion (PIR) sensors and door switches, a central unit makes a decision about the current occupancy state. They further use *historical* data to establish a baseline occupancy-sleep-absence pattern,

serving the purpose to heat the home before the occupants are expected to return home. Another use case for activity recognition is ambient assisted living. Sprint et al. [SCFS16] set out to detect health or behavior change in a residential apartment. For this, they use motion, lighting, temperature sensors, and door switches to detect activities of daily living. Likewise, Cicirelli et al. [CFG+16] developed a framework allowing the recognition of activities in a residential home. For better resolution, they complement PIR sensors and switches with wearables. Using temperature, humidity, small particulates, and PIR sensors, door switches, wearables, and HVAC events, Fan et al. [FXL+17] offer activity recognition for ambient assisted living as a service.

2.2.4 This Work

In contrast to previous work, our results rely exclusively on temperature[1] and relative humidity. Previously published experimental results involved other or additional types of sensors, such as CO_2, acoustics, motion, or lighting (the latter three are referred to as AML in Table 2.1), door switches or states of appliances (also gathered with the help of switches), such as water taps or WC flushes. For this reason, our detection results are not directly comparable to these works.

2.3 Threat Model

The overall aim of our work is to understand the potential privacy implications if room climate data is accessed by other parties. Consequently, we consider a threat model where the attacker is an outsider, i. e., a party that is not present in the room, that gains access to the room climate data. The goal of the attacker is to extract information from the data that would otherwise not be known to her. Obviously, the term "information" is subject to manifold interpretations. Moreover, the type and granularity of measurements are most likely to have an impact on what kind of information can be extracted. Therefore, we need to make more explicit what information an attacker is aiming for and what kind of measurements are at her disposal.

2.3.1 Attack Scenario and Goals

There are various scenarios in which an attacker has incentives to collect and analyze room climate data. For example, the management of a company wants to track the presence and working practices of employees in the offices. In another case, a provider of private spaces (hotels, dormitories, etc.) wants to reveal lifestyle and intimate activities in these spaces. This information may be utilized for targeted advertising or sold to insurance companies. In any case, the evaluation of room climate data provides the attacker with the possibility to undermine the occupants' privacy.

In our work, we focus on the question whether basic information are extractable. That is, we consider as attack goal to gain information about the *occupancy state*, i. e., if and how many persons are present in the room, as well as the *activity* of the occupants without

[1] Note that additional classification results using temperature as the only predictor are reported in the range of 67 % to 87 % by Candanedo and Feldheim [CF16]. The underlying distribution of absence (79 %, 64 %, and 79 %) and presence (21 %, 36 %, and 21 %) in their three datasets is unbalanced and, thus, may therefore have biased the classifier.

their consent. With respect to the available data, it is obvious that the more information an attacker can gather, the more likely she can deduce privacy-harming information from the measurements. However, in practice, only few sensors may be installed within one room, especially as one sensor node per room is sufficient to monitor the room climate, and these sensors may measure a very limited set of data only. Therefore, we base our analysis on an attacker model that considers a room climate system where only *few* sensor nodes are used and only *basic data* is measured.

That is, the main focus of our experiments is on the case that only *one* single sensor is present. If it turns out that a single sensor already provides sufficient data to derive privacy-harming information about the occupants, the situation may get even worse when more sensors are accessible. To examine this, we also consider the case that an attacker can make use of the measurements of *two* different sensors to analyze whether and by how much the attacker performs better compared to the single-sensor case.

In our threat model, we assume that a sensor node takes only the two most basic measurements, temperature and relative humidity. These data are the fundamental properties to describe room climate. Note that our restricted data is in contrast to existing work (cf. Table 2.1 and Section 2.2) that based their experiments on more types of measurements or used data that is less common to characterize room climate.

2.3.2 Attacker Model and Capabilities

We consider a sensor system that measures the climate of a room, denoted as *target location*. At the target location, temperature and relative humidity sensors are installed that report the measured values at regular intervals to a central database. We consider an attacker model where the attacker has access to this database and aims to derive information about the occupants at the target location. Furthermore, we assume that the attacker has access to either the target location itself, or rooms that are similar in size, layout, sensor positions, and furniture. Such situations are given, for example, in office buildings, hotels, cruise ships, and student dormitories. These locations, denoted as *training location*, are used to train the classifier, which is a machine learning algorithm learning the input data labeled with the ground truth. As the attacker has full control over the training location, she can freely choose what actions take place during the measurements. For example, she could do measurements while no persons are present at the training location, or some persons are present and execute predefined activities.

2.3.3 Attack Procedure

The attack procedure is as follows: First, the attacker collects training data at a training location, which might be the actual target location or other rooms similar in size, layout, sensor positions, and furniture. The attacker also records the ground truth for all events that shall be distinguished, e. g., the number of persons present in the room (including the case that the room is empty), or different activities such as working, walking, and sleeping. The training data is recorded with a sample rate of a few seconds and split into *windows* (i. e., a temperature curve and a relative humidity curve) of the same time length, usually one to three minutes. Using the collected training data, the attacker trains a machine learning classifier. After the classifier is trained, it can be used to classify windows of climate data from the target location to determine the events. The classifier works on previously collected

data, thus reconstructing past events, and also on live-recorded data, thus determining current events "on-the-fly" at the target location.

2.3.4 Scope

Note that we do not claim that analyzing room climate data is the most effective attack to gain the information specified above. It is our ambition to demonstrate that room climate data potentially reveals privacy-harming information and hence needs to be protected. In other words, we are *not* trying to present the best attack for occupancy detection (or estimation) or activity recognition, but rather to determine whether room climate data should be considered *private* data. Note that our experiments show that occupancy and activities can be determined in almost real-time, i. e., after measuring the room temperature and humidity for a small time window. Even if an attacker may get access to such data only at some later point in time, the risk of privacy violation remains, e. g., the measurements may allow one to determine the user profile.

2.4 Experimental Design and Methods

We conducted a study to investigate the feasibility of detecting and estimating occupancy as well as inferring activities in an office environment from temperature and relative humidity data alone. From March to April 2016, we performed experiments at two locations simultaneously, *Location A* and *Location B*, with a distance of approximately 200 km between them. In addition, from January to February 2017, we conducted further experiments at a third location, denoted as *Location C*, which is in the same building as Location B.

2.4.1 Experimental Setup and Tasks

The experimental spaces at the three locations are different in size, layout, and sensor positions (see Figure 2.1). Thus, each target location also serves as the training location in our study. At Location A, the room has a floor area of 16.5 m^2 and was equipped with room climate sensors at four positions as shown in Figure 2.1a. At Location B, the room has a floor area of 30.8 m^2, i. e., roughly twice as much as Location A, and had room climate sensors installed at three positions as illustrated in Figure 2.1c. Location C has a floor area of 13.9 m^2 and was equipped with room climate sensors at five positions as shown in Figure 2.1b. In all locations, the room climate sensors measured temperature and relative humidity. The number of deployed sensors varied due to limitations of hardware availability.

Our goal was to determine the extent to which the presence and activities of occupants influence the room climate data. Therefore, we measured temperature and relative humidity during phases of absence as well as phases of presence. If occupants were present, these persons had to perform a single task or a sequence of tasks. We defined the following four experimental tasks (see also Figure 2.2):

- **Read:** Sit on an office chair next to a desk and read.
- **Stand:** Stand in the middle of the room, try to avoid movement.
- **Walk:** Walk slowly and randomly through the room.
- **Work:** Sit on an office chair next to a desk and use a laptop located on the desk.

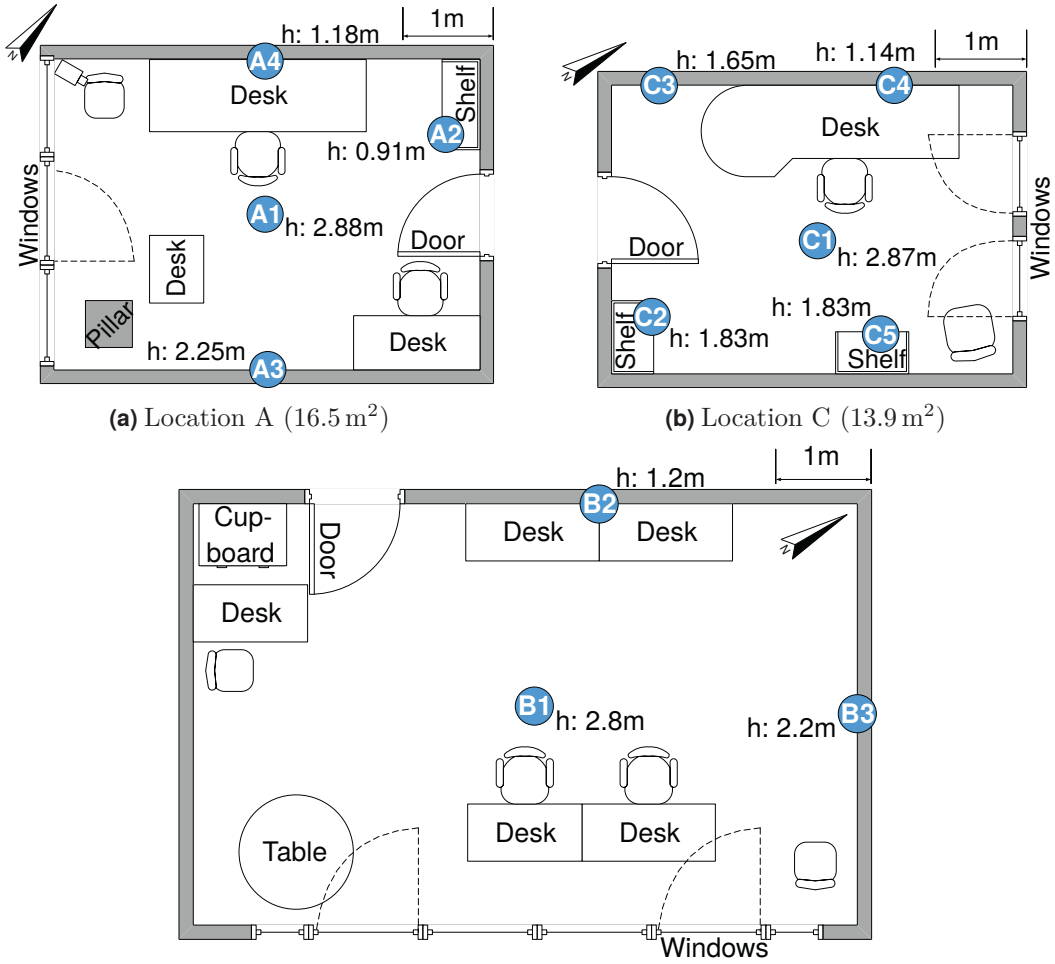

(a) Location A ($16.5\,\mathrm{m}^2$)

(b) Location C ($13.9\,\mathrm{m}^2$)

(c) Location B ($30.8\,\mathrm{m}^2$)

Figure 2.1: Floor plans of the three experiment spaces including sensor node locations, h indicates the respective node's height.

To eliminate confounding factors, we defined location default settings applying to all locations. Essentially, all windows were required to remain closed, and no person was allowed in the room when not in use for the experiment. The rooms have radiators for heating, which were adjusted to a constant level. At Locations A and B, we used shutters fixed in such positions that enough light was provided for reading and working.

2.4.2 Sensor Data Collection

We used a homogeneous hardware and software setup at all locations for data collection, which is described in the following.

2.4.2.1 Hardware

At each location, we set up a sensor network consisting of several Moteiv Tmote Sky sensor nodes with an integrated IEEE 802.15.4-compliant TI CC2420 radio transceiver [Mot06;

(a) Read **(b)** Stand **(c)** Walk **(d)** Work

Figure 2.2: The four defined tasks performed by participants at Location A.

Tex13] as well as an integrated temperature and relative humidity sensor. The nodes have the Contiki operating system [DGV04] version 2.7 installed. In addition, we deployed a webcam that took pictures in a 3-second interval at Location A. These were used for verification during the data collection phase only, and were not given to the classification algorithms.

2.4.2.2 Software

For sensor data collection, we customized the *Collect-View* application included in Contiki 2.7, which provides a graphical user interface to manage the sensor network. For our purposes, we implemented an additional control panel offering a customized logging system. The measurement settings of the Collect-View application were set to a report interval of 4 seconds with a variance of 1 second, i. e., each sensor node reported its current values in a time interval of 4 ± 1 seconds. The variance is a feature provided by *Collect-View* to decrease the risk of packet collisions during OTA transmissions.

2.4.2.3 Collected Data

We structured data collection in *units* and aimed for a good balance between presence and absence as well as the different tasks among all units, as this is needed for the later analysis using machine learning. Each unit has a fixed time duration, t, where exactly one or two persons were present ($t \in \{10, 30, 60\}$, in minutes) who executed predefined activities. In the two-occupants case, both performed the same activities simultaneously, which allows investigation whether it is possible to distinguish the number of persons in a room. If the presence time was t minutes, then the absence time before and after it, respectively, was determined as $\frac{t}{2} + 5$ minutes, where 5 minutes served as buffer. This accounts for both the equal distribution of presence time and absence time, respectively, and the fact that temperature and humidity settle within a 15-minute period after the 60-minute presence of one person.

Overall, we collected around 115 hours of sensor data, 66 hours with at least one person being present. A more extensive overview of the amount of measured sensor data is shown in Table 2.2. To encourage replication and further investigations, all collected sensor data is available as open data sets on GitHub.[2]

[2] https://github.com/IoTsec/Room-Climate-Datasets

Table 2.2: Measured sensor data of all three locations (in hours), detailing the duration of occupancy levels (0, 1, or 2), and tasks (Read, Stand, Walk, or Work) performed by one or two persons per location. At Location C, only one-person experiments were conducted.

Variable	Value	Recorded Time		
		Location A	Location B	Location C
Occupancy	0	20:38:26	15:21:00	13:21:42
	1	14:41:56	11:33:06	13:44:29
	2	14:51:05	11:41:22	–
Task (one person)	Read	4:46:13	2:56:44	3:19:47
	Stand	2:45:27	2:34:20	3:28:27
	Walk	2:43:53	2:37:12	3:20:05
	Work	4:03:33	3:00:20	3:20:52
Task (two persons)	Read	4:26:43	2:58:43	–
	Stand	2:51:08	2:37:11	–
	Walk	2:43:23	2:38:20	–
	Work	4:27:57	2:58:24	–

2.4.3 Experimental Procedure

The participants were assigned to at least one experimental unit with fixed presence times and tasks, and provided with a script for their actions (that is, for how long and in which order the tasks should be performed). Every participant performed each unit twice, with the same tasks, but possibly on different days and in a permuted chronological order. Tasks were performed in blocks of 10, 20, or 30 minutes. Thus, 10-minute units contained only one task of 10 minutes; 30-minute units consisted of either three tasks of 10 minutes or one task of 10 minutes plus one of 20 minutes; 60-minute units were composed of either two tasks of 20 minutes plus two of 10 minutes, or one task of 10, 20, and 30 minutes each.

At the beginning of the presence time for each unit, i. e., the time period where one or two persons had to be present, the experimental supervisor unlocked the room door to let the participants in. The participants started with the first task and were instructed by phone (at Locations A and C) or through a glass pane (at Location B) when it was time to change activities or to leave the room.

2.4.4 Participants and Ethical Principles

For participating in the experiment, 14 subjects volunteered at Location A, 12 subjects at Location B, and 10 subjects at Location C as shown in Table 2.3. Demographic data of participants was collected in order to facilitate replication and future experiments. All subjects gave their written informed consent after the study protocol was approved by the data protection office.[3] We assigned each participant to a random ID. All collected sensor data as well as the demographic data is only linked to this ID.

[3]Ethical review boards at both locations only consider medical experiments.

Table 2.3: Demographic data of participants per location. μ denotes the average, σ the standard deviation. Weight is reported in kg, height in cm, and age in years.

Characteristic		Location		
		A	B	C
Gender	f:	3	2	5
	m:	11	10	5
Weight	μ:	74.9	81.7	63.1
	σ:	8.0	12.1	10.0
Height	μ:	175.9	178.4	170.7
	σ:	9.2	5.3	9.3
Age	μ:	33.7	30.3	25.6
	σ:	8.2	4.8	2.8

2.4.5 Classifier Design

To predict occupancy and activities in the rooms, we used a classification-based method. We adopt an approach that has successfully been used in several applications of biosignal processing, namely extraction of a number of statistical descriptors with subsequent feature selection [HPP+16; JBKE16].

2.4.5.1 Data Segmentation

The features use measurements from short time windows. We experimented with windows of different lengths, namely 60 s, 90 s, 120 s, 150 s, and 180 s. The offset between two consecutive windows was set to 30 s. We excluded all windows where only a part of the measurements belongs to the same activity.

2.4.5.2 Feature Extraction

The feature set was composed from a number of statistical descriptors that were computed on temperature and humidity measurements within the time windows. These are

- mean,
- variance,
- skewness,
- kurtosis,
- number of nonzero values,
- entropy,
- difference between maximum and minimum value of the window (i. e., value range),
- correlation between temperature and humidity,

- mean of the regression line for the preceding measurement window, and
- slope of the regression line for the preceding measurement window.

Additionally, we subtracted from the measurements their least-squares linear regression line, and computed all of the listed statistics on the subtraction residuals as well.

2.4.5.3 Feature Selection and Classification

Feature selection was performed using a sequential forward search [WFH11, Ch. 7.1 & 11.8], with an inner leave-one-subject-out cross-validation [HTF09, Ch. 7] to determine the performance of each feature set. For classification, we used the Naïve Bayes classifier. To avoid a bias in the results, we randomly selected identical numbers of windows per class for training, validation, and testing. For implementation, we used the ECST software [RJKE12] which wraps the WEKA library [HFH+09].

Note that neither the features nor the rather simple Naïve Bayes classifier are particularly tailored to predicting privacy leakage. However, we also show that such an unoptimized system is able to correctly predict occupancy and activities and hence produce privacy leakage. Higher detection rates can be expected if more advanced classifiers are utilized.

2.4.5.4 Evaluation Metrics

As performance measures, we use accuracy (i. e., the number of correctly classified windows divided by the number of all windows), and per-class sensitivity (i. e., the number of correctly classified windows for a specific class divided by the number of all windows of this class). Classification accuracy was deemed statistically significant if it was significantly higher than random guessing, which is the best choice if the classifier could not learn any useful information during training. For each experiment, a binomial test with significance level $p < 0.01$ was carried out using the R software [RCo14].

2.5 Results

In this section, we present the experimental results. First, a visual inspection of the collected data is presented, followed by the machine learning-aided occupancy detection and estimation, and activity recognition using one or two sensors, respectively.

2.5.1 Visual Inspection

We started our evaluation by analyzing the *raw* sensor data. Hence, we implemented a visualization script in MATLAB, which plots the data. The visualizations of two measurements are shown as examples in Figure 2.3.

The visualizations show an immediate rise of the temperature and humidity as soon as an occupant enters the room. Furthermore, variations in temperature and humidity increase rapidly and can be clearly seen. Thus, one can visually distinguish between phases of occupancy and non-occupancy. One can also notice different patterns during the performance of the tasks. As Figure 2.3a shows, an occupant performing the task Walk in the experimental space causes a constant increase of temperature and humidity with only small variations. In contrast, an occupant performing the task Stand in the room causes the largest variations

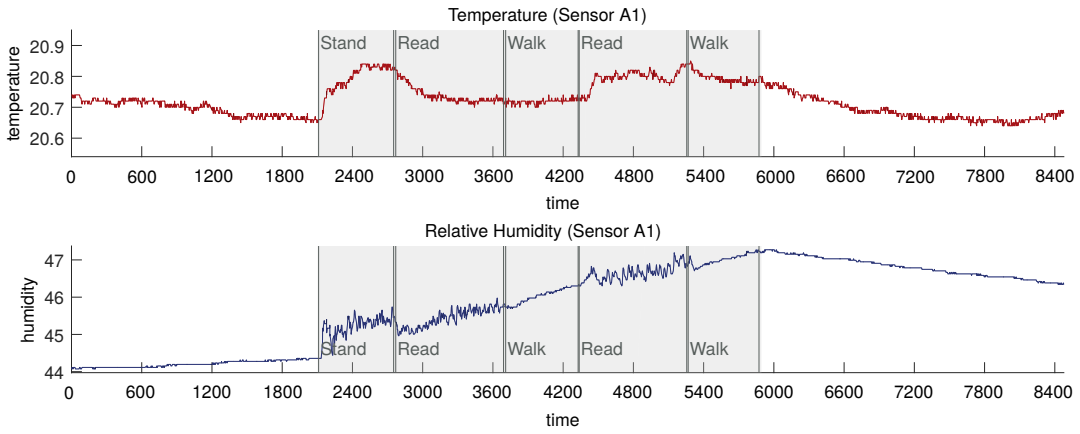

(a) One occupant is present for 60 minutes at Location A.

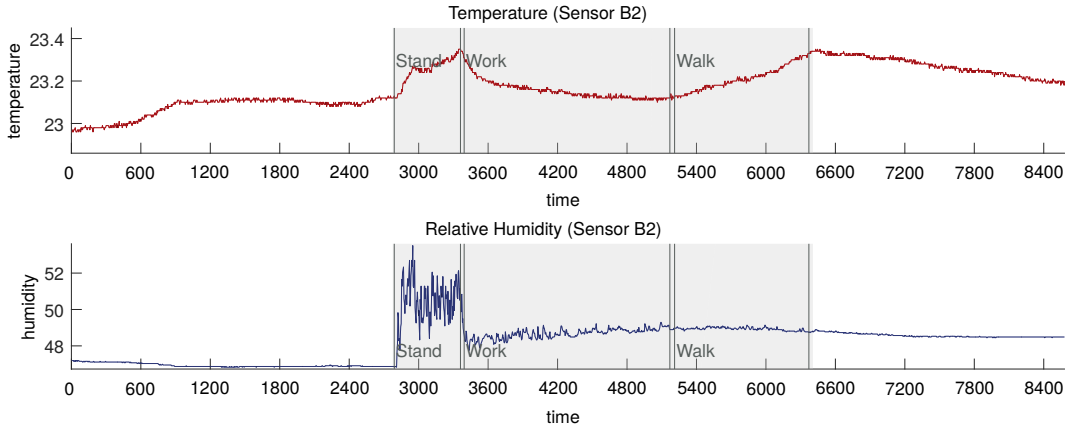

(b) One occupant is present for 60 minutes at Location B.

Figure 2.3: Visualization of two examples of room climate measurements, i. e., temperature and relative humidity. The gray background indicates the presence and performed activity of the occupant inside the experimental space. Time is shown in seconds, temperature in °C (upper red-colored plot in a and b, respectively), relative humidity in % (lower blue-colored plot in a and b, respectively).

of humidity compared to the other defined tasks (cf. Figure 2.3b). The effects of the tasks Read and Work on temperature and humidity in the depicted figures are very similar: both variables tend to increase showing medium variations. For further analysis of the data, we used machine learning as outlined in Section 2.4.5.

2.5.2 Occupancy Detection

Occupancy detection describes the binary detection of occupants in the experimental space based on features from windows with a length of 180 seconds (cf. Section 2.4.5). This is a two-class task, namely to distinguish whether an occupant is present (true) or not (false). We only considered training and testing data within the same room (but separated training and testing both by the days and participants of the acquisition). We randomly selected the same number of positive and negative cases from the data. Thus, simply guessing the state

Table 2.4: Classification accuracy for *occupancy detection*. Notation: 'Occup.' – sensitivity for class *occupancy*, 'No Occup.' – sensitivity for class *no occupancy*, 'Guess' – probability of correct guessing, 'Acc.' – classification accuracy. All numeric values in %.

Sensor	Sensitivity		Guess	Acc.
	Occup.	No Occup.		
A1	94.1	93.0	50.0	93.5
A2	94.5	85.0	50.0	89.7
A3	92.0	76.4	50.0	84.2
A4	77.8	79.1	50.0	78.4
B1	91.9	85.1	50.0	88.5
B2	85.3	77.2	50.0	81.3
B3	69.7	63.9	50.0	66.8
C1	92.9	89.2	50.0	91.0
C2	89.9	87.4	50.0	88.6
C3	90.0	82.0	50.0	86.0
C4	89.8	87.6	50.0	88.7
C5	92.5	88.8	50.0	90.7

has a success probability of 50 %. However, our classification results are considerably higher than that. Table 2.4 shows that the highest accuracies per location were 93.5 % (Location A), 88.5 % (Location B), and 91.0 % (Location C). Considering all sensors of all three locations, detection accuracy ranges between 66.8 % (Sensor B3) and 93.5 % (Sensor A1) as shown in Figure 2.4a. All classification accuracies were statistically significantly different from random guessing. This indicates that an attacker can reveal the presence of occupants in a target location with high probability.

2.5.3 Occupancy Estimation

Results for detecting multiple persons are shown in Table 2.5. Here, we first restate the results for deciding whether a single person is in a room or not, denoted as "0–1". Then, we report the results for deciding whether one or two persons are in a room, denoted as "1–2". In the third column, we report results on the joint classification problem, denoted as "0–1–2". While the first two columns cover two-class classification tasks (i. e., with a guessing chance of 50 %), the third column covers a three-class classification task (with a guessing chance of 33.3 %). Decisions whether a room is empty or occupied by a single person achieve relatively high accuracies, for some sensors, close to or even above 90 %. However, distinguishing between one or two persons is much harder. At Location A, three out of four sensors perform around guessing chance, whereas at Location B, performances range between 60 % and 70 %. The three-class task confirms these two initial results, lying between the well-discernible case of an empty room and the much harder case of determining the exact number of occupants. From an attacker's perspective, deciding whether one or two persons are present in a target location is an ambiguous task.

Table 2.5: Occupancy detection and estimation accuracies for two- and three-class tasks at Locations A and B (in %) on a sensor level. 0–1 refers to distinguishing between no and one person present, 1–2 between one and two persons present, 0–1–2 between no, one, and two persons present.

Sensor	0–1	1–2	0–1–2
A1	93.5	47.7	59.7
A2	89.7	47.8	58.7
A3	84.2	68.0	66.4
A4	78.4	48.2	54.9
B1	88.5	67.6	72.2
B2	81.2	60.7	60.3
B3	66.8	67.3	56.6

2.5.4 Activity Recognition

Activity recognition reports the current activity of an occupant in the experimental space. The four activities are described in Section 2.4.1. The recognition results for these tasks are shown in Figure 2.4.

ACT4 classifies between the activities Read, Stand, Walk, and Work. As shown in Figure 2.4b, the accuracy of recognizing activities achieved by the machine learning pipeline ranged from 23.9 % (Sensor C1) to 56.8 % (Sensor A1). Overall, the accuracy of ACT4 was statistically significantly better than the probability of guessing the correct task (25 %) for 8 out of 12 sensors. Thus, the distinction between multiple activities is possible, but depends on the target location and the position of the sensor.

In the next step, we investigated whether an attacker can increase the recognition accuracies by distinguishing between a smaller set of activities. To this end, we combined two tasks to a *meta task*, e. g., the tasks Read and Work became Sit. The model ACT3 classifies between the tasks Sit, Stand, and Walk. The probability of correct guessing is thus 33.3 %. This model is typical to represent activities of an occupant in a private space or an office room. For ACT3, the achieved accuracy ranged from 31.8 % (Sensor C1) to 81.0 % (Sensor A1). Our results were statistically significant for 10 out of the 12 sensors deployed in the three locations (see Figure 2.4c). Assuming a known layout of the target location, the attacker might be able to determine the position of the occupant in the space and infer activities such as watching TV, exercising, cooking, or eating.

The model ACT2 classifies between the tasks Sit and Upright, whereby Sit is as previously Read or Work, and Upright combines Stand and Walk. In this classification, the attacker distinguishes whether an occupant is at a certain posture. The model ACT2a classifies between the tasks Read and Work, and the model ACT2b classifies between the tasks Stand and Walk. ACT2a indicates that an attacker can even distinguish between the sedentary activities, such as reading a book or working on the laptop. In contrast, ACT2b shows that an attacker can differentiate between standing and moving activities. Thus, an attacker can detect movements at the target location. For ACT2, ACT2a, and ACT2b, the probability of guessing the correct class is 50 %. Using these models, the attacker can infer various work and life habits.

For ACT2, accuracies vary between 54.6 % (Sensor C2) and 82.1 % (Sensor A1), and all

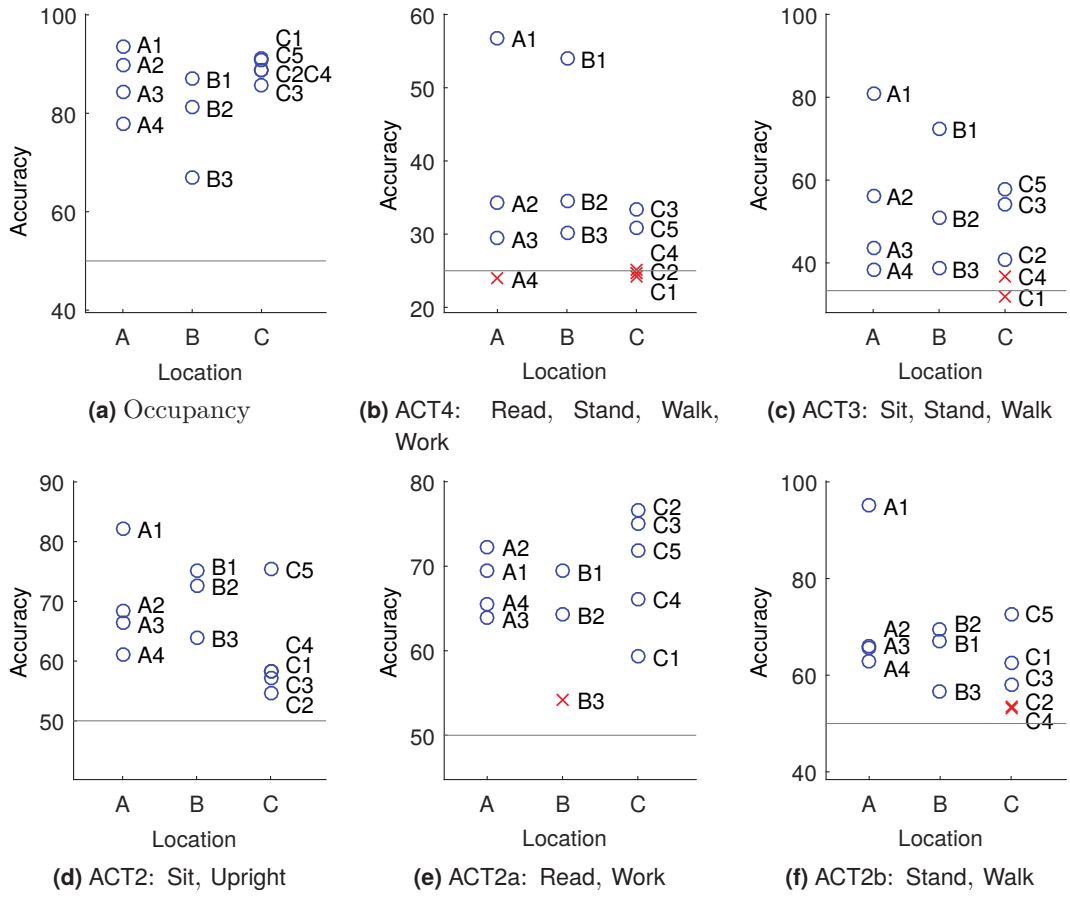

Figure 2.4: Classification accuracy (in %) for occupancy detection and activity recognition. In each diagram, the guessing probability is plotted as a solid line. Each symbol represents the accuracy achieved with a single sensor. A '○' marks a statistically significant result, '×' a statistically *insignificant* result.

accuracies are statistically significant (see Figure 2.4d). For ACT2a, the lowest and highest accuracies were 54.2 % (Sensor B3) and 76.6 % (Sensor C2), respectively, with statistical significance for 11 out of 12 sensors (see Figure 2.4e). For ACT2b, the accuracy ranged from 53.3 % (Sensor C4) to 95.1 % (Sensor A1), and the results for 10 out of 12 sensors are statistically significant (see Figure 2.4f).

2.5.5 Multi-Sensor Classification

Another interesting question is whether the availability of data from multiple sensors within a room can improve the prediction performance. To assess this question, we concatenated the feature vectors of two sensors prior to the classification. Consistently, the evaluation process was kept identical to the previous analyses.

Table 2.6 shows the results for occupancy estimation by using a pair of sensors. Overall, the detection accuracies improve by several percent when combining the feature vectors of two sensors. However, distinguishing between one and two persons in a room is still a hard task, with accuracies ranging approximately from 50 % to 70 %. The three-class problem

Table 2.6: Occupancy detection and estimation accuracies denoted 'acc.' (in %) for two- and three-class tasks at Locations A and B based on two sensors. The differences to the accuracies of single sensors are reported as ΔSm and ΔSn, respectively. For example, in the occupancy detection case (0–1), the accuracy when measuring with the two sensors A1 and A2 increases to 96.3 %. When measuring with A1 only, the accuracy is 93.5 % (cf. Table 2.4), which constitutes the difference of +2.8 % (ΔSm). The difference to measuring with A2 only is +6.6 % (ΔSn).

Sensors	0–1			1–2			0–1–2		
Sm-Sn	acc.	ΔSm	ΔSn	acc.	ΔSm	ΔSn	acc.	ΔSm	ΔSn
A1-A2	96.3	+2.8	+6.6	60.8	+13.1	+13.0	68.1	+8.4	+9.4
A1-A3	93.9	+0.4	+9.7	71.5	+23.8	+3.5	75.7	+16.0	+9.3
A1-A4	94.2	+0.7	+15.8	56.6	+8.9	+8.4	69.0	+9.3	+14.1
A2-A3	90.1	+0.4	+5.9	70.9	+23.1	+2.9	74.5	+15.8	+8.1
A2-A4	87.8	−1.9	+9.4	53.7	+5.9	+5.5	61.2	+2.5	+6.3
A3-A4	84.8	+0.6	+6.4	69.8	+1.8	+21.6	70.6	+4.2	+15.7
B1-B2	91.7	+3.2	+10.5	72.0	+4.4	+11.3	75.1	+2.9	+14.8
B1-B3	89.6	+1.1	+22.8	67.9	+0.3	+0.6	71.7	−0.5	+15.1
B2-B3	83.4	+2.2	+16.6	61.4	+0.7	−5.9	64.8	+4.5	+8.2

also slightly improves, but still, by far the most reliable classification task is to determine whether there is a person in the room at all, with almost all accuracies now in the range of 90 % and above.

Table 2.7 shows the results for activity recognition. Results for a single sensor are shown on the left side, results for pairs of two sensors on the right side. Interestingly, classification of activities appears to be a similarly hard task for pairs of sensors as it is for single sensors: using the information from multiple sensors does not really improve the results for activity classification. One reason might be that the information between multiple sensors is correlated for this task, such that the addition of multiple sensors only enlarges the feature space, but does not add significant information. From a privacy perspective, however, this is somewhat soothing: although the achieved accuracies are in almost all cases well above guessing chance, it is apparently not a straightforward task to accurately predict the activity of a person in the room.

2.6 Further Observations

This section discusses observations regarding our evaluation, i. e., we elaborate on considered window sizes, selected features, room size and layout, and sensor positions as well as their respective influence on classification performance.

2.6.1 Length of Measurement Windows

The length of the measurement windows influences the accuracy of detection. We evaluated window sizes in the range between 60 and 180 seconds. As an example, we analyzed the average accuracy of occupancy detection depending on the window size for all three locations.

Table 2.7: Activity Recognition accuracies (in %) for all three locations based on a single sensor and on two sensors. Values are reported on a sensor level and a sensor pair level, respectively, classes are as defined in Section 2.5.4.

One sensor						Two sensors					
Sensor	ACT4	ACT3	ACT2	ACT2a	ACT2b	Sensors	ACT4	ACT3	ACT2	ACT2a	ACT2b
A1	56.8	81.0	82.6	57.5	96.3	A1-A2	50.2	75.4	83.1	74.8	91.0
A2	39.7	56.1	68.5	64.1	65.9	A1-A3	50.1	76.7	80.4	72.4	88.3
A3	32.7	47.2	66.7	60.3	68.7	A1-A4	54.5	79.0	83.0	67.3	89.4
A4	26.8	39.8	62.8	65.6	69.0	A2-A3	32.6	50.0	65.9	74.8	67.1
						A2-A4	41.1	55.1	71.0	75.5	69.1
						A3-A4	31.1	43.8	68.3	75.3	46.8
B1	55.1	76.0	75.6	63.4	71.5	B1-B2	52.7	77.0	80.0	70.1	86.0
B2	35.2	56.4	74.5	62.1	73.1	B1-B3	55.6	74.5	76.4	64.2	79.6
B3	33.3	39.5	63.9	50.9	62.2	B2-B3	39.3	54.5	72.1	59.8	62.7
C1	28.3	40.0	58.4	59.3	43.5	C1-C2	29.5	46.8	55.8	65.2	62.5
C2	25.8	41.7	60.0	71.6	55.3	C1-C3	38.8	52.5	60.2	63.8	76.9
C3	34.3	54.2	62.4	66.3	60.7	C1-C4	28.6	44.6	67.8	64.4	53.3
C4	25.1	36.8	62.6	66.1	53.3	C1-C5	36.6	52.4	68.5	64.1	65.4
C5	30.9	57.7	75.5	65.8	72.5	C2-C3	32.2	52.5	62.4	76.5	72.8
						C2-C4	29.0	48.9	64.5	77.3	59.3
						C2-C5	30.5	49.3	72.3	75.5	67.1
						C3-C4	40.2	54.7	67.9	72.4	69.7
						C3-C5	42.4	60.5	70.9	74.7	77.4
						C4-C5	35.3	57.1	72.6	72.4	64.5

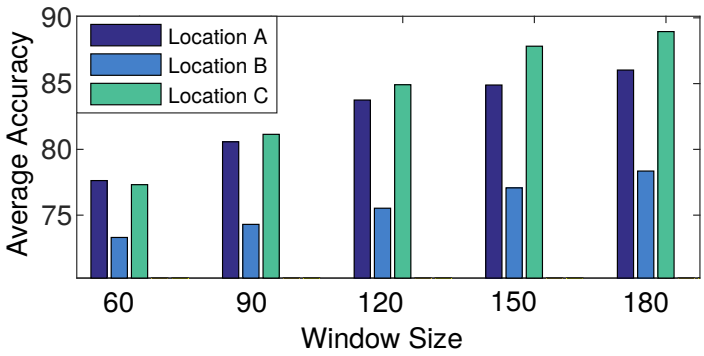

Figure 2.5: Average accuracy (in %) over all sensors from each location for occupancy detection depending on the window size (in s).

As shown in Figure 2.5, the accuracy increases with a larger window. We achieved the best results with the largest window size of 180 seconds, indicating that the highest accuracies are possible if longer time periods are considered. From a practical perspective, it is not advisable to extend the window size to a much larger duration than a few minutes since we assume that the performed activity is consistent for the whole duration of the window.

2.6.2 Selected Features

To assess the feasibility of an attacker that has access only to either temperature data or relative humidity data, we evaluated whether it might be enough to solely collect one type of room climate data. In the classification process, an attacker derives a set of features from temperature and relative humidity data and selects the best-performing features for each sensor and classification goal automatically (cf. Section 2.4.5). Analysis shows that features computed from temperature and relative humidity are of similar importance. In our evaluation, 57.9 % of the selected features are derived from temperature measurements, and 52.3 % from relative humidity measurements.[4]

We also compared the features in terms of differences between the three locations as well as differences between occupancy detection and activity recognition. In all these cases, there are no significant differences between the importance of temperature and relative humidity. An attacker restricted to either temperature or relative humidity data will perform worse than with both data.

2.6.3 Size and Layout of Rooms

All our locations are office-like rooms, which have a similar layout (rectangular) but differ in size and furnishing. In our evaluation, the accuracy correlates (negatively) with the size of the target location. As shown in Figure 2.5, we had the highest average accuracy in occupancy detection with Location C, which also has the smallest floor area of 13.9 m^2. Location A has a floor area of 16.5 m^2, and has a slightly lower average accuracy. Location B is almost twice as large (30.8 m^2) and shows the worst average accuracy compared to the other locations. Thus, our experiment indicates that an increasing room size leads to

[4]Note that some features are based on both, temperature and relative humidity, which is why the sum of both numbers exceeds 100 %.

decreasing accuracy on average. An attacker achieves higher accuracies by monitoring target locations of a small size compared to target locations of larger sizes.

2.6.4 Position of Sensors

According to our threat model in Section 2.3, the attacker controls the target location's layout. Thus, we assume an attacker who can decide where to install room climate sensors in the target location. We consider how the position of a room climate sensor influences the accuracy of derived information. For occupancy detection, we had the best accuracy with a sensor node that is located in the center point at the ceiling of the target location (Sensors A1, B1, C1). In this position, the sensor has the largest gathering area to measure the climate of the room. Sensors mounted to the walls or on shelves perform differently in our experiments. For occupancy estimation, the best sensors differ per location, i. e., wall-mounted A3 outperforms all others at Location A, whereas B3 (similar to A3) and B1 are almost on par. For activity recognition, the central sensor nodes performed best at Locations A and B, but not at Location C.

From the attacker's perspective, the best position to deploy a room climate sensor is at the ceiling in the center of the target location. In large rooms, multiple sensors at the ceiling could be installed, each covering a subsection of the room.

2.7 Discussion

As our experiments reveal, knowing the temperature and relative humidity of a room allows detection of the presence of people and recognition of certain activities with a significantly higher probability than guessing. By evaluating temperature and relative humidity curves of the length of 180 seconds, we were able to detect the presence of an occupant in one of our experimental spaces with an accuracy of 93.5 % using a single sensor. Occupancy estimation results show no definitive trend as three out of four sensors perform around guessing chance for deciding between one or two persons, yet, the best accuracies range around 68 %. The classification accuracies for the three-class problem, deciding between zero, one, and two persons being present, lie consequently between the occupancy detection and one-versus-two-person results with 72.2 % as the best result. In terms of activity recognition, we distinguished between four activities with an accuracy of up to 56.8 %, between three activities of up to 81.0 %, and between two activities of up to 95.1 %. Thus, an attacker focusing on the detection of a specific activity is more successful than an attacker that aims to classify a broader variety of activities. In the following, we discuss relevance, implications, and limitations of our results.

2.7.1 Privacy Implications

We show that an attacker might be able to infer occupants' work and life habits from room climate data. Thus, the attacker is able to distinguish between sitting, standing, and moving, which already might reveal the position and activities of an occupant within a room. Moreover, the attacker can distinguish between upright and sedentary activities, between moving and standing, and between working on a laptop and reading a book.

Given the limited amount of recorded sensor data, the achieved accuracies in occupancy detection and activity recognition give a clear indication that occupants are subject to pri-

vacy violations according to the threat model described in Section 2.3. However, occupancy estimation respectively activity recognition are not straightforward since the achieved accuracies are low respectively differ between the different sensor positions and locations.

On the bright side, it is also reassuring that simply increasing the number of sensors is not as ominous as one might fear since the relative accuracy increase is rather slim. In other words, the most restrictive scenario, i. e., one deployed sensor, is sufficient. Hence, for an attacker, the benefit of deploying and exploiting more than one sensor is at least questionable.

Further experiments are required for a better assessment of the privacy risks induced by the room climate data. Our work provides promising directions for these assessments. For example, we demonstrated the existence of the information leak with the Naïve Bayes classifier. Naïve Bayes is arguably one of the simplest machine learning classifiers. In future work, it would be interesting to explore upper bounds for the detection of presence/absence, occupancy estimation, and different activities by using more advanced classifiers such as the recently popular deep learning algorithms.

2.7.2 Policy Implications

Smart heating is one of the most desirable smart home applications. For example, in a representative smart home survey of German consumers from 2015, 34 % of the participants stated that they are interested in technologies for intelligent heating or are planning to acquire such a system [Del15]. Another survey with 1,000 US and 600 Canadian consumers found that for 72 % of them, the most desired smart home device would be a self-adjusting thermostat, and 37 % reported that they were likely to purchase one within the next twelve months [ico15].

While consumers understand the benefits of smart heating systems, their privacy implications may be difficult to fathom. For example, in a recent representative survey with 461 American adults by Pew Research [RD16], the participants were presented with a scenario of installing a smart thermostat "in return for sharing data about some of the basic activities that take place in your house like when people are there and when they move from room to room" [RD16, Section 7]. Whereas 55 % of respondents said that this scenario was not acceptable for them, 27 % said that it was acceptable, with the remaining 17 % answering "it depends" [RD16, Section 7]. Furthermore, in a worldwide survey with 9,000 respondents from nine countries (Australia, Brazil, Canada, France, Germany, India, Mexico, the UK, and the US), 54 % of respondents said that "they might be willing to share their personal data collected from their smart home with companies in exchange for money" [Int16].[5]

According to the results of the above surveys, quite a high percentage of consumers might be willing to share their room climate data. On the other hand, our experiments show that this sharing may have serious privacy implications. Moreover, our findings indicate that room climate data measurements should be classified as personal data that can be used to profile the user, as specified in Art. 4 of the General Data Protection Regulation (GDPR) [Eur16]. For example, for single households, room climate data may leak personal data about when the person is at home, and whether this person is in company of some other person at a certain time. As a consequence, the collection and processing of room climate data is subject to the GDPR. More precisely, it is our understanding that before such data can be processed,

[5] Methodological details, such as representativeness, breakdown by country and the exact formulation of the questions, are not known about this survey.

a Data Protection Impact Assessment (DPIA) needs to be conducted. Citing from Art. 35 of the GDPR: "Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data." Note that the EU Commission already developed DPIAs for similar IoT applications: Smart Metering[6] and Radio-Frequency Identification (RFID)[7]. These may serve as a starting point for smart heating applications.

Moreover, according to Art. 25 GDPR, data protection by design and by default should be applied to smart heating systems. Our experiments show that the amount of personal information that can be inferred from room climate data is highly depending on the granularity of the measurements, and also on the placement of sensors. In this way, our findings may provide a starting point for data minimization guidelines, as they indicate that privacy in smart heating systems can be protected by adjusting data collection granularity and sensor placement.

We think that the idea of sharing smart home data for various benefits will continue to be intensively discussed in the future, and therefore, consumers and policymakers should be made aware of the level of detail inferable from smart home data. As such, it should be clarified which rewards would actually be beneficial for consumers, and which kind of data sharing would be ethically permissible in the first place. Only by addressing these open questions it would be possible to design fair policies and, for instance, to establish beneficial personal data markets [SABH15]. We believe that our work is the first step towards informing the policy for the smart heating scenario.

2.7.3 Location-Independent Classification

An important question is whether it is possible to perform location-independent classification, i. e., to train the classifier with sensor data of one location and then use it to classify sensor data at the target location that is not similar to the training location in size, layout, and sensor positions. If this was possible, the service providers of smart heating applications would be able to detect occupancy and to recognize activities without having access to the target locations.

According to their privacy statements, popular smart thermostats from Nest [Nes16], Ecobee [Eco15], and Honeywell [Hon15] send measured climate data to the respective service provider's databases. To evaluate these privacy threats, we used the room climate data of the best-performing sensor of a location as training data set for other locations. For example, to classify events of an arbitrary sensor of Location A, we trained the classifier with room climate data collected by Sensor B1 or Sensor C1. We gained statistically significant results for a few combinations in occupancy detection, but the majority of our occupancy detection results was not significant. Since discriminating between one and two persons proved unreliable, occupancy estimation was consequently excluded. For activity recognition, we were not able to gain statistically significant results.

However, the possibility of location-independent attackers cannot be excluded. Absence of significant results in our experiments may be merely due to the limited amount of data. Fu-

[6] https://ec.europa.eu/energy/en/content/dpia-template-smart-grid-and-smart-metering-systems
[7] https://ec.europa.eu/digital-single-market/en/news/privacy-and-data-protection-impact-assessment-framework-rfid-applications

ture studies should be conducted to gather data from various rooms up to a point where the combined results hold for arbitrary locations. Having more data from a multitude of rooms available would help the machine learning classifiers to recognize and ignore data characteristics that are specific to either of the experimental rooms. Consequently, the algorithms could better identify the distinct data characteristics of the different classes in occupancy detection and activity recognition. This would enable location-independent classification of room climate data, in which the training location is not similar to the target location regarding size, layout, furnishing, and positions of the sensors.

2.8 Conclusion

We investigated the common belief that data collected by room climate sensors divulge private information about the occupants. To this end, we conducted experiments aiming to reflect realistic conditions, i. e., considering an attacker who has access to typical room climate data (temperature and relative humidity) only. Our experiments revealed that knowing a sequence of temperature and relative humidity measurements already allows the detection of the presence of people and recognition of certain activities with high accuracy. Contrarily, the distinction between the presence of one or two persons is evidently harder, while using data from two different sensors slightly improves occupancy detection, but activity recognition remains a hard task. Nonetheless, our results confirm that the need for protection of room climate data is justified: the leakage of even such 'inconspicuous' sensor data as temperature and relative humidity can seriously violate privacy in smart spaces.

Future work is required to determine the level of privacy invasion in more depth and to develop appropriate countermeasures. Apart from data minimization techniques, e. g., regulating the frequency of measurements or the placement of sensors, additional protective measures could be explored. For instance, it would be interesting to examine whether (artificial) noise added to the measured sensor data could significantly reduce the information gain. Another open question is whether the influence that human activities have on the room climate could be artificially recreated, e. g., to simulate occupancy of an unoccupied space or a specific target activity different from the actual one. Future work could address these and further conceivable countermeasures as well as their effort, costs, and impact.

I operate by investigating the fundamental interconnectedness of all things.

Dirk Gently (Douglas Adams)

3 Speak Touchlink and Enter

Abstract

As the previous chapter demonstrated, even basic information like temperature and (relative) humidity emanating from a smart home should be protected from illegitimate or third party access throughout their lifecycle. Otherwise, occupants risk sharing insights into the state of their homes and rooms, potentially without knowing the exact granularity of deducible state information. Consequently, IoT devices inside a smart home should only communicate securely, e. g., by using an established communication protocol offering at least a *secure* encryption mechanism.

A Hundred million IoT devices implement Zigbee, a low-power mesh network standard, and the number is expected to be growing. To facilitate an easy integration of new devices into a Zigbee network, touchlink commissioning was developed. It was adopted in the latest specifications, Zigbee 3.0, released to the public in December 2016, as one of two commissioning options for Zigbee devices. Zigbee 3.0 products can be used in various applications, also including security-critical products such as door locks and intruder alarm systems.

The aim of this chapter is to warn about a further adoption of this specific commissioning mode. We analyze the security of the touchlink commissioning procedure and present novel attacks that make direct use of the standard's features, showing that this commissioning procedure is insecure by design. We release an open-source penetration testing framework to evaluate the practical implications of these vulnerabilities. Evaluating our tools on popular Zigbee-certified products, we demonstrate that a passive eavesdropper can extract key material from a distance of 130 meters. Furthermore, an active attacker is able to take over devices from a distance of up to 190 meters. Our analysis concludes that even a single touchlink-enabled device is sufficient to compromise the security of a Zigbee 3.0 network, and therefore, touchlink commissioning should not be supported in any future Zigbee products.

Declaration of Origin: This chapter is based on the paper *Insecure to the touch: attacking ZigBee 3.0 via touchlink commissioning* [MMB+17], written together with *Philipp Morgner, Stephan Mattejat, Zinaida Benenson, and Frederik Armknecht*, and presented at *WiSec '17*.

3.1 Introduction

Zigbee is a popular standard for wireless low-power communication in the IoT, especially in the domain of smart home networks. The Zigbee Alliance, a non-profit organization of more than 400 member companies maintaining the Zigbee specifications, lists more than 1,300 certified products [Zig17b], and claims to have the largest base of installed IoT devices worldwide [Zig16b] with more than a hundred million devices [Zig17a].

In December 2016, the latest Zigbee specifications, denoted as Zigbee 3.0, were released to the public. These specifications define function clusters for several smart home applications including security-critical applications such as door locks, window shades, and intruder alarm systems. To prevent manipulation and unauthorized control of these applications, appropriate security measures are crucial for Zigbee networks. One of the most critical parts in the security design is commissioning, which is the procedure of either starting a new network or integrating a new node into an existing network. During the process of joining a new node to an existing network, this node needs to receive the network key in a secure manner, which is a challenging task for heterogeneous IoT networks that interconnect products of multiple manufacturers.

Zigbee 3.0 provides two different commissioning procedures to accomplish this task: EZ-Mode commissioning and touchlink commissioning. In this chapter, we focus on touchlink commissioning, which was originally developed to easily integrate devices in connected lighting systems that follow the (legacy) Zigbee Light Link (ZLL) standard. The basic idea of touchlink commissioning is to facilitate close physical proximity instead of cryptographic authentication for joining new nodes to a network. The Zigbee 3.0 specifications inherited the touchlink commissioning procedure as a commissioning option for Zigbee 3.0 products without giving guidelines regarding an application's suitability for touchlink commissioning.

3.1.1 Contribution

Our contribution is twofold. As first contribution, we provide a security analysis of the Zigbee touchlink commissioning procedure, which, to the best of our knowledge, has not been part of a comprehensive security analysis before. During our investigations, we analyzed the specifications and learned that the touchlink communication relies on inter-Personal Area Network (PAN) frames, which are neither secured nor authenticated. Furthermore, the transport of the network key to a joining device is protected solely by a global master key, the so-called touchlink preconfigured link key. This key is distributed to manufacturers of touchlink-enabled products under a Non-Disclosure Agreement (NDA) but was leaked in March 2015 and cannot be renewed due to the backward compatibility demands towards legacy ZLL products. In addition, we learned from the specifications that the distance check between the joining node and the initiator is based on a simple signal strength threshold.

As second contribution, we developed and evaluated a real-world attack system to eavesdrop on and inject packets into Zigbee networks. In this context, we release the open-source penetration testing framework Z3sec[1] that is able to create arbitrary touchlink commands and provides an interface to control Zigbee-certified devices once the network key is known. We evaluated our penetration testing framework on popular Zigbee-certified, touchlink-enabled products of four different manufacturers. In our evaluation, we demonstrated the extraction of the current network key from a distance of 130 meters through

[1] https://github.com/IoTsec/Z3sec

passively eavesdropping on a touchlink commissioning procedure. In the domain of active attacks, we were able to permanently disconnect nodes from the legitimate network, or to reset them to factory-new. We could also trigger the so-called identify action, e. g., causing light bulbs to blink for several hours. Furthermore, we demonstrated that we can remove nodes from their legitimate networks and join them to the attacker's network. In our evaluation, we were able to perform such an active attack from a distance between 15 and 190 meters depending on the tested products. Due to limitations of the experimental setup, even greater distances might be possible.

In conclusion, our evaluation shows that the support of touchlink commissioning is sufficient to compromise the security of Zigbee 3.0 applications. In our threat scenarios, we outline that already a single touchlink-enabled device allows attackers to take control of arbitrary devices in the Zigbee network, also including security-critical applications.

3.1.2 Organization of This Chapter

The remainder of this chapter is organized as follows. We introduce the Zigbee 3.0 standard in Section 3.2. In Section 3.3, we present the threat model for our security analysis. Section 3.4 performs the security analysis of the touchlink commissioning procedure, which contains the description of novel attacks as well as their practical evaluation. In Section 3.5, we describe the disclosure to the manufacturers and their responses. We discuss the results of our evaluation, consequences, and mitigations in Section 3.6. Related work is presented in Section 3.7, while this chapter concludes with Section 3.8.

3.2 Zigbee

Zigbee is a wireless low-power standard that connects embedded devices via Wireless Personal Area Networks (WPANs). Compared to Wireless Local Area Network (WLAN), Zigbee-certified devices send smaller packets and consume far less energy, while Zigbee has a larger wireless range than Bluetooth.

The Zigbee specifications are maintained by the Zigbee Alliance, a global non-profit organization that comprises over 400 member companies, which defines the network, security, and application layers, and supervises the conformance and interoperability of Zigbee-certified products. The *Zigbee 3.0* specifications, ratified in December 2015 and released to the public in December 2016, replace the *Zigbee Pro* specifications [Zig12c] from 2012. The main difference between Zigbee 3.0 and Zigbee Pro is that the Zigbee Pro specifications facilitated several *application profiles* comprising customized sets of features and protocols for specific application areas. Examples of such profiles are Zigbee Home Automation [Zig13] for applications in residential environments, Zigbee Smart Energy [Zig12b] for smart metering, or ZLL [Zig12a] for connected lighting systems. The fragmentation in the Zigbee Pro standard resulted in interoperability problems between Zigbee-certified products of different profiles such that the Zigbee Alliance decided to merge these profiles into one standard, Zigbee 3.0. An exception is the Smart Energy profile, which remains independent due to special requirements of smart metering applications. The Zigbee 3.0 standard is defined in multiple specification documents, of which the Zigbee 3.0 Base Device Behavior specification [Zig16a] and the Zigbee 3.0 Cluster Library specification [Zig16c] are publicly available.

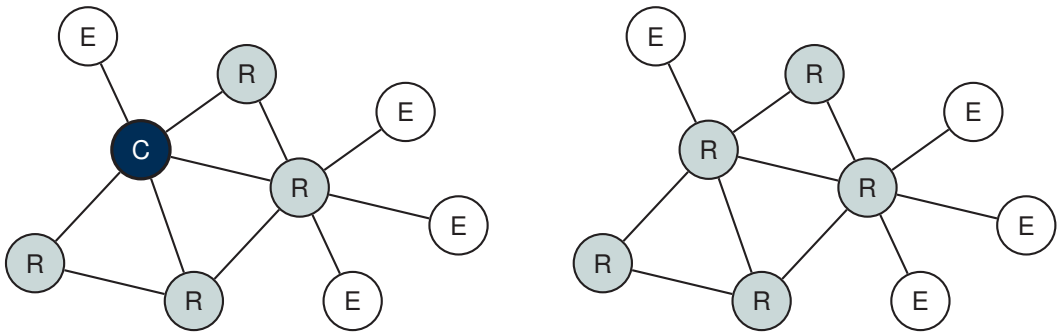

(a) Centralized security **(b)** Distributed security

Figure 3.1: Zigbee's security network models. C denotes a coordinator, R a router, and E an end device.

3.2.1 System Model

The Zigbee 3.0 specifications describe three logical types of nodes: *coordinator*, *router*, and *end device*. Each node can comprise one or more devices, and at any point in time is designated to only one of the logical types. Coordinators and routers are usually devices that have permanent power supply, in contrast to end devices, which are typically battery-powered.

Each Zigbee 3.0 network is either a distributed or a centralized security network (see Figure 3.1). As illustrated in Figure 3.1a, a *centralized security network* is managed by a coordinator that includes the trust center. This coordinator authenticates new nodes and joins them to the network. In contrast, a *distributed security network* is formed by a router and has no coordinator as shown in Figure 3.1b. A new node is authenticated and joined to the network by an arbitrary router, which then becomes the new node's parent node.

3.2.2 Security

The Zigbee 3.0 stack sits on top of the Physical Layer (PHY) and the Media Access Control Layer (MAC) defined in the IEEE 802.15.4 specifications [IEE11]. Security measures in Zigbee applications are only applied to the network layer and the application layer. Although the MAC of the IEEE 802.15.4 standard specifies multiple encryption and authentication mechanisms, these mechanisms are not used in Zigbee applications. Zigbee-certified devices facilitate the AES-CCM* authenticated encryption scheme[2] with a 128-bit *network key*. This network key is shared between all devices of a network and used to secure the communication.

3.2.3 Commissioning

Commissioning is the process in which either a new Zigbee network is started, or a new node is joined to an existing Zigbee network. The Zigbee 3.0 standard specifies two commissioning procedures: *EZ-Mode commissioning* and *touchlink commissioning*. While the support of EZ-Mode commissioning is mandatory for each Zigbee 3.0 device, manufacturers can decide whether to enable touchlink commissioning in their products. The specifications

[2] Compared to AES-CCM (without asterisk), this specific mode allows also encryption-only or integrity-only variants.

do not provide any guidance for whether touchlink commissioning is appropriate for a certain application. In addition, touchlink is supported by all legacy devices that follow the ZLL specifications.

3.2.3.1 EZ-Mode Commissioning

The EZ-Mode must be invoked with a user action, e. g., through pushing a button on the device. After this mode is activated, the node is put into EZ-Mode for a time frame of 3 minutes, which can be extended through further user actions.

In EZ-Mode, a node that is not joined to a network, scans for open networks within its wireless range. In case the node finds a suitable network, it attempts to join this network using MAC association as specified in IEEE 802.15.4. If the network allows the node to join, the node waits to become authenticated and receives the network key. In a centralized security network, the network key is encrypted using either the publicly known *default global Trust Center link key*, which is provided in the Zigbee specifications, or via a pre-configured link key that is derived from an *install code*, which is a unique code printed on the node in a manufacturer-specific fashion. In a distributed security network, the network key is transmitted after being encrypted using the NDA-protected *distributed security global link key*.

3.2.3.2 Touchlink Commissioning

The touchlink commissioning procedure was first introduced in the ZLL standard, and later adopted by the Zigbee 3.0 specifications. Touchlink commissioning is patented by Philips [Lee14] and was specifically designed to make connected lighting systems easy to deploy and use for consumers. Compared to other commissioning options, touchlink commissioning provides an extended functionality that goes beyond the plain joining of devices. The objective was to enable use cases in which commissioning is performed between a bulb and a low-function device, e. g., a remote control. For such scenarios, touchlink commissioning offers the possibility to manage network features, such as reset to factory-new or channel switch, with the so-called *touchlink commands*.

Figure 3.2 describes the commissioning protocol for joining a device, denoted as target or end device, to an existing Zigbee network. The initiator is usually a remote control or a bridge device that is connected to the Internet. First, the initiator starts the device scan procedure by sending a `scan request` on multiple specific channels as defined in the specifications. This request includes a randomly generated `transaction identifier`. The target replies with a `scan response` containing the same `transaction identifier`, a random `response identifier`, and further information. The device scan may yield multiple potential nodes from which the user can select one for the next steps. The user has the option to send an `identify request` to a device, upon which the target performs a pre-defined identify action, e. g., a bulb blinks for a few seconds. An `identify request` contains the corresponding `transaction identifier` as well as the duration of the identify action.

To join a new node to a network, the initiator encrypts the current network key with the *touchlink preconfigured link key*, builds a `network join end device request` containing the encrypted network key, `transaction identifier` as well as further network information, and then sends this command frame to the selected node. On receiving the frame, the joining node decrypts the network key using the touchlink preconfigured link key, and replies

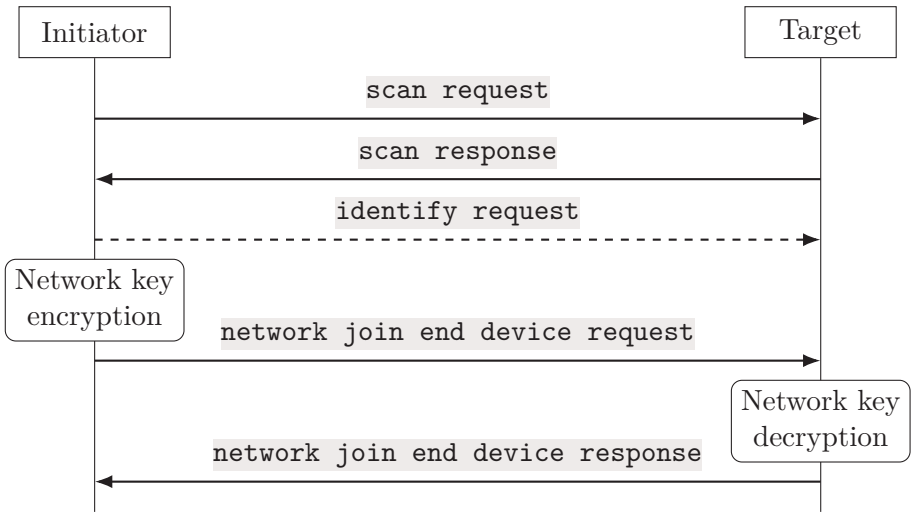

Figure 3.2: Touchlink commissioning protocol.

with a `network join end device response` indicating success or failure.

3.3 Threat Model

In this section, we describe the considered security goals, attacker model, and threat scenarios. This allows for defininining what constitutes an attack based on the attacker's capabilities and goals, potential threats, and ultimately, to identify whether a specific attack is successful.

3.3.1 Security Goals and Attacker Model

The Zigbee specifications describe security assumptions [Zig12c, p. 426], such as safekeeping of key material and proper implementation of security protocols, but do not define security goals [MABM16]. For this reason, we define the following *security goals* that apply to Zigbee networks:

- Confidentiality: Only legitimate entities are allowed to access data and commands sent within the network.
- Integrity: Data and commands sent within the network are not tampered with.
- Authenticity: The receiver is able to reject commands and data sent by illegitimate entities.
- Availability: The functionalities and data of the devices in the network are continuously available to all legitimate entities.

The threat model is determined as follows. The user of the Zigbee network is trusted and honest, and installs the network as required by the manufacturer. Online account credentials for remote access to the Zigbee network are not disclosed. All nodes of the network are certified by the Zigbee Alliance and follow the protocols described in the ZLL or the Zigbee 3.0 specifications. In particular, all touchlink-enabled Zigbee devices are equipped with the touchlink preconfigured link key.

The attacker's goal is to violate any of the security goals mentioned above. We assume that attackers have neither physical access to the Zigbee devices nor to the LAN or WLAN to which a Zigbee device might be connected. The attacker's only capability is to eavesdrop on and inject packets into the wireless communication of at least one touchlink-enabled node of the targeted Zigbee network.

Thus, the attacker controls an IEEE 802.15.4 radio transceiver, which is present within the wireless range of the targeted node. Note that the radio transceiver can be mounted on a drone such that no close physical proximity of the person that performs the attack is required. This potential scenario was demonstrated in [RSWO17].

3.3.2 Threat Scenarios

There exists a wide variety of Zigbee-certified devices in the smart home domain, from sensors like smoke detectors and hygrometers, controllers for building services to appliances like washing machines. The Zigbee 3.0 specifications also define function clusters for security-critical applications, such as front door locks, window shade controllers, and intruder alarm systems. Thus, the attack surface, which we evaluate on touchlink-enabled connected lighting systems, potentially affects all future Zigbee-certified products that are compliant to either the ZLL or Zigbee 3.0 specifications.

Security-critical Zigbee devices can be attacked in two ways: either directly, in case the targeted Zigbee device supports touchlink commissioning, or indirectly, in case there exists at least one touchlink-enabled device in a Zigbee network that contains security-critical devices. We outline the following three threat scenarios.

1. We assume an application that authorizes door access to a restricted area. The network is organized as a centralized security network and communicates based on Zigbee 3.0. One of the nodes of the network is a Zigbee-certified door lock that implements touchlink commissioning. As we show in Section 3.4.3.3, the attacker can reset arbitrary touchlink-enabled nodes to factory-new without knowing any cryptographic secrets. The attacker applies this reset attack to the door lock, which puts the lock in factory-new state and most probably clears the way for the attacker.

2. We assume a smart home network comprising various applications including an intruder alarm system. This network is organized as a Zigbee 3.0 distributed security network. We further assume that an attacker can eavesdrop on the touchlink commissioning procedure of an arbitrary touchlink-enabled device, which is joined to this network, e. g., a new light bulb. From this captured communication, the attacker extracts the network key as described in Section 3.4.4.2. As a consequence, the attacker can decrypt all further network communication as well as inject commands into the network. Hence, the attacker can reset the intruder alarm system to factory-new through sending spoofed network leave commands, and break into the house without triggering the alarm.

3. A household intended for elderly living deploys a distributed security network consisting of numerous touchlink-enabled devices. In Sections 3.4.3.2 and 3.4.3.4, we present attacks that allow to permanently disconnect touchlink-enabled devices from their legitimate network and to trigger a manufacturer-specific identify action for a duration of up to 18 hours. For example, in case of light bulbs, this action is blinking but

for other devices this can also be making sounds or moving. The attack proceeds as follows. First, the attacker permanently disconnects all touchlink-enabled devices, and then makes them blink, beep, or move for several hours. Trying to extort money, the attacker promises to stop the attacks and to reconnect all affected devices to the legitimate network as soon as the ransom money is paid. The residents can decide whether they want to manually reset and recommission each device or to pay the demanded ransom. Since the recovery of Zigbee devices from the permanent disconnect attack can be an extremely cumbersome task that requires dexterity and precise timing, as described in Section 3.4.6, residents might prefer to pay the ransom.

3.4 Security Analysis of the Touchlink Commissioning Procedure

We divide our attacks into passive and active attacks. Passive attacks only eavesdrop on the wireless communication of nodes in the targeted network, while active attacks require the attacker to interact with the targeted node via wireless communication. Additionally, we also categorize our attacks according to the attacker's goal. In the first category, we describe DoS attacks that exploit security weaknesses in the concept of so-called inter-PAN frames. These attacks require no knowledge of any cryptographic material. In the second category, we show attacks that allow the attacker to control devices in the network. These attacks require knowledge of the touchlink preconfigured link key that has been leaked in March 2015 [RSWO17]. In Table 3.1, we provide an overview of all attacks that are described in this section.

Table 3.1: Overview of attacks.

Attack	Type of Attack		Attacker Goals	
	Active	Passive	DoS	Control
Identify Action	●	○	●	○
Reset to Factory-New	●	○	●	○
Permanent Disconnect	●	○	●	○
Hijack	●	○	●	●
Network Key Extraction	○	●	○	●

The attacks in this section outline the procedures to compromise a single touchlink-enabled device. All attacks can be easily extended to target multiple devices at the same time by running the attack procedures for different target devices simultaneously.

3.4.1 Penetration Testing Framework Z3sec

For our research, we developed the penetration testing framework Z3sec in Python to evaluate the security of Zigbee 3.0 devices. These tools and their documentation are available as open-source software on GitHub[3]. The Z3sec framework consists of three major components. First, a *touchlink library* to build arbitrary touchlink packets and to keep track of source addresses and sequence numbers. Second, a *crypto module* that provides the functionality to encrypt

[3] https://github.com/IoTsec/Z3sec

(a) Philips Hue **(b)** Osram Lightify **(c)** GE Link **(d)** IKEA Trådfri

Figure 3.3: The four evaluated Zigbee-certified connected lighting systems.

and decrypt Zigbee packets. This component also handles key transport frames, especially decrypting the encrypted network key, and vice versa. Third, the *radio interface module* enables the communication between the radio transceivers and the touchlink library.

As radio transceiver, we utilize the Universal Software Radio Peripheral (USRP) B200 from Ettus, a Software-Defined Radio (SDR) covering the radio-frequency range between 70 MHz and 6 GHz. The USRP features a field-programmable gate array and connects to a host computer via USB 3.0. We use Scapy-radio [PLD14] as interface to send and receive Zigbee packets with the USRP. Scapy-radio itself uses capabilities of GNU Radio and an IEEE 802.15.4 GNU Radio flowchart implementation [BLDS13].

In addition, we implemented a command line tool that performs the attack procedures described in the following sections, and a module that is able to send or spoof control commands to the devices once the network key is disclosed.

3.4.2 Testbed

In total, four different systems were evaluated, depicted in Figure 3.3. More precisely, we analyzed a Philips Hue starter set including one bridge and three white and color ambiance (LCT001) LED bulbs, exemplarily shown in Figure 3.3a. Furthermore, we deployed an Osram Lightify gateway with a classic A60 tunable white LED bulb as depicted in Figure 3.3b. Our third system is a GE Link starter pack containing a Link hub and a Link A19 soft-white LED light bulb as shown in Figure 3.3c. The last system is the IKEA Trådfri system consisting of a Trådfri LED 980 lm light bulb and a Trådfri remote control as depicted in Figure 3.3d. All these systems implement the ZLL standard, for which touchlink commissioning is mandatory (cf. Section 3.2). Thus, these systems are representative for arbitrary touchlink-enabled Zigbee devices. Due to the novelty of the Zigbee 3.0 specifications, only one manufacturer released Zigbee 3.0-certified products so far to the best of our knowledge (as of March 2017). These released Zigbee 3.0-certified products by Ubisys do not support touchlink commissioning. Nevertheless, the presented attacks apply to all future Zigbee 3.0 products that enable touchlink commissioning.

Before starting our evaluation, we updated the Philips Hue firmware to the then-latest version 01031131 as well as the API to version 1.12.0. We updated the Osram Lightify gateway WLAN to version 1.1.2.101 and the gateway Zigbee to version 1.2.0.67. For the GE Link firmware, we could not find any possibility to update it by using the manufacturer-recommended Wink app. At the time of the evaluation, there existed neither a mobile device app for IKEA Trådfri nor a possibility to update the firmware of the Trådfri bulbs or remote control[4].

The attacker equipment comprises a laptop on which our penetration testing framework

[4]IKEA released a Trådfri gateway and a Trådfri mobile device app in April 2017.

Z3sec is installed. A radio transceiver, the Ettus USRP, is connected to the laptop. We started the evaluation of each attack with the default settings, in which the lighting system works as intended and the system is not compromised.

3.4.3 Denial-of-Service Attacks

Our DoS attacks exploit the concept of inter-PAN frames, which are a special type of Zigbee frames that allow the communication between different PANs. In 2008, inter-PAN frames were introduced in the Zigbee Smart Energy application profile. In the purpose description, the Zigbee Smart Energy standard states that inter-PAN transmissions allow Zigbee devices to "perform limited, insecure, and possibly anonymous exchange of information" [Zig12b, Section B.1]. In the context of smart metering, for which the Zigbee Smart Energy standard was intended, the mandate for such a transmission mechanism is the "market requirement to send pricing information to very low-cost devices" [Zig12b, Section B.1], e. g., refrigerator magnets showing the current energy consumption or prices. The ZLL standard adopted the inter-PAN transmission mechanism to enable the commissioning of networks with constraint devices, e. g., remote controls. In the touchlink commissioning procedure, inter-PAN frames are used to transmit touchlink commands and their responses between initiator and target device.

Since there exists no shared key material between different PANs, inter-PAN frames are neither secured nor authenticated. Hence, all attacks presented in this section are performed *without* requiring any knowledge of the touchlink preconfigured link key or of any other cryptographic (key) material relating to these devices.

3.4.3.1 Active Device Scan

The active device scan searches for Zigbee devices in the wireless range of the attacker's equipment. It is a mandatory step in the preparation of any further attack.

Procedure. The attacker builds a `scan request`, then sends this inter-PAN command frame on all Zigbee channels consecutively and, on each channel, listens for a few milliseconds for any `scan response`. Through the reception of `scan response` frames, the attacker learns about all Zigbee devices that are also listening on this channel. Zigbee uses 16 channels in the 2.4 GHz Industrial, Scientific, and Medical (ISM) band, designated channels 11 to 26, while channels 1 to 10 are located in other ISM bands. The Zigbee 3.0 specifications define four primary channels on which devices are listening for touchlink scans: 11, 15, 20, and 25. These channels are used for commissioning and normal operations, while all remaining channels can be used as backup.

Evaluation. The active device scan works with all four lighting systems. In general, all light bulbs responded to the `scan request`, while also the Osram Lightify gateway answers each time. The GE Link hub does not respond to the `scan request`, and the Philips Hue bridge only replies if the button on the hub was pushed within the previous 30 seconds.

ACK Spoofing. All following attacks start with an active device scan and then send further inter-PAN command frames. When target devices with Osram Lightify or GE Link implementations receive a `scan request`, they answer with the `scan response` and then

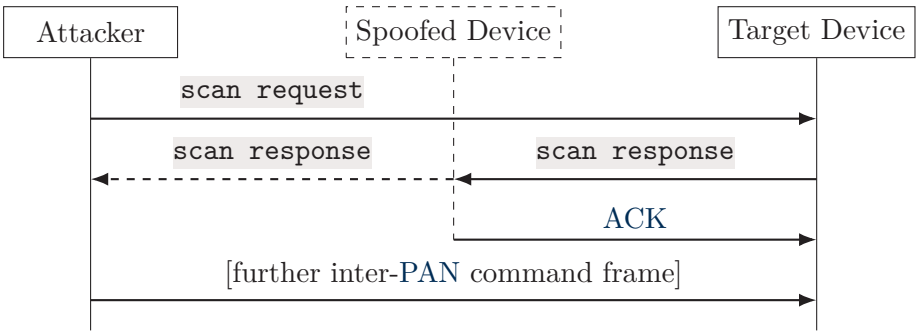

Figure 3.4: Acknowledgment spoofing: the attacker spoofs a third device to send an ACK to the target device timely.

wait for a MAC Acknowledgment (ACK). If they do not receive an ACK within a specified time frame of 864 µs, they drop further communication. Z3sec is too slow in sending this ACK since the processing of the received Zigbee frames is not performed on the hardware platform but in software, and therefore delayed for a few milliseconds. However, we can impersonate an existing Zigbee device, referred to as spoofed device in Figure 3.4, by setting the extended source address of the `scan request` to the extended address of the spoofed device. As a result, the spoofed device sends an ACK upon reception of a `scan response`, even if this device had never sent any `scan request`. This is an inherent feature of the IEEE 802.15.4 MAC, and we can leverage this mechanism for "providing" ACKs in a timely manner. In contrast, Philips Hue and IKEA Trådfri devices do not require any ACK to a `scan response`, making our attacks easier to implement.

3.4.3.2 Identify Action Attack

The touchlink commissioning procedure provides the possibility to request a Zigbee device to identify itself via a pre-defined identify action, e. g., blinking, dimming, or beeping. Originally, the identify procedure is intended to give the user a possibility to select and identify a certain node, which should be added to the network, but the identify action can also be abused by attackers.

Procedure. After an active device scan, the attacker can send an `identify request` to the targeted Zigbee device. The `identify request` contains the `transaction identifier` and `identify duration`. The identify duration can be at maximum 0xFFFE, which converts to a time duration of 18 hours 12 minutes 14 seconds. If the identify duration is set to 0, a previously started identify procedure is aborted before the specified duration elapsed. Setting the identify duration to 0xFFFF requests the device to perform the identify procedure for a device-specific default period of time, usually a few seconds. At the reception of the `identify request`, the targeted Zigbee device starts its identify action for the defined period of time.

Evaluation. All lighting systems are vulnerable to the identify action attack. During the blinking of the lights, the users can neither turn off nor control the light bulb using the apps provided by the manufacturers. The only way to shut down the lights is to physically disconnect the bulb from the power supply. An exception are IKEA Trådfri bulbs, which dim

Table 3.2: Maximum duration (in hours) of the identify action attack for each evaluated lighting system.

System	Max. duration
Philips Hue	18:12:14
Osram Lightify	9:12:53
GE Link	9:06:31
IKEA Trådfri	0:01:00

their lights up and down instead of blinking. Also, the identify action can be immediately aborted by pressing an arbitrary button on the remote control.

The attacker can abort the attack anytime by sending another `identify request` with the field duration set to zero. The maximum duration of blinking that can be triggered with a single `identify request` is shown in Table 3.2. We assume that the duration depends on the manufacturer's implementation of touchlink commissioning.

After performing the identify action attack, the Philips Hue bulb and the IKEA Trådfri bulb return to the pre-attack state and color, while the Osram Lightify bulb and the GE Link bulb change to the default state and color. This attack also works if the device is turned off but supplied with power.

3.4.3.3 Reset to Factory-New Attack

In this attack, the attacker resets the configuration of a Zigbee device to the factory-new state.

Procedure. To perform the attack, a `reset to factory new request` inter-PAN command frame is sent after a prior active device scan. The payload of the `reset to factory new request` only contains the `transaction identifier`. On the reception of a valid `reset to factory new request`, the light bulb discards the current configuration. The color and brightness of the light bulb changes to the default states.

Evaluation. Our evaluation showed that all four lighting systems are vulnerable to the reset to factory-new attack. Interestingly, we are also able to reset the Lightify gateway (at any time) as well as the Philips Hue bridge (if the button of the bridge was pushed within the last 30 seconds) to a factory-new state. After a reset to factory-new attack, the legitimate user would have to reintegrate the bulb into the legitimate network by either searching for new devices via the mobile device app or by using a remote control. This operation has to be initiated manually by the user. In the meantime, i. e., before the user initiates a recommissioning, an attacker has the chance to hijack the reset device using the classical commissioning in ZLL or EZ-Mode commissioning in Zigbee 3.0, each in combination with the publicly known default global Trust Center link key as demonstrated by Ronen et al. [RSWO17] and Zillner and Strobl [ZS15].

3.4.3.4 Permanent Disconnect Attack

In the permanent disconnect attack, the user loses control over the touchlink-enabled device. This attack differs from the reset to factory-new attack in the process of recovery: after a reset to factory-new attack, the user can simply recommission the attacked bulb to the network again. In the aftermath of a permanent disconnect attack, the user needs to recover the bulb first, as we describe in Section 3.4.6, before a recommissioning to the legitimate network is possible again.

Procedure. We present two approaches to perform a permanent disconnect attack. In the first approach, we force a targeted Zigbee device to change the current channel to another channel determined by the attacker. In the second approach, we join the targeted device to a non-existent network.

A change of the wireless channel can be enforced by sending a `network update request` inter-PAN command frame. The command must include a `network update identifier` that is higher than the current update identifier of the targeted network, which is a counter that is incremented each time the network settings are updated. The current `network update identifier` can be retrieved from the `scan response` of the target device. After receiving the `network update request`, which includes the new channel, the target device switches to this channel. The legitimate network does not recognize the shift. As a consequence, the targeted device does not receive legitimate user commands anymore.

Using a `network join end device request` inter-PAN command frame (instead of the `network update request`), an attacker can manipulate additional network settings like the PAN ID and the current network key. The attacker sets the encrypted network key field to an arbitrary 128-bit value and then sends the `network join end device request` to the targeted Zigbee device. On reception of a `network join end device request`, the Zigbee device leaves its current network and sets the internal parameters according to the new configuration. Since the encryption of the network key is not authenticated, the device decrypts the arbitrary 128-bit value to a garbage network key, which also is not known to the attacker. The transaction is confirmed by sending a `network join end device response`.

Evaluation. In the evaluation, all four presented lighting systems are vulnerable to both permanent disconnect attacks. These attacks do neither change the color nor the state of the bulb but after performing the attack procedures, the targeted bulbs cannot be controlled by the legitimate user anymore.

3.4.4 Attacks to Gain Control

The authenticity and integrity of the Zigbee touchlink commissioning procedure relies on the *touchlink preconfigured link key*, also denoted as *ZLL master key* in the ZLL specifications, that is used to encrypt the current network key before this key is transmitted to the joining device. The procedure of the network key encryption starts by expanding the `transaction identifier` and the `response identifier` from `scan request` and `scan response`, respectively, to a 128-bit string. This bit string is the input to the AES encryption function, while the touchlink preconfigured link key is used as encryption key. The resulting output is denoted as transport key. In the next step, the actual network key is encrypted with the transport key using AES encryption again.

The touchlink preconfigured link key is distributed to manufacturers of Zigbee-certified products under an NDA. However, in March 2015, the touchlink preconfigured link key was leaked on Twitter[5]. In the following, we present attack procedures in which the knowledge of the touchlink preconfigured link key is facilitated to take full control over Zigbee networks.

3.4.4.1 Hijack Attack

The hijack attack extends the permanent disconnect attack described in Section 3.4.3.4. Instead of sending arbitrary bytes as the encrypted network key, the attacker forces the Zigbee device to use an attacker-chosen network key.

Procedure. Again, the attacker builds the `network join end device request` inter-PAN command frame as described in Section 3.4.3.4. The attacker-chosen network key is encrypted using the leaked touchlink preconfigured link key, the `transaction identifier` from the `scan request` and the `response identifier` from the `scan response` of the targeted device. This encrypted network key is embedded in a `network join end device request` and sent to the targeted device. On reception of a `network join end device request`, the device updates its internal parameters according to the received values and confirms the transaction by sending a `network join end device response`. The targeted device is now commissioned to the network of the attacker, who has full control over this device.

Evaluation. In the evaluation, we were able to force Zigbee devices of all four connected lighting systems to accept an attacker-chosen network key. This attack paves the way to send further application-specific commands to the targeted devices.

3.4.4.2 Network Key Extraction

An attacker is able to extract the current network key by eavesdropping on the `scan response` and the `network join end device request`[6] of an initial touchlink commissioning. All these command frames must belong to the same transaction, i. e., contain the same `transaction identifier`.

Procedure. The legitimate user can be motivated to perform a touchlink commissioning procedure as a result of a prior reset to factory-new attack (cf. Section 3.4.3.3). Then, the user is forced to commission the node to the legitimate network again. After eavesdropping on the encrypted network key from the `network join end device request`, the network key is decrypted using the leaked touchlink preconfigured link key. The `response identifier` is known from the `scan response`, while the `transaction identifier` is included in all packets belonging to the same transaction.

Evaluation. For this attack, the legitimate user of a connected lighting system has to perform the touchlink commissioning procedure. Our investigations conclude that the Philips Hue and the IKEA Trådfri lighting systems can be targeted with this attack since only a few

[5] https://x.com/mayazigbee

[6] Instead of capturing the `network join end device request`, this attack can also be performed by capturing a `network join router request` or a `network start request`.

(a) Evaluated light bulbs. **(b)** Attacker equipment.

Figure 3.5: Outdoor testbed to measure the maximum distance of successfully attacking a Zigbee-certified device.

Philips Hue third-party apps as well as the IKEA Trådfri remote control trigger touchlink commissioning. To the best of our knowledge, there exist neither apps nor Zigbee-certified devices by Osram or GE that can initiate the touchlink commissioning procedure.

In our evaluation, we showed that all four lighting systems can be controlled once the network key is exposed. We were able to send commands to turn the bulbs on and off and to change the light color of the Philips Hue bulbs to any arbitrary color.

3.4.5 Evaluation of Wireless Range

In the Zigbee specification, the manufacturers are advised to limit the wireless range of touchlink commands such that only Zigbee devices in close proximity are able to perform the touchlink commissioning procedure. This limitation, denoted as *proximity check*, should be implemented in a way that the received signal strength of an initiator device must be above a certain threshold.

To get a baseline for our attacks, we measured the maximum distance to successfully perform the touchlink commissioning procedure with a Philips Hue bridge since the Osram Lightify gateway and the GE Link hub provide no possibility to trigger the touchlink commissioning procedure. The maximum distances to successfully commission a Hue bulb or a Link bulb to the Hue bridge are 1.8 meters, and 1.6 meters for a Lightify bulb. For IKEA Trådfri, the maximum distance to trigger the touchlink commissioning procedure between the remote control and a bulb are 1.5 meters.

Since the touchlink commissioning procedure is intended to require close proximity, we investigated whether the attacks work for longer distances. We set up an outdoor testbed on a sports ground, in which a line-of-sight between the USRP and the attacked bulbs was given. At the USRP, we mounted rod antennas with 8 dB gain according to the manufacturer. The setup is shown in Figure 3.5. We decided on an outdoor measurement with line-of-sight between the attack equipment and the target device since it is hard to generalize statements about the wireless range of these attacks through buildings. The propagation of radio waves depends on many variables, e. g., the ground plan (reflections), structure and thickness of walls, wall openings, electrical installations, as well as interference with other deployed wireless networks. Thus, we conducted outdoor measurements eliminating as many types of distortions as possible and thus providing *ground truth* for further investigations.

Table 3.3: Maximum range (in meters) for an active attack in our evaluation for each evaluated lighting system.

System	Legitimate Range	Active Attack Range	
		Regular	Disclosed bug
Philips Hue	1.8	36	(patched)
Osram Lightify	1.6	15	> 190
GE Link	1.8	28	> 190
IKEA Trådfri	1.5	> 190	> 190

3.4.5.1 Active Attacks

In the evaluation of the active attacks, we measured the maximum distance from which we are able to trigger an identify action attack, which is described in Section 3.4.3.2. The identify action attack requires to pass the proximity check and to receive a response from the targeted node, and therefore, this attack is representative for active attacks. The results are shown in Table 3.3. The maximum distance of successfully attacking the Osram Lightify system is 15 meters, the maximum distance of the GE Link system is 28 meters and the maximum distance of the Philips Hue system is 36 meters. These distances depend on the current noise on the channel as well as the orientation of the bulbs and the antennas of the USRP. We experimented with different gain and antenna settings and also with different positions and directions of the bulbs. From our measurement results, we estimate that the received signal strength of inter-PAN command frames has to be stronger than −40 dBm.

We also measured the maximum distance to perform the same attack on IKEA Trådfri bulbs. In our evaluation, we were able to trigger the identify action attack from a distance of more than 190 meters. This was the maximum measurable distance due to space restrictions of the outdoor testbed. Since the range to actively attack an IKEA Trådfri bulb is much larger compared to the other systems, we assume that the proximity check is not enforced. For all these tests, the `transaction identifier` was chosen randomly.

Ronen et al. [RSWO17] disclosed a bug in the ZLL implementation of Philips Hue, in which the proximity check can be circumvented by setting the `transaction identifier` of the touchlink command to zero. This bug was patched by Philips in October 2016. In our evaluation, we confirmed that the same bug also affects Osram Lightify, GE Link, and IKEA Trådfri. We were able to trigger the identify action attack from the maximal measurable distance of 190 meters for each of these systems. Again, we assume that this distance can be enlarged.

3.4.5.2 Passive Attacks

We also evaluated the ranges of passively eavesdropping on the touchlink commissioning procedure to extract the current network key. We were able to extract the network key, that was established using an IKEA Trådfri system consisting of a bulb and a remote control, from a distance of 42 meters. In addition, we extracted the network key, established between a Philips Hue bridge and a Philips Hue bulb, from a distance of 130 meters. We did not evaluate Osram Lightify or GE Link systems since they do not provide an interface for consumers to trigger the touchlink commissioning procedure.

3.4.6 Recovery

The recovery from an identify action attack or a reset to factory-new attack is easy to accomplish. For the identify action attack (cf. Section 3.4.3.2), the user has to disconnect the node from the power source and reconnect again. For the reset to factory-new attack (cf. Section 3.4.3.3), the user needs to rejoin the node to the network.

Recovery from permanent disconnect and hijack attacks (cf. Sections 3.4.3.4 and 3.4.4.1), which alter the configuration of the nodes, is more challenging but all evaluated lighting systems possess functions to regain control over the attacked devices. However, these procedures are not obvious at first sight. In any case, a recovery entails manual effort for the user, as we explain below.

For GE Link, Osram Lightify, and IKEA Trådfri, the only way to recover the attacked bulbs is a physical reset. The physical reset is not specified in the ZLL standard[7], but can be achieved by powering the bulbs on and off in a certain manufacturer-specific pattern, e. g., an Osram Lightify A60 bulb must be turned on 3 seconds, off 5 seconds, and this procedure must be repeated five times. This is a cumbersome task and might not be successful at the first try. Since the physical reset mechanisms are not obvious, either a mobile device app (GE Link, Osram Lightify) or a manual (IKEA Trådfri) guides the process of performing a physical reset. The Philips Hue system lacks a physical reset, to the best of our knowledge. However, the Hue system supports an additional commissioning mechanism, *manual search*, which is not specified in any Zigbee standard. Manual search works by entering a code that is printed on a bulb into the Hue app. The manual search fails if the channel of the attacked device was altered to a secondary channel (cf. Section 3.4.3.1).

Touchlink commissioning can be applied as an alternative recovery procedure for Philips Hue and IKEA Trådfri bulbs. However, the IKEA Trådfri remote control does not search on secondary channels, therefore, if the channel of a bulb was changed to a secondary channel, it must be reset before recommissioning. In Philips Hue, touchlink commissioning can be performed by using either the Hue API debug tool or a third-party app. After the recommissioning with touchlink, an interesting effect can be observed. Instead of reintegrating the attacked bulb into the former network, the Philips Hue bridge detects that the `network update identifier` of the discovered device is higher than its own. The Hue bridge adapts to these "newer" network settings and switches to the attacker-defined channel. Consequently, the bridge loses the connection to all other bulbs, which remain on the former channel. Afterwards, all other bulbs of the former network have to be recommissioned to the new network using touchlink. This is a time-consuming task because all devices have to be moved in close proximity (1-2 meters) to the Hue bridge in order to perform the touchlink commissioning.

3.5 Disclosure and Response

We reported the results of our security analysis to GE, IKEA, Philips, Osram, and the Zigbee Alliance. The manufacturers IKEA, Philips, and Osram responded to our outreach. In contrast, the GE Product Security Incident Response Team confirmed the reception of our notification but did not comment on our report. We discussed the results of our security

[7] The Zigbee 3.0 specifications recommend supporting a reset via local action in a manufacturer-specific fashion.

analysis with representatives and members of the Zigbee Alliance and received their feedback that the analysis is accurate and complete. According to the Zigbee Alliance, the main reason for the development of touchlink was to reduce the complexity of commissioning procedures for Zigbee Pro devices. Touchlink offered a low entry level for the consumers to set up connected lighting systems that are configured via remote controls. However, at the same time as the ZLL standard was developed (2010–2012), the popularity of smartphones increased rapidly. Because of this development, the Zigbee Alliance decided to introduce a bridge device that translates TCP/IP traffic sent by smartphones into Zigbee commands, thus simplifying the handling of the classical Zigbee commissioning procedure and providing an alternative method to overcome the complexity of commissioning.

3.6 Discussion

Summarizing the results of the security analysis, we can see that all tested Zigbee-certified products are susceptible to passive and active attacks. An attacker is able to thwart the availability and can take complete control over any touchlink-enabled device. It is irrelevant whether the targeted Zigbee network was initially set up using the touchlink or another commissioning procedure. Furthermore, we showed that close proximity is not required. In our evaluation, we successfully performed active attacks from a distance between 15 and 190 meters depending on the targeted product. Also, we were able to (passively) eavesdrop on the touchlink commissioning procedure from distances between 42 and 130 meters. We assume that these distances could be further extended if the attacker uses directional antennas. Four different Zigbee-certified connected lighting systems facilitating the ZLL standard have been tested by us. Since the touchlink commissioning procedure in Zigbee 3.0 has not been changed compared to ZLL, all presented attacks also apply to arbitrary Zigbee 3.0 products that enable touchlink commissioning. In summary, we state that all three threat scenarios outlined in Section 3.3.2 are realistic and exploit security weaknesses that exist *by design*. In addition, we explored that the recovery of attacked devices is quite a cumbersome task.

3.6.1 Usage of Touchlink

Since touchlink commissioning is an optional feature in Zigbee 3.0, we recommend disabling this commissioning option in all future Zigbee products. Already a *single* touchlink-enabled device in the network can expose the network key and thus lead to insecurities of other nodes. In our communication with the Zigbee Alliance, they suggested putting the enabling of the touchlink features under application control, for example to enable touchlink only a few minutes after power-up. Although this restriction limits the vulnerability time frame, the users can be motivated by social engineering techniques to power up devices at predictable times. For example, jamming of Zigbee communication may motivate the consumers to disconnect a device from the power source and power it up again. Furthermore, the recommendation of putting touchlink commissioning under application control is not included in the specifications and so, it is not quite clear how the manufacturers should become aware of this.

3.6.2 Manufacturer-Specific Mitigation

No immediate mitigation of the attacks presented in Section 3.4 is possible since the security weaknesses result from legitimate features in the specification, especially from the concept of unauthenticated inter-PAN frames. If touchlink commissioning is required, manufacturer-specific changes can be made to contain the effects of the attacks. For example, the identify action should be limited to a reasonable duration (like in the IKEA Trådfri implementation), which would decrease the impact of the identify action attack significantly.

3.6.3 Renewal of the Touchlink Preconfigured Link Key

Since the touchlink preconfigured link key was leaked in March 2015, the touchlink commissioning procedure is considered compromised. On the one hand, the replacement of the touchlink preconfigured link key would circumvent the attacks presented in Section 3.4.4, and therefore protect against take-over attacks, in which the attacker gains control of the targeted devices. On the other hand, the renewal of this key would render the integration of ZLL-based connected lighting systems and complementary equipment into Zigbee 3.0 networks impossible. This would most likely lead to public resentment, as the following incident illustrates. In December 2015, an update of the Hue app by Philips locked out light bulbs of other vendors if these vendors did not participate in the so-called *Friends of Hue* certification program. The public outcry made Philips revert this decision after a few days through providing a non-scheduled update [Phi15]. In addition to compatibility problems, the non-disclosure of a renewed key cannot be guaranteed since the current touchlink preconfigured link key was also protected by an NDA but leaked anyway.

3.6.4 EZ-Mode Commissioning

EZ-Mode commissioning is an alternative commissioning procedure supported by all Zigbee 3.0 products. This commissioning procedure neither relies on inter-PAN frames nor a proximity check. Therefore, the attack procedures described in Section 3.4 cannot be adapted. EZ-Mode commissioning offers three options of securing the network key transport. In centralized security networks, the network key transport can be either protected by the publicly known default global Trust Center link key, or a link key derived from an install code. In distributed security networks, the network key transport is encrypted using the NDA-protected distributed security global link key.

All three EZ-Mode commissioning options have serious drawbacks in terms of security and usability. Using the default global Trust Center link key makes the networks susceptible to attack scenarios where users are forced (e. g., by jamming) to recommission a node to the network in presence of the attacker. In this case, the attacker would be able to recover the network key through eavesdropping on the commissioning procedure. Using the NDA-protected distributed security global link key is only secure as long as the key is not leaked. Thus, in the long run, the install code option is the only EZ-Mode commissioning option that is secure against the local attacker model described in Section 3.3.2. We note, however, that this option requires an extra effort from the users in terms of scanning (or entering) the install code using a mobile device app. This might constitute a serious usability problem, especially if already installed devices have to be recommissioned. For example, this can happen if a broken coordinator node has to be replaced. If the devices are difficult to reach, e. g., bulbs installed on the ceiling, the recommissioning might become quite cumbersome.

3.6.5 Limitations of Our Work

The Zigbee 3.0 specification warns about supporting the touchlink command to join an end device to an existing centralized security network [Zig16a, p. 64], which means that the hijack attack and network key extraction possibly cannot be performed on *centralized security networks*, if manufacturers heed this warning. However, these attacks work for distributed security networks. All other attacks work for both security network models.

3.7 Related Work

The security of the Zigbee specification as well as the underlying IEEE 802.15.4 standard attracted much less attention in the academic research community compared to other wireless standards, such as WLAN [IEE21], Bluetooth [Blu16], or mobile telephony like 4G [3GP16].

Sastry and Wagner [SW04] analyzed the security mechanisms of the IEEE 802.15.4 protocol. However, these mechanisms are not used in Zigbee. Wright [Wri09] published the penetration testing tool KillerBee, which allows sniffing and analyzing traffic of Zigbee and other IEEE 802.15.4-based networks. Wright also exposed that the network key of the then-current Zigbee standard was sent in clear text OTA. He demonstrated successful replay attacks using previously captured Zigbee traffic. The Zigbee Pro specification, released in 2012, addressed these security weaknesses. Goodspeed et al. [GBM+12] developed exploration tools to analyze the wireless attack surface of IEEE 802.15.4 networks. Müller et al. [MABM16] present a formal security model for the Zigbee touchlink commissioning. Further papers by Olawumi et al. [OHA+14] and Vidgren et al. [VHP+13] cover security issues of Zigbee networks, but these papers refer to security weaknesses concerning outdated Zigbee specifications and have not been evaluated with Zigbee-certified products.

Since the emergence of connected lighting systems in 2012, these systems have been subject to a number of security investigations. Dhanjani [Dha13] published implementation weaknesses of the command authentication in the Philips Hue lighting system. He discovered that the secret whitelist token, which is required to authenticate the commands sent from the app (or website) to the bridge, is a hash of the MAC address of the controlling device. Chapman [Cha14] obtained the firmware of LIFX light bulbs via a JTAG debugger and extracted cryptographic key material through reverse engineering of the firmware. Heiland [Hei16] exposed vulnerabilities in the Osram Lightify system. Through reverse-engineering, he discovered that WLAN credentials are stored in plaintext in the iOS Lightify Home app. Zillner and Strobl [ZS15] exposed security weaknesses in the Zigbee Pro specification. They showed that Zigbee-based lighting systems use publicly known fallback keys in the classical commissioning procedure for the initial key exchange, which allows the extraction of the network key. Ronen and Shamir [RS16] used the Philips Lux lighting system, which is the white-color variant of Hue, to build a covert channel for the exfiltration of data from an isolated environment. Also, Ronen et al. [RSWO17] exploited an implementation bug in Philips Hue bulbs that allowed to reset and to control these bulbs from a distance of a few hundred meters. In addition, they extracted cryptographic material, which secured the update process of the bulb's firmware, from the hardware using correlation power analysis. As a result, Ronen et al. were able to install a manipulated firmware image on Hue bulbs, and discussed the threat of a self-spreading IoT worm.

In contrast to related work on connected lighting systems, we do not analyze the products of a certain manufacturer but investigate the underlying standard. Also, and to the best of

our knowledge, we are the first to investigate the security mechanisms of the latest Zigbee standard, Zigbee 3.0.

3.8 Conclusion

Millions of IoT devices, including security-critical products that should be secured against local attackers, such as door locks and intruder alarm systems, use Zigbee for wireless low-power communication. In this chapter, we investigated the touchlink commissioning procedure, which is a commissioning option in the latest Zigbee specification, Zigbee 3.0. We performed a security analysis of the touchlink commissioning procedure, described active and passive attacks, and evaluated their impact using our penetration testing framework Z3sec.

Our results show that attackers can thwart the availability of touchlink-enabled devices and gain control over all nodes in the network. A single touchlink-enabled Zigbee device suffices to expose the network key to an attacker, and thus compromises the security of all nodes in the network, no matter how these nodes were originally added to the network. Therefore, we warn against the adoption of touchlink commissioning in all future Zigbee 3.0 devices. To prevent these attacks, we recommend manufacturers of Zigbee-certified products use EZ-Mode commissioning in combination with install codes.

Future work is required to analyze the security of EZ-Mode commissioning in more depth. In addition, the design of a robust and suitable authentication infrastructure for IoT networks that interconnects function-constrained low-power nodes of multiple manufacturers without initial trust but with high usability for non-expert users, is still a future challenge.

Because I invent, transform, create, and destroy for a living, and when I don't like something about the world, I change it.

Rick (Rick and Morty)

4 A Cryptographic Model for Wireless Communication Is Never Late

Abstract

The previous chapter showed that established and widespread communication standards may be susceptible to attacks, probably due to a lack of appropriate communication or attacker models as well as misconceptions about the nature of wireless communication and attackers' capabilities. It is unsurprising that considering models designed for wired communication may either lead to unrealistic requirements in a wireless setting or not be applied at all. Since IoT devices or, in a more general sense, lightweight devices communicating in a multi-party, wireless setting are continuing to proliferate, having adequate models for analyses at one's disposal is becoming more and more important.

In multi-party communication scenarios, the MitMM is a widely used model for formally representing attackers. It assumes that an attacker has unrestricted control over the communication, including the ability to intercept, block, and modify messages at her discretion. However, we show that the MitMM is too strong for accurately capturing the peculiarities of wireless communication. Consequently, the MitMM cannot accommodate security mechanisms that leverage unique wireless properties, such as friendly jamming or DB protocols. Moreover, the absence of a specialized model hinders the design and formal analysis of new wireless security schemes.

To address these limitations, we introduce a novel formal model tailored to the wireless domain, dubbed the WCM. Building on the MitMM, our model incorporates three key elements missing in the traditional MitMM — communication channels, signals, and locality. These aspects allow for a more realistic representation of wireless communication. To illustrate the model's utility, we provide a concrete example of representing Quadrature Phase-Shift Keying (QPSK) modulation within our model. With the WCM as a basis and to formally capture the constraints faced by attackers in wireless environments, we define an attacker model, dubbed the WAM.

Using our model, we formalize friendly jamming and propose a new security notion, IND-CSA. We further demonstrate the utility of our model by providing a formal security proof for DB protocols.

Declaration of Origin: This chapter is based on the paper *Towards a Cryptographic Model for Wireless Communication: Extended Version* [AM25], written together with *Frederik Armknecht* which is an extended journal version of *Towards a Cryptographic Model for Wireless Communication* [AM24], written together with *Frederik Armknecht*, presented at *SECRYPT 2024*. First published in *Security and Cryptography*, Communications in Computer and Information Science, volume 2588, pp. 215–242, 2025 by Springer Nature. Reproduced with permission from Springer Nature.

4.1 Introduction

Security models are fundamental for unambiguously analyzing or proving the security of cryptographic schemes by formalizing security goals and attacker models. While the security goals may vary between different use cases, the attacker model is usually more universal. For typical scenarios involving two or more communicating parties, the so-called MitMM, introduced by Dolev and Yao [DY83], is commonly utilized. There, the attacker is assumed to have full control over the communication, e. g., message eavesdropping, delaying, or modification.

While we do not question the general usability or relevance of this model, this chapter is motivated by the observation that the MitMM can be too strong in the case of *wireless* communication. Indubitably, considering a very strong attacker model is often necessary and useful, i. e., if certain capabilities of an attacker cannot be excluded, then, assume the worst case; if a mechanism provides security against a strong attacker, then it also does so for weaker (and potentially more realistic) attackers.

For wireless communication, however, we claim that the situation is different. First, academic and industrial research has developed various security mechanisms which, regarding the security, cannot be adequately addressed in the MitMM. The probably best known example are DB protocols (e. g., Brands and Chaum [BC94] and Hancke and Kuhn [HK05]). But also other *existing* schemes cannot be analyzed adequately in the model, e. g., the shield for Implantable Medical Devices (IMDs) by Gollakota et al. [GHR+11], or other friendly jamming approaches (e. g., Berger et al. [BGF+14] and Shen et al. [SNHD13]).

Second, the development of new schemes is impeded, e. g., mechanisms may be dismissed as insecure in the MitMM, even though they could be secure *in practice*. More specifically, certain assumptions made in the MitMM regarding the attacker's capabilities often do not apply in wireless networks due to physical laws. For example, a transmitted message spreads like a wave in all directions from the sender. Typically, an attacker cannot affect the wave that is on the opposite side of the sender. Additionally, altering messages directly is not feasible, as the attacker can only send her own signals, which then interfere with the signals of other parties. For instance, Pöpper et al. [PTDC11] demonstrated that reliable and targeted manipulation of messages is challenging in practice within wireless communication networks. Similarly, Avoine et al. [ABG+21] criticize existing formal models for making impractical assumptions.

Given these shortcomings, we see the benefits of a model tailored to wireless communication as an alternative to the MitMM. This is particularly true given the ever-increasing relevance of wireless communication as, e. g., the IoT, with a forecast of 27.0 billion IoT devices by the end of 2025 [Has22], connects a plethora of devices in wireless networks and over the Internet. There is a strong demand for appropriate security measures to protect IoT devices and users against misuse and malicious behavior.

4.1.1 Contribution

Our key contributions are as follows:

- We identify several key properties of wireless networks that are not (or not correctly) covered by the MitMM and categorize these as *communication channels*, *signals*, and *locality*.

- Starting from the MitMM, we introduce a new model for wireless communication, dubbed WCM. Considering how the WCM may impact an attacker's capabilities, we develop the WAM.
- We demonstrate how WCM and WAM can be used to express QPSK, FHSS as well as for formalizing and analyzing several existing security mechanisms that could not be properly represented in the MitMM, namely DB protocols (including a formal security proof) and friendly jamming.

We hope the new model allows for founded and comparable security research in this important area and helps to develop novel security mechanisms that would be unthinkable in the traditional model.

4.1.2 Organization of This Chapter

The remainder of this chapter is organized as follows. Section 4.2 provides an introduction to wireless communication, with a focus on the IEEE 802.15.4 standard for Low-Rate Wireless Personal Area Networks (LR-WPANs), and jamming in the form of adversarial (reactive as well as selective) and friendly jamming. In Section 4.3, we identify and categorize key properties of wireless networks as *communication channels*, *signals*, and *locality*. Our new model for wireless communication, dubbed WCM, is introduced in Section 4.4. In Section 4.5, we discuss how the WCM may impact an attacker's capabilities, resulting in the WAM. We use the WCM and WAM to express FHSS as well as to formalize DB protocols and friendly jamming in Section 4.6. Section 4.7 summarizes related work and Section 4.8 concludes this chapter.

4.2 Background

For wireless communication, several standards coexist which target different ranges of device capabilities and communication distances, e. g., 5G [3GP22b], NB-IoT [3GP22a], LoRa and LoRaWAN [LoR20; LoR22], WLAN [IEE21], IEEE 802.15.4 [IEE11], or Bluetooth [Blu16]. Some are traditionally used in non-lightweight devices, e. g., 5G or WLAN, while others specifically target lightweight or IoT-devices, e. g., NB-IoT or IEEE 802.15.4 which, among others, forms the basis for ZLL as part of Zigbee 3.0 (cf. Section 3.2).

4.2.1 IEEE 802.15.4

The IEEE 802.15.4 standard [IEE11] defines the PHY and MAC for LR-WPANs and forms the basis for higher-layer specifications, e. g., Zigbee [zig17] (see also Section 3.2), WirelessHART [IEC16], or 6LoWPAN [RFC4944]. Besides for the 2.4 GHz band, the standard also defines communication mechanisms for several sub-GHz bands, e. g., the 780 MHz, 868 MHz, and 915 MHz bands. Depending on the specific band, it is divided into several non-overlapping communication channels, each with a dedicated center frequency. In addition to the available frequency bands, the specification includes different modulation schemes, e. g., Binary Phase-Shift Keying (BPSK), Offset Quadrature Phase-Shift Keying (OQPSK), or Chirp Spread Spectrum (CSS). However, not all possible combinations are considered valid. For instance, while BPSK is mandatory and OQPSK is optional in the 868 MHz band, in the 2.4 GHz band, BPSK is not available, OQPSK is mandatory and CSS

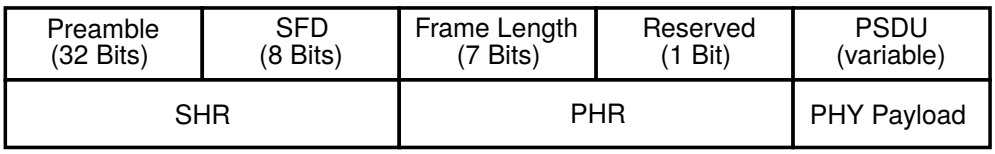

Figure 4.1: The PPDU structure according to IEEE 802.15.4 [IEE11].

is optional [IEE11]. Due to the proliferation, in this chapter, we focus on the 2.4 GHz band with OQPSK modulation.

Naturally, the IEEE 802.15.4 standard defines the PHY and MAC packet structures for the various combinations of frequency bands and modulation mechanisms. Figure 4.1 presents a PHY Protocol Data Unit (PPDU) for OQPSK in the 2.4 GHz band detailing the included header fields and their respective length.[1] The first two fields, i. e., preamble and Start Frame Delimiter (SFD), form the Synchronization Header (SHR) and contain fixed values to allow for receivers synchronizing with the sender and to identify the standard this transmission conforms to. Then, the PHY Header (PHR), essentially consisting of only the frame length, i. e., the length of the PHY payload, is included. The next higher layer information, i. e., MAC protocol data unit or MAC frame, is included as the PHY payload. Such a frame consists of MAC Header (MHR), MAC payload or MAC service data unit, and MAC Footer (MFR). While address information or, e. g., a sequence number, are part of the MHR, the MAC payload contains the frame payload, and the MFR contains the so-called Frame Check Sequence (FCS) which is essentially a Cyclic Redundancy Check (CRC) value computed over MHR and MAC payload.[2]

4.2.2 Adversarial Jamming

In general, opposed to unintentional interference like general activity by other communicating parties (not necessarily using the same communication protocol), jamming introduces *deliberate* interference in order to prevent reliable communication. For this, different strategies exist with differing energy requirements, activity levels, and risks for detection. The more basic forms, e. g., constant or so-called random jamming where the attacker jams regardless of actual messages being sent, require a lot of energy and bear a high detection risk, whereas more sophisticated strategies may be harder to detect, e. g., reactive jamming or selective jamming [GLY14; NN23; XTZW05; ZCLY17]. On the one hand, the constant jammer would keep the communication channel busy and, ultimately, hinders legitimate parties from sending messages at all. On the other hand, the random jammer interrupts her activity randomly, i. e., jams only some messages of all transmitted. Due to their sophistication, stealthiness, and damage potential, this Thesis focuses on reactive and selective jamming. For more in-depth information and other types of jamming, we refer the reader to Pirayesh and Zeng [PZ22], and Jameel et al. [JWKD19].

[1] Note that this structure generally applies to most of the other modulation and frequency bands defined by the standard, however, the fields' lengths may differ for other modulation techniques.

[2] Implicitly, the PHR (or more precisely, the frame length field) from the PHY is also part of the CRC as it controls at which offset within the received bytes the CRC value is to be expected.

4.2.2.1 Reactive Jamming

In contrast to the aforementioned basic variants, i. e., constant and random jamming, the *reactive jammer* is only minimally active, e. g., whenever a node starts sending data according to the attacked communication protocol, the jammer might interfere by sending a short signal leading to bit flips at the receiver. For instance, the jammer waits until an SHR is detected and sends another signal which overlaps with the original one, influencing some parts of the transmitted PPDU so that receivers discard it due to a non-matching CRC value. Usually, the jammer only needs to alter few bits to succeed if there are no error correction measures in place, e. g., as is the case for IEEE 802.15.4. The jammer might even directly target the CRC value only and leave the rest of the packet untouched [OFl11]. This allows the attacker to be minimally active as she only transmits for a short time. Comparable to constant jamming, the reactive jammer's goal is thus to impede *all* communication attempts within range, while, at the same time, being minimally active (in terms of actively transmitting).

4.2.2.2 Selective Jamming

If the attacker's decision to jam ongoing communication is based on additional criteria, e. g., a target party's network address, then, the attacker is a *selective jammer*, and thus, her decision is based on actual *content*. Compared to the reactive jamming strategy, the selective jammer's decision to interfere is hence based on actual transmission *content*. That is, the ongoing transmission is observed until at least the MHR is available as it contains information beyond the transmission's length, e. g., address information about sender and receiver. Suppose that the jammer targets a single network node that she would like to cut off, then, her jamming criterion could be to jam whenever the target's address is in the destination or source field of the MHR. Naturally, this leads to a delayed reaction compared to the reactive jammer. Nonetheless, this jammer also only needs to send signals that minimally overlap with the targeted transmission, thereby producing at least one bit flip so that receivers discard the transmission due to the non-matching CRC. With this strategy, the jammer is only actively transmitting whenever the *targeted* node is involved, which might lead to even lower sending activity compared to the reactive jammer.

4.2.3 Friendly Jamming

Interestingly, jamming could also be used as a defensive mechanism. This form of jamming is called *friendly jamming* and has been explored in several works, e. g., [BGF+14; BGF+16; GHR+11; JLS+22; JZJ22; JZZ21; LDW+20; LLD+22; SNHD13]. There, jamming is used to block unauthorized messages (*authorization*) or to hide the content of (potentially unencrypted) communication from illegitimate parties (*confidentiality*).[3] For instance, Gollakota et al. [GHR+11] exploit the proximity between two devices to mitigate eavesdropping and, furthermore, utilize jamming to achieve their goal.

That is, deliberately sending another signal whenever a transmission is going on renders the original message undecodable for other parties *except* for the — in this case — benign jammer, who, using the advantage of knowing the jamming signal, recovers the original message after all. Commonly, when discussing data privacy in communications, encryption is the primary method employed; however, friendly jamming may help in scenarios where such

[3] In this thesis, we focus on confidentiality only.

traditional methods might fail, e. g., a device only capable of unencrypted communication cannot be upgraded due to lack of any update mechanism or since it is placed in a location such that it is hard to physically access it, e. g., an IMD.

4.2.4 Generating Jamming Signals

An important aspect of jamming, be it adversarial or friendly, is how a jammer generates signals that effectively jam. Generally speaking, a jammer has basically three options. One option is to send randomly generated signals on the same carrier frequency, i. e., the jamming signal would not necessarily be correctly modulated regarding some protocol specification, e. g., when emitting white Gaussian noise. The second option is to send correctly modulated signals, i. e., the jammer follows the modulation as mandated by the targeted protocol's specification, whereas the third option is to send signals that are modulated according to a *different* modulation scheme than is used by the targeted protocol. However, when targeting QPSK or BPSK, the most effective method is to actually send correctly QPSK- or BPSK-modulated signals, respectively, according to Amuru and Buehrer [AB14], and to Han and Chen [HC21].

4.3 Motivation

This section motivates the need for a dedicated security model for wireless networks. First, in Section 4.3.1, we recall the established MitMM. Afterwards, we highlight shortcomings regarding peculiarities of wireless communication inherent to this model. We categorize these as communication channels (Section 4.3.2), signals (Section 4.3.3), and locality (Section 4.3.4). For each of these, we describe the situation in wireless networks, discuss why it is not represented in the MitMM and why it needs to be integrated, and summarize these findings as "lessons learned". The latter form the basis for our WCM shown in Section 4.4.

4.3.1 The Man-in-the-Middle Model

In the following, we recall the MitMM introduced by Dolev and Yao [DY83], using the formulation as given by Katz [Kat02]. Note that we are only interested in modeling the capabilities of an attacker to intercept and modify messages that are exchanged between communicating parties.

The system in the MitMM consists of a finite set of parties, denoted as Π, which are modeled as Interactive Turing Machines (ITMs). In a nutshell, an ITM enhances the traditional probabilistic Turing machine by incorporating, in addition to the work tape, random tape, and auxiliary tapes, a read-only communication-in tape and a write-only communication-out tape. These additional tapes enable the ITMs to receive messages from and send messages to one another, respectively. We refer to Katz [Kat02, Definition 2.1] for a full definition. The MitMM assumes that all communication between the parties in Π is under exclusive control of an attacker $\mathcal{A}$. This is captured by the notion of being *linked via* $\mathcal{A}$, formally defined by Katz [Kat02, Definition 2.2] as follows.

Definition 4.3.1: Linkage via $\mathcal{A}$

Let $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$ and $\mathcal{A}$ be ITMs. We say the parties in Π are *linked via* $\mathcal{A}$ if:

1. Each $\mathcal{P}_i$ has a single communication-out tape.
2. $\mathcal{A}$ has n communication-out tapes labeled $1, \ldots, n$.
3. The public-auxiliary-input tapes of $\mathcal{P}_1, \ldots, \mathcal{P}_n$ and $\mathcal{A}$ coincide.
4. The i-th communication-out tape of $\mathcal{A}$ coincides with the communication-in tape of $\mathcal{P}_i$. The communication-out tapes of each $\mathcal{P}_i$ coincide with the communication-in tape of $\mathcal{A}$.
5. All other tapes are distinct.

The distinguishing feature of the above model is that there is no direct communication between any of the parties in Π. Instead, all communication is "routed" through $\mathcal{A}$, and she may decide for any incoming message whether it will be forwarded and, if so, whether it is modified beforehand.

4.3.2 Communication Channels

Wireless communication utilizes a shared and open medium composed of one or several physical channels, i. e., ultimately, any communication channel is realized through *at least* one physical channel.

Prominent examples for systems allocating a single *physical channel* for communication are WPANs, e. g., those that are based on IEEE 802.15.4 [IEE11] (cf. Section 4.2.1), and WLANs, e. g., WLAN with Direct-Sequence Spread Spectrum (DSSS) [IEE97]. On the other hand, mobile telephony networks, such as those based on GSM or 5G, and the satellite telephony services based on Inmarsat, have separate channels for sending and receiving data [Elb08].

4.3.2.1 Discussion

In the MitMM, each party has access to exactly *one* communication-out tape (cf. Definition 4.3.1 Property 1) and each of these is *directly linked* to a communication-in tape of the adversary (cf. Definition 4.3.1 Properties 4 and 5), thus forming the *only* available communication channel to any party.

A wireless communication channel is established using at least one physical channel. Therefore, if multiple physical channels are available, *several* communication channels can be established between two parties. In fact, techniques like FHSS, utilized in systems like Bluetooth [Blu16] and WLAN with FHSS [IEE97], cannot be represented by a single communication channel alone.

Additionally, these physical channels are *accessible* to anyone. In principle, any party listening on the same channel can receive messages transmitted over these channels. As we elaborate further in Section 4.3.3.1, this restricts the possibilities of an attacker.

4.3.2.2 Lessons Learned

A wireless communication model should allow parties to establish several communication channels. Note that these are open to everyone.

4.3.3 Signals

In wireless communication, messages are encoded into sinusoidal waves using various modulation schemes, e. g., Frequency-Shift Keying (FSK) or Phase-Shift Keying (PSK). Additionally, these waves are transmitted with a certain amount of power, which the receiver perceives as signal strength, often simply represented as the Received Signal Strength Indicator (RSSI).

4.3.3.1 Discussion

Whenever two parties send messages at the same time on the same physical channel, these messages can *collide*. This can lead to cancellation, amplification, or other modification of the resulting signal. In fact, jamming (cf. Sections 4.2.2 and 4.2.3) represents the only technical possibility for an attacker to *modify* messages.

In the MitMM, a user cannot detect whether messages sent to him have been blocked. This is no longer true for wireless communication where jamming *detection* is possible. A simple form of it relies on the RSSI, e. g., if packet errors occur during reception but the corresponding RSSI is high, then the transmission was probably jammed. This form of jamming detection has been extensively studied, e. g., by Xu et al. [XTZW05] or by Grover et al. [GLY14] in the form of a systematic overview of jamming and detection methods.

4.3.3.2 Lessons Learned

The modulation of messages onto physical signals influences an attacker's capabilities, potential countermeasures, and the creation of new schemes. Thus, a wireless communication model must encompass physical signals, including signal collisions and jamming.

4.3.4 Locality

When a party transmits a signal at a given point in time, it reaches another party with a *time delay* that is determined by the distance between the sender and the receiver. Moreover, the distance affects the *signal strength* of the perceived signal which eventually determines whether the recipient gets the message at all.

4.3.4.1 Discussion

In the MitMM, the notions of time and space are absent. However, the situation is different for wireless communication. For example, while in transit, signals experience path loss and fading, particularly when passing through obstacles. If the signal strength drops below the noise level, the transmitted data may become irrecoverable. This affects both the connectivity between parties *and* an attacker's ability to jam messages. For instance, the attacker might fail to receive the message she intends to jam, or her jamming signal may be too weak to reach the targeted party. In fact, the properties of time and space have been explored as potential foundations for new security measures in wireless systems.

In the time domain, a compelling example is the area of DB protocols, e. g., [AH13; AMSP17; BC94; BMV15a; DM07; GLO15; HK05; RC10; RCHC09; RDC15; RTŠ+12; SP07; TČ09]. These build on ideas presented by Desmedt et al. [DGB88]. DB protocols consider the round trip delay of (two) communicating parties, thus, providing an upper bound to their distance.

The property of space can also be utilized as a security measure based on physical properties. For example, ZLL [Zig12a] leverages signal strength measurements to determine *proximity* of two parties. Beyond the available power for sending and receiving, the RSSI is primarily influenced by physical factors such as path loss and distance between the sender and the receiver. In practice, distance can be easily measured, whereas path loss fluctuates over time, impacted by environmental conditions and actual physical factors, even if the distance remains unchanged. This means that when two parties establish a communication channel, an eavesdropper's view of that channel quickly de-correlates as distance increases [ESWM12; ZPP16]. Hershey et al. [HHY95] suggested using the physical environment to generate shared information between two parties, which is unique to them at that specific moment in time and space.

4.3.4.2 Lessons Learned

In the MitMM, the attacker is omnipresent, i. e., she controls all communication. This does not apply to wireless networks where an attacker's position in relation to the other parties is relevant. Even if an attacker comprises several parties at different locations, the communication between these are likewise subject to the restrictions discussed above. That is, a model should integrate the concept of relative positions of parties and the influence on the communication.

4.4 Wireless Communication Model

Next, we introduce a formal model for wireless communication on the PHY of the OSI model, dubbed WCM. To this end, we explain how the model addresses each of the three identified aspects, i. e., communication channels (cf. Section 4.3.2), signals (cf. Section 4.3.3), and locality (cf. Section 4.3.4).

We stress that the aim of the model is to represent wireless communication only, i. e., parties may have further means to communicate, e. g., being directly wired. In particular, in some scenarios, it might seem reasonable to add further communication tapes to the model in order to also cover non-wireless communication. However, such communication channels can be modeled "classically" and are hence out of scope of the model.

4.4.1 Communication Channels in the WCM

We assume a set of parties Π and an attacker $\mathcal{A}$, all modeled as probabilistic ITMs (see Katz [Kat02, Definition 2.1]) like in the MitMM. To model the wireless communication between different parties (including the attacker $\mathcal{A}$), we likewise adopt the concept of *communication tapes*. However, there are several key differences:

- Cells contain physical signals denoted by σ (cf. Section 4.4.2).
- To uniquely reference parallel activities and incorporate the concept of time, we assume a global discrete timer that segments the flow of time into discrete time slots $t \in \mathbb{N}$ with $t + 1$ following t and so on.
- A system-wide parameter, chs, represents the total number of *publicly* accessible physical channels, along with an ordering on these channels.

Contrary to Properties 1 and 2 stated in Definition 4.3.1:

- We assume that *any* party $\mathcal{P}$ (including any attacker $\mathcal{A}$) has *exactly* chs communication-in and chs communication-out tapes, which are denoted by $(\mathsf{CT}_{\text{in}})_1^{\mathcal{P}}, \ldots, (\mathsf{CT}_{\text{in}})_{\mathsf{chs}}^{\mathcal{P}}$ and $(\mathsf{CT}_{\text{out}})_1^{\mathcal{P}}, \ldots, (\mathsf{CT}_{\text{out}})_{\mathsf{chs}}^{\mathcal{P}}$, respectively.
- The i-th communication-in tape $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}}$ and the i-th communication-out tape $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}$ are connected to the i-th physical channel.
- The number of communication-out tapes available to the attacker solely depends on chs.

For any communication-in or -out tape CT, $\mathsf{CT}[j]$ denotes the j-th *cell* of this specific tape. Communication-out tapes are write-only and communication-in tapes are read-only tapes. That is, a party can write to a cell of any of its communication-out tapes and read from a cell of any of its communication-in tapes. However, reading from and writing to any cell but the current one are not possible. That is, any party $\mathcal{P}$ (including $\mathcal{A}$) can only read from the t-th cell where t denotes the current time slot. Likewise, any such party can only write to the t-th cell.[4]

To further elaborate, imagine that these tapes continuously "fly by" the parties and a party can only access the currently visible entry. The reason for this design decision is that the tapes model the physical channels used for communication. Reading or writing essentially correspond to eavesdropping on and demodulating signals or sending modulated signals on these frequencies. Once a signal passed a party, it cannot be eavesdropped on anymore. Once a signal is sent, this action cannot be taken back.

Additionally, each party can access only one cell at a time, meaning it can either read from *or* write to a *single* communication tape. The reasoning behind this design decision is as follows. In cryptography, parties are typically modeled as Turing machines, which model single-task algorithms. From a technical perspective, separate devices are required for simultaneously sending *and* receiving, or for accessing multiple physical channels. Therefore, in our model, a party represents the smallest processing unit capable of independent operation. Note that we do *not* exclude that, in practice, an attacker may indeed access multiple physical channels simultaneously. However, this would be formally represented as multiple interacting parties, potentially using non-wireless communication channels. As in the MitMM, this is not explicitly modeled here, but it is straightforward to incorporate if needed.

We emphasize that, in our model, each physical channel is assumed to be *loosely* synchronized with the global timer, meaning that each cell on a communication tape corresponds exactly to the duration of one signal period of the underlying carrier frequency. In other words, all communication channels advance with the same rate.

Strictly speaking, this is not completely accurate as some channels may appear to be ticking faster than others. However, the actual difference is considered negligible, e. g., the duration of a period for the lowest ($f_c = 2.412\,\text{GHz}$) and highest ($f_c = 2.484\,\text{GHz}$) WLAN channel in the 2.4 GHz band differ by approximately $1.2 \cdot 10^{-11}\,\text{s}$ which is a factor of 33.5 or 34.5 smaller compared to the duration of a period of the carrier frequency of the highest or lowest channel, respectively. Also note that electromagnetic waves' propagation speed is

[4] Note that there is no restriction on *storing* data that was read from a tape in the past. Likewise, there is no restriction on preparing data to be written in the future.

independent of their frequency. Therefore, we consider time slots of different channels to be loosely synchronized such that parallel and subsequent activities can still be uniquely referenced by the global timer.

4.4.2 Signals in the WCM

Communication parties are usually oblivious to the fact that, on the PHY, signals are used to transport messages. However, as discussed in Sections 4.3.3 and 4.3.4, physical signals may interfere or be delayed which eventually impacts the messages received by the parties. Given that the main intention of the proposed model is to express the capabilities of an attacker to impact wireless communication, we discuss the connection between physical signals on the one hand, and the messages sent and received by parties on the other hand in the following.

4.4.2.1 Physical Signals

Physical signals are commonly represented by sinusoidal waves, e. g., formalized as

$$\psi(t) = \mathbf{A}(t) \cdot \Big(\cos\big(2\pi f t + \boldsymbol{\phi}(t)\big) + i \cdot \sin\big(2\pi f t + \boldsymbol{\phi}(t)\big)\Big) \tag{4.1}$$

with the following parameters:

- the *amplitude* $\mathbf{A} : \mathbb{N} \to \mathbb{R}_{\geq 0}$ which is correlated to the signal strength, with $A = \mathbf{A}(t)$ denoting an output value for any fixed t,
- the *frequency* $f \in \mathbb{R}_{>0}$ where $1/f$ gives the signal's period, and
- the *phase shift* $\boldsymbol{\phi} : \mathbb{N} \to [0; 2\pi)$ that describes the phase shift of the signal compared to a non-shifted sine signal with $\varphi = \boldsymbol{\phi}(t)$ denoting an output value for any fixed t.

Consequently, we model a signal σ by this triple of parameters.

Definition 4.4.1: Signal Space and Signals

The *signal space* is denoted by $\Sigma = \mathbb{R}_{\geq 0} \times \mathbb{R}_{>0} \times [0; 2\pi)$, a *signal* σ is defined as $\sigma = (A, f, \varphi) \in \Sigma$. The term 0 refers to the signal with $A = 0$.

In our model, σ refers to the information about a physical signal stored in one slot of a communication tape. That is, it represents the state of a signal during a particular time slot. We consider all components of σ to be constant for the duration of one time slot, which is the reason for the amplitude A and the phase shift φ being represented as scalars instead of time-dependent functions. Practically, 0 refers to the case that no signal is present. In particular, an *empty cell* contains this term.

As signals are considered to be periodic, it is possible to phase-shift any signal by shifting it along the time axis. Formally, we introduce the Shift operation for this.

Definition 4.4.2: The Shift Operation

The Shift operation changes the phase of a given signal by a given parameter. It is defined as $\mathsf{Shift} : [0; 2\pi) \times \Sigma \to \Sigma$. Let $\varphi_0 \in [0; 2\pi)$ be a phase shift value and

$\sigma = (A, f, \varphi) \in \Sigma$ a signal, then,

$$\mathsf{Shift}\,(\varphi_0, \sigma) := \mathsf{Shift}_{\varphi_0}\,(\sigma) := \left(A, f, \varphi + \varphi_0 - \left\lfloor \frac{\varphi + \varphi_0}{2\pi} \right\rfloor\right). \tag{4.2}$$

4.4.2.2 Modulation and Demodulation

When two parties are communicating, they typically exchange messages consisting of *message symbols* $\mu \in M$ where M represents the message alphabet. Technically, this involves converting message symbols into physical signals (modulation) and vice versa (demodulation), as defined below.

Definition 4.4.3: Modulation and Demodulation

The *modulation function* Mod converts a message μ into a sequence of ρ signals to be sent OTA. It takes as input a message symbol $\mu \in M$ and the signal strength parameter $s \in \mathbb{R}_{>0}$, and encodes these into ρ signals $\sigma_1, \ldots, \sigma_\rho$. Note that this factor, ρ, is inherent to and depending on the chosen modulation scheme. That is, we have

$$\mathsf{Mod} : M \times \mathbb{R}_{>0} \to \Sigma^\rho, (\mu, s) \mapsto (\sigma_1, \ldots, \sigma_\rho)\,.$$

Correspondingly, the *demodulation function* Demod takes ρ signals $\sigma_1, \ldots, \sigma_\rho$ as input, and outputs a message symbol $\mu \in M$ (or, in case that demodulation is not possible, it outputs $\perp$):

$$\mathsf{Demod} : \Sigma^\rho \to M \cup \{\perp\}\,,\ (\sigma_1, \ldots, \sigma_\rho) \mapsto \mu\,.$$

Demod is essentially the inverse function of Mod if the signal strength is sufficiently strong. That is, it holds for all $\mu \in M$ that

$$\mathsf{Demod}\big(\mathsf{Mod}\,(\mu, s)\big) = \mu$$

if $s \geq \tau$ for some threshold τ.

Depending on the specific modulation scheme in use, not all signal parameters directly impact the demodulation process in terms of extracting the message. For example, in PSK, amplitude and frequency are only necessary for signal *identification*, while the encoded *information* is determined solely by the observed phase shift relative to the underlying carrier frequency. However, if the signal amplitude is too low, the receiver may be unable to successfully demodulate the signals into messages after all.

Quadrature Phase-Shift Keying. As an example of modulations, we show how to express QPSK (e. g., see Semenov [Sem11]) pertaining to Definition 4.4.3. QPSK, a variant of which is used in the IEEE 802.15.4 standard with the 2.4 GHz band [IEE11], modulates two-bit combinations onto the carrier signal. That is, the modulation function differentiates between the four phase values $\frac{\pi}{4}, \frac{3\pi}{4}, \frac{5\pi}{4}, \frac{7\pi}{4}$ conveying the four different two-bit values

$00, 01, 11, 10$ with respect to the underlying carrier frequency.[5] Concretely, the modulation and demodulation functions for QPSK with respect to Definition 4.4.3 are defined as follows.

Definition 4.4.4: QPSK Modulation and Demodulation

Let the QPSK message space $M_{\text{QPSK}} = \left\{\mu^{(i)}\right\}_{i=0,\dots,3}$ with $\mu^{(0)} = 00$, $\mu^{(1)} = 01$, $\mu^{(2)} = 11$, and $\mu^{(3)} = 10$. For $s \in \mathbb{R}_{>0}$ and $n \in \{0,1,2,3\}$, the message symbol $\mu^{(n)} \in M_{\text{QPSK}}$ is modulated by

$$\mathsf{Mod}_{\text{QPSK}}\left(\mu^{(n)}, s\right) = (\sigma, \sigma)$$

with $\sigma = \left(A, f, \frac{(2n+1)\pi}{4}\right)$ and $\rho = 2$.

For the demodulation function, let $\sigma_1 = (A_1, f_1, \varphi_1), \sigma_2 = (A_2, f_2, \varphi_2) \in \Sigma$ with $A_1, A_2 \geq \tau, f_1 \approx f_2$, and M_{QPSK} as before. Then,

$$\mathsf{Demod}_{\text{QPSK}}(\sigma_1, \sigma_2) = \begin{cases} \mu^{(n)}, & \text{if } \frac{2n\pi}{4} \leq \varphi_1, \varphi_2 < \frac{2(n+1)\pi}{4}, n \in \{0,1,2,3\} \\ \perp, & \text{else.} \end{cases}$$

4.4.2.3 Sending and Receiving Message Symbols

Sending a message symbol μ involves modulating it into a sequence of signals $\sigma_1, \dots, \sigma_\rho$ based on a given signal strength s, using the Mod function. These signals are then transmitted over one of the available communication channels by writing each signal σ_i on a communication-out tape. This procedure is formalized as the Send operation.

Definition 4.4.5: The Send Operation

The Send operation takes as input a triple (CT, μ, s) where CT is a communication-out tape of $\mathcal{P}$, $\mu \in M$ is a message symbol, and $s \in \mathbb{R}_{>0}$ represents a signal strength. It involves executing $\mathsf{Mod}(\mu, s)$ first to get a sequence of ρ signals $\sigma_1, \dots, \sigma_\rho$. These are written step-by-step onto the tape CT. Note that Send does not specify the intended recipient of the message symbol μ but the communication channel. Recall that during one time slot, a party $\mathcal{P}$ can only write to the t-th cell. That is, the writing procedure is formally equivalent to setting ρ consecutive cells of CT as follows:

$$\mathsf{CT}[t] := \sigma_1, \dots, \mathsf{CT}[t + \rho - 1] := \sigma_\rho.$$

Therefore, the Send operation influences the current and the next $\rho - 1$ cells of the communication-out tape CT, which, in total, takes ρ time slots to perform.

Also, we extend the notation to the case of sending messages: for a message $m = (\mu_1, \dots, \mu_\ell)$, we define by $\mathsf{Send}(\mathsf{CT}, m, s)$ the sequence of operations $\mathsf{Send}(\mathsf{CT}, \mu_i, s)$, $i = 1, \dots, \ell$.

[5] The order is chosen according to Gray encoding, i. e., neighboring values only differ in a single bit.

Likewise, parties can *receive* message symbols from their communication-in tapes and only from these.

Definition 4.4.6: The Receive Operation

The Receive operation takes as input a communication-in tape CT. It reads out the current ρ entries of CT, i. e., $\mathsf{CT}[t], \ldots, \mathsf{CT}[t+\rho-1]$. Given this, $\mathsf{Demod}(\mathsf{CT}[t], \ldots, \mathsf{CT}[t+\rho-1])$ is executed to get a message symbol $\mu \in M \cup \{\perp\}$. Then, μ represents the output of Receive.

Since the Demod function requires ρ signals to extract a message, and since it is not possible to receive future signals ahead of time, this definition implies that Receive must wait for ρ time slots before being able to returning a message μ. In other words, Receive must wait until the demodulation function has received enough signals to produce a valid message symbol $\mu \neq \perp$.

4.4.2.4 Signal Collision

As the parties may share the same physical channel, different signals could *collide*. Formally, this allows for parallel writing on communication tapes, unlike in the MitMM. Additionally, because accessing communication cells "from the past" is not possible, this leads to the conclusion that the WCM is not message-driven.

For two signals $\sigma, \sigma' \in \Sigma$, we denote by

$$\sigma + \sigma'$$

the resulting signal. That is, if two different signals σ and σ' are simultaneously received on the same tape CT, then $\sigma + \sigma'$ appears on the tape instead (see Section 4.4.3 for more details). The concrete working of the collision depends on the underlying physical channel and the selected modulation/demodulation procedures.

Nonetheless, there are some properties which hold in all cases, e. g.,

$$\sigma + 0 = 0 + \sigma = \sigma$$

for all signals σ. Furthermore, collision is commutative and associative:

$$\sigma + \sigma' = \sigma' + \sigma \quad \text{and} \quad \sigma + \left(\sigma' + \sigma''\right) = \left(\sigma + \sigma'\right) + \sigma'' .$$

Collision is also commutative with respect to Shift (cf. Definition 4.4.2), i. e.,

$$\mathsf{Shift}_\varphi\left(\sigma + \sigma'\right) = \mathsf{Shift}_\varphi\left(\sigma\right) + \mathsf{Shift}_\varphi\left(\sigma'\right) .$$

In certain special cases, the signals amplify or cancel each other out, depending on their relative phase shift. More precisely, it holds

$$(A, f, \varphi) + \left(A', f, \varphi\right) = \left(A + A', f, \varphi\right), \tag{4.3}$$

$$(A, f, \varphi) + \mathsf{Shift}_\pi\left(A', f, \varphi\right) = \begin{cases} \left(A - A', f, \varphi\right), & \text{if } A \geq A' \\ \mathsf{Shift}_\pi\left(A' - A, f, \varphi\right), & \text{else.} \end{cases} \tag{4.4}$$

In particular, it holds that $(A, f, \varphi) + \mathsf{Shift}_\pi(A, f, \varphi) = 0$. Consequently, we define for a given signal $\sigma = (A, f, \varphi)$ its inverse as

$$-\sigma := \mathsf{Shift}_\pi(A, f, \varphi) \ . \tag{4.5}$$

Channels are characterized by signals that share the same frequency f within a specific bandwidth, typically chosen to minimize interference between neighboring channels. For instance, in IEEE 802.15.4, channels in the 2.4 GHz band are non-overlapping as they are 5 MHz apart with a bandwidth of 2 MHz [IEE11]. Therefore, we focus on signal collisions occurring on the same channel. Moreover, for a given frequency f, we use the term Σ_f to refer to the set of signals that approximately have f as their frequency. In other words, signals belong to the same channel if and only if they are elements of the same space Σ_f for a selected frequency f. For any frequency f, it holds that

$$\sigma, \sigma' \in \Sigma_f \Rightarrow \sigma + \sigma' \in \Sigma_f \ .$$

Formally, this means that Σ_f is closed under $+$. In combination with the properties mentioned above, it follows that $(\Sigma_f, +)$ forms a commutative group. This perspective is, for example, helpful when describing the effects of jamming, e. g., see Section 4.6.3.

4.4.3 Locality in the WCM

In wireless communication, signals are *physical* objects. This has a number of consequences that are *not* covered by the MitMM, e. g., multiple signals can collide, resulting in different signals (cf. Section 4.4.2.4). Further relevant aspects are:

1. Once a party sends a signal, it takes *time* until it reaches another party.
2. During transmission, the signal strength may *decrease*.

To represent these aspects, we adopt and extend the notion of being linked (cf. Definition 4.3.1). Due to the fact that all parties rely on *public* physical channels, any two parties are linked in the sense that, potentially, messages can be exchanged. Here, we also have to take into account their relative locations. Therefore, we propose the following formal definition of linkage.

Definition 4.4.7: Linkage between two parties

Consider two parties $\mathcal{P}, \mathcal{P}' \in \Pi \cup \{\mathcal{A}\}$ with $\mathcal{P} \neq \mathcal{P}'$. The *linkage from $\mathcal{P}$ to $\mathcal{P}'$*, denoted by $\mathsf{Link}(\mathcal{P}, \mathcal{P}')$, is defined by a tuple

$$\mathsf{Link}\left(\mathcal{P}, \mathcal{P}'\right) = (\delta; \lambda_1, \ldots, \lambda_{\mathsf{chs}}) \tag{4.6}$$

where

- $\delta \in \mathbb{N}_{\geq 0}$ is a non-negative integer, representing the time delay, and
- $\lambda_i : \Sigma \to \Sigma$ is a probabilistic procedure, dubbed linkage procedure, expressing changes of a signal sent by $\mathcal{P}$ (using $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}$) and received by $\mathcal{P}'$ on $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}$.

The linkage between two parties defines when and what kind of signal arrives at $\mathcal{P}'$ when $\mathcal{P}$ sends a signal. For instance, the linkage procedures λ_i could represent a reduction in amplitude A, e. g., due to expected path loss based on the distance between the sender and the recipient. This can be expressed as $\lambda_i(A, f, \varphi) = (A', f, \varphi)$, where $A' < A$.

Another example involves a channel-induced phase shift on transmitted signals. Assume that $\mathcal{P}$ and $\mathcal{P}'$ are separated by a *physical* distance d, which takes the form $d = \delta \cdot \frac{c}{f_i} + r = \delta \cdot \frac{c}{f_i} + \varphi_i \cdot \frac{c}{2\pi f_i}$, where $\delta \geq 0$ is an integer, c is the speed of light constant, f_i is the carrier frequency of the i-th communication channel, and r is the remainder, with $0 \leq r < \frac{c}{f_i}$. In this case, d is *not* a multiple of the carrier frequency's wavelength. The linkage can then be described as $\lambda_i(A, f, \varphi) = (A', f', \varphi + 2\pi - \varphi_i) = (A', f', \varphi')$.

Recall that we have one linkage procedure λ_i per channel with a total number of chs available channels and a fixed order (cf. Section 4.4.1), such that the i-th communication-in and communication-out tapes refer to the i-th physical channel and that for any index $i \in \{1, \ldots, \mathsf{chs}\}$ the i-th communication-out tape of some party is linked to every i-th communication-in tape of any other party.

Now, assume that some party $\mathcal{P}$ writes a signal σ on one of its communication-out tapes $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}$, i. e., $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}[t] := \sigma$. Writing on CT automatically affects *all* communication-in tapes of parties that are linked to $\mathcal{P}$. That is, let $\mathcal{P}' \neq \mathcal{P}$ be some other party and let $\mathsf{Link}(\mathcal{P}, \mathcal{P}') = (\delta; \lambda_1, \ldots, \lambda_{\mathsf{chs}})$ be the linkage between $\mathcal{P}$ and $\mathcal{P}'$. The procedure of $\mathcal{P}$ writing σ into the cell $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}[t]$ affects the content of the corresponding communication-in tape $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}$ for party $\mathcal{P}' \neq \mathcal{P}$ as follows:

$$(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}[t + \delta] := \lambda_i(\sigma) \ + \ (\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}[t + \delta]\,. \tag{4.7}$$

That is, the signal already residing in $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}$, expressed by $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}[t + \delta]$, is updated to $\lambda_i(\sigma) \ + \ (\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}[t + \delta]$. This definition allows covering the aspects mentioned above:

1. When party $\mathcal{P}$ sends a signal σ at a specific time t, it reaches $\mathcal{P}'$ after a time delay δ. In other words, $\mathcal{P}$ writes on cell $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}[t]$, i. e., with index t, but this affects $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}[t + \delta]$, which occurs δ time slots later.

2. Moreover, the signal σ physically travels through the space between $\mathcal{P}$ and $\mathcal{P}'$ and arrives at $\mathcal{P}'$ as the signal $\lambda_i(\sigma)$.

3. If there is already a signal present in the targeted cell $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}[t + \delta]$, then a collision occurs between $\lambda_i(\sigma)$ and the already existing signal (cf. Equation (4.7)).

A schematic overview is depicted in Figure 4.2.

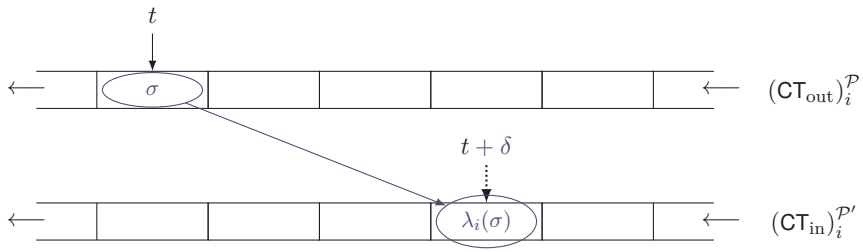

Figure 4.2: Linkage from $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}$ to $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}$ with delay δ.

Finally, we introduce two notions with respect to the time delay.

Definition 4.4.8: Time Delay Symmetry and Triangle Inequality

Consider a set of parties Π and their linkages $\mathsf{Link}(\mathcal{P}, \mathcal{P}')$ for any $\mathcal{P} \neq \mathcal{P}' \in \Pi$. We denote by $\delta_{\mathcal{P},\mathcal{P}'}$ the time delay in $\mathsf{Link}(\mathcal{P}, \mathcal{P}')$.

We say that the time delay parameter is symmetric with respect to these linkages if it holds for any $\mathcal{P} \neq \mathcal{P}' \in \Pi$ that

$$\delta_{\mathcal{P},\mathcal{P}'} = \delta_{\mathcal{P}',\mathcal{P}} \,. \tag{4.8}$$

Moreover, we say that the time delay parameter fulfills the triangle inequality if it holds for any pairwise distinct $\mathcal{P}, \mathcal{P}', \mathcal{P}'' \in \Pi$ that

$$\delta_{\mathcal{P},\mathcal{P}''} \leq \delta_{\mathcal{P},\mathcal{P}'} + \delta_{\mathcal{P}',\mathcal{P}''} \,. \tag{4.9}$$

If not stated otherwise, in the following, we assume that the delays of the different linkages are symmetric and fulfill the triangle inequality.[6]

4.5 Wireless Attacker Model

Recall that in the MitMM, an attacker is mainly characterized by her ability to stop, forward, and modify messages at her discretion. In the following, we discuss if and to what extent this is still given in the WCM. To this end, we present the WAM. First, in Section 4.5.1, we explain an attacker's alleged capabilities. Based on this, we investigate the possibilities for jamming (tampering with) signals in Section 4.5.2.

4.5.1 Capabilities

Analogous to the MitMM, for the WAM, we consider an attacker completely characterized by her abilities to access her communication tapes. This means she can send signals by writing to her communication-out tapes and receive signals by reading from her communication-in tapes. Furthermore, the attacker is non-invasive in the sense that she does *not* tamper with the parties or the environment, but only sends or receives signals. However, the attacker is not limited in her choice of signals, i. e., the signals do not necessarily have to be related to the modulation of a message symbol. Naturally, she may still choose to use the Mod and Send as well as the Demod and Receive functions for communication.

Additionally, we assume that the attacker has complete knowledge of the communication channels between herself and the other parties. Formally, this means the attacker $\mathcal{A}$ knows $\mathsf{Link}(\mathcal{A}, \mathcal{P})$ and $\mathsf{Link}(\mathcal{P}, \mathcal{A})$ for all $\mathcal{P} \in \Pi$. However, the linkages between any two parties $\mathcal{P}, \mathcal{P}' \in \Pi$ that do *not* involve the attacker are only partially known to her. As explained in Section 4.3.4.1, $\mathcal{A}$ knows only the delay parameter δ of $\mathsf{Link}(\mathcal{P}, \mathcal{P}')$, while the λ_i remain secret.

4.5.2 Modification of Signals Through Jamming

As discussed in Section 4.2.2, jamming refers to any intentional interference with signals and, as a possible result, modifications thereof. Consequently, we use the term *jamming* to

[6] In fact, we cannot think of any contradictory practical scenario.

express any change of the signal incurred by an attacker. Before we discuss the capabilities of an attacker in more detail, we have to point out one important difference between tampering in the MitMM and jamming in the WCM. In the latter, messages are composed of message symbols which in turn are modulated into a sequence of ρ signals. Thus, whether the alteration of signals results in any changes in the message received by $\mathcal{P}'$ depends at least on the demodulation procedure Demod and how the semantic message is derived from the message symbols. For example, a single bit flip might render a message unreadable if this change violates some checksum. However, if some error correction codes are in use, such a bit flip might have no effect at all. Thus, we focus only on the question if and how an attacker may change *signals* and leave the discussion on the impact on messages to the individual scenarios.

For the discussion on jamming signals, consider two parties $\mathcal{P}, \mathcal{P}'$. Now, at time t, $\mathcal{P}$ sends a signal σ to $\mathcal{P}'$, i. e., $\mathcal{P}'$ writes σ on one of its communication-out tapes $(\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}$. Let $\mathsf{Link}(\mathcal{P}, \mathcal{P}') = \left(\delta_{\mathcal{P},\mathcal{P}'}; \lambda_1, \ldots, \lambda_{\mathsf{chs}}\right)$ denote the linkage between $\mathcal{P}$ and $\mathcal{P}'$. If no other signals are present on the same channel, then,

$$(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}\left[t + \delta_{\mathcal{P},\mathcal{P}'}\right] = \lambda_i\left((\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}[t]\right)$$

represents the signal received by $\mathcal{P}'$.

For $\mathcal{A}$, the only option to influence σ is to send her own signal such that it collides with the signal received by $\mathcal{P}'$. Let $\mathsf{Link}(\mathcal{A}, \mathcal{P}') = \left(\delta_{\mathcal{A},\mathcal{P}'}; \lambda_1^*, \ldots, \lambda_{\mathsf{chs}}^*\right)$ be the linkage between $\mathcal{A}$ and $\mathcal{P}'$. Then, in essence, jamming translates to provoking a collision $(\mathsf{CT}_{\text{in}})_i^{\mathcal{P}'}\left[t + \delta_{\mathcal{P},\mathcal{P}'}\right] = \sigma + \sigma^*$ with $\sigma = \lambda_i\left((\mathsf{CT}_{\text{out}})_i^{\mathcal{P}}[t]\right)$ and $\sigma^* = \lambda_i^*\left((\mathsf{CT}_{\text{out}})_i^{\mathcal{A}}\left[t + \delta_{\mathcal{P},\mathcal{P}'} - \delta_{\mathcal{A},\mathcal{P}'}\right]\right)$. Hence, $\mathcal{A}$ has to write a signal on her communication-out tape no later than $t + \delta_{\mathcal{P},\mathcal{P}'} - \delta_{\mathcal{A},\mathcal{P}'}$. Obviously, if $\delta_{\mathcal{P},\mathcal{P}'} < \delta_{\mathcal{A},\mathcal{P}'}$, i. e., the distance between $\mathcal{P}$ and $\mathcal{P}'$ is smaller than the distance between $\mathcal{P}'$ and $\mathcal{A}$, an attacker would have to send her signal even *before* $\mathcal{P}$ sent σ. In certain cases, this may indeed be possible — such as when σ is part of a longer message and $\mathcal{A}$ can predict that σ (or some other signal) will be sent. However, even in these cases, $\mathcal{A}$ cannot *react* to signals from $\mathcal{P}$. The same applies when $\delta_{\mathcal{P},\mathcal{P}'} < \delta_{\mathcal{P},\mathcal{A}}$. In this scenario, σ from $\mathcal{P}$, i. e., the signal that $\mathcal{A}$ aims to jam, would reach $\mathcal{P}$' *before* it reaches $\mathcal{A}$.

In addition to sending the jamming signal in time, another challenge is selecting and transmitting a signal that effectively jams. Recall that, in some cases, two signals can amplify or cancel each other out (cf. Section 4.4.2.4). However, an amplified signal does not necessarily affect a modulated message. That is, the effectiveness of a jamming signal heavily depends on the modulation chosen.

Summing up, to manipulate signals received by other parties, a wireless attacker has to decide *when* and *what* signals she sends. Depending on her jamming strategy, this can be less or more challenging. For example, an attacker could simply resort to less advanced jamming strategies, e. g., constant or random jamming (cf. Section 4.2.2). Formally, a constant jammer would keep the channel busy by constantly writing signals with a sufficiently high signal strength on her communication-out tapes. An attacker representing a random jammer, who ultimately only jams *some* messages of all transmitted, would only randomly write signals on her communication-out tapes. While this somewhat relaxes the timing and modulation restrictions discussed above, the effect on the jammed signal is hard to predict. Note that both of these less advanced strategies incur higher activity and risk of being detected, located, and eventually neutralized. In contrast, more sophisticated jammers,

e. g., random or selective jammer (cf. Sections 4.2.2.1 and 4.2.2.2), only target certain communication and messages, i. e., these would constantly read from their communication-in tape and, once (targeted) ongoing communication is noticed, they would then write signals on their communication-out tape. While these types of jammers are stealthier, and thus, harder to detect, they face all those problems mentioned above.

Interestingly, Pöpper et al. [PTDC11] reported that intentional and *deterministic* signal manipulation of QPSK-modulated (cf. Section 4.4.2 and Definition 4.4.4) signals is not easily accomplished in the real world. They show that an attacker must be extremely precise to successfully flip a targeted bit instead of a neighboring or *random* bit flip. Furthermore, the considered *strong* attacker knows transmitted signals in advance and, hence, can prepare her own based on the targeted signals. Yet even then, the attacker needs her signal's carrier phase to be within 25 % of a signal duration relative to the targeted signal at the *receiver* [PTDC11].[7] On the other hand, changing at least any single bit during an ongoing transmission is reported to be much easier. Note that, however, a receiving party would discard such a manipulated message due to the invalid CRC which would indicate a transmission error. This implies that the attacker needs to successfully modify the targeted part of a message *and* the CRC for a receiver to accept the manipulated message.

4.6 Applications

To demonstrate the applicability of the proposed model, in this section, we revisit FHSS (Section 4.6.1), DB protocols (Section 4.6.2), and friendly jamming (Section 4.6.3), all of which cannot be represented in the MitMM, and we explain how the WCM and WAM can be used for analyzing these. Among others, we provide a formal proof for the security of DB protocols and formulate a security notion for friendly jamming. Note that each one of the three different components identified and listed in Section 4.3 are required at least once, showing that at least these three are required for a comprehensive wireless communication model.

4.6.1 Frequency-Hopping Spread Spectrum

The core idea of the FHSS technique is to avoid using a single physical channel during communication by hopping across available channels in a (pseudo-)random manner. This enhances robustness against interference on a specific frequency (or channel) and makes it more difficult for an attacker to target the correct channel if the hop sequence is secret.

In our model, incorporating FHSS is naturally possible by switching between different communication tapes when sending messages, as tapes essentially represent different channels with distinct carrier frequencies. We illustrate this with a variant of FHSS where the messages sent by $\mathcal{P}$ are indexed, i. e., $m_1, m_2, \ldots$, and for each message, a separate communication channel is selected. To achieve this, sender $\mathcal{P}$ and receiver $\mathcal{P}'$ agree on a secret key k to initialize a pseudorandom function $f_k : \mathbb{N} \to \{1, \ldots, \mathsf{chs}\}$. For each message m_i, f_k determines which communication tape should be used. More precisely, sending the i-th

[7] More precisely, for the considered modulation scheme, QPSK, the attacker needs to be within 13.5 % of a signal duration for changing targeted signals *arbitrarily*, i. e., performing two-bit inversions instead of only one-bit inversions.

message means that $\mathcal{P}$ executes

$$\mathsf{Send}\left((\mathsf{CT}_{\mathrm{out}})^{\mathcal{P}}_{f_k(i)}, m_i, s\right) .$$

Likewise, $\mathcal{P}'$ uses f_k to determine the channel to listen to next.

4.6.2 Distance Bounding Protocols

DB protocols were introduced by Brands and Chaum [BC94] with the goal of allowing a prover P to authenticate themselves to a verifier V and simultaneously demonstrate their proximity to V. To achieve this, the verifier measures the round-trip time and, using the propagation speed of electromagnetic waves, calculates an upper bound for the distance, ensuring that the verifier is confident the prover is no farther from the verifier than this calculated bound.

In our model, the timing aspect is represented by the delay parameter δ in the linkage between two parties (cf. Definition 4.4.7). More specifically, consider the two parties $\mathcal{P}, \mathcal{P}'$, where $\mathcal{P}$ wants to determine the distance to $\mathcal{P}'$. Assume the linkages $\mathsf{Link}(\mathcal{P}, \mathcal{P}')$ and $\mathsf{Link}(\mathcal{P}', \mathcal{P})$ (cf. Definition 4.4.7 and Equation (4.6)) with a symmetric delay δ (cf. Definition 4.4.8 and Equation (4.8)). In this case, the round-trip time between $\mathcal{P}$ and $\mathcal{P}'$ is at least $2 \cdot \delta$. Therefore, measuring the round-trip time provides a lower bound on δ, which in turn gives an upper bound on the *geographical* distance between them.

In the following, we formally show the security of a DB protocol within our model. Our goal is to prove a statement analogous to the "fundamental lemma" [BMV15b, Lemma 3.1] by Boureanu et al.: when a verifier V sends a challenge, any response that arrives at the verifier before the round trip time between verifier V and prover P is completed, is either independent of the prover P or of the challenge. While the original lemma [BMV15b, Lemma 3.1] has been formulated directly for challenge-response protocols and the event of accepting responses, we opt for a more general statement. In a nutshell, it says that in the scenario mentioned above, if $\mathcal{P}$ sends some signal σ to $\mathcal{P}'$ at some point in time t, then the concrete value of σ cannot impact $\mathcal{P}'$ before time $t + \delta$. We make this statement, in particular the meaning of "having impact", more precise in the following.

Assume a set of parties $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_\ell\}$ (that may also include some attacker). To make the description more readable, we restrict to the case of $\mathsf{chs} = 1$, we abbreviate $(\mathsf{CT}_{\mathrm{in}})^{\mathcal{P}_i}$ by $(\mathsf{CT}_{\mathrm{in}})^{i}$, and likewise for $(\mathsf{CT}_{\mathrm{out}})^{\mathcal{P}_i}$. All the parties are linked, and we refer to the delay from $\mathcal{P}_i$ to $\mathcal{P}_j$ with $\delta_{i,j}$. We refer to all these parties and the content of their tapes (including their random tapes) as a system $\mathcal{S}$.

Now assume that one party, say $\mathcal{P}_1$, sends some signal σ, at point in time t_0, that is,

$$(\mathsf{CT}_{\mathrm{out}})^{1}[t_0] := \sigma . \tag{4.10}$$

Intuitively, we want to express that the fact of sending σ cannot have any impact on any other party $\mathcal{P}_i$ before $t_0 + \delta_{1,i}$. To make the concept of "impact" more formal, we consider a copy of the system $\mathcal{S}$ that we refer to as the system copy $\widetilde{\mathcal{S}}$. This $\widetilde{\mathcal{S}}$ allows us to express possible impact of signals by inspecting potential differences between $\mathcal{S}$ and $\widetilde{\mathcal{S}}$.

In many aspects, $\widetilde{\mathcal{S}}$ is a 1:1 copy of $\mathcal{S}$. Both systems, $\mathcal{S}$ and $\widetilde{\mathcal{S}}$, comprise the same set of parties Π, including exactly the same random tapes. As a consequence, a party $\mathcal{P}$ has the same deterministic behavior in $\mathcal{S}$ and $\widetilde{\mathcal{S}}$ if the values on the communication-in tapes are the same. To distinguish between the communication tapes of $\mathcal{S}$ and $\widetilde{\mathcal{S}}$, we employ a different

notation for the latter and use $\left(\widetilde{\mathsf{CT}_{\text{in}}}\right)^i$ and $\left(\widetilde{\mathsf{CT}_{\text{out}}}\right)^i$ to refer to the communication-in and communication-out tapes of $\mathcal{P}_i$ in $\widetilde{\mathcal{S}}$, respectively.

To study the potential impact of sending σ (cf. Equation (4.10)), we inject one change into $\widetilde{\mathcal{S}}$, namely

$$\left(\widetilde{\mathsf{CT}_{\text{out}}}\right)^1[t_0] := \widetilde{\sigma} \tag{4.11}$$

with $\widetilde{\sigma}$ being an arbitrary signal different from σ. That is, both systems are exactly identical up to t_0 (except for the change expressed in Equation (4.11)), but they may evolve differently from then on. Note that $\widetilde{\mathcal{S}}$ is "inconsistent" in the sense that $\mathcal{P}_1$ would not send $\widetilde{\sigma}$ (as it is deterministic, and we already know that it would send σ). One could imagine that the signal "magically" changed after being submitted. This inconsistency is not a problem, however, as we are only interested in what potential differences Equation (4.11) may induce in $\widetilde{\mathcal{S}}$ from t_0 on. Those parts that are the same in both systems are definitely independent of the concrete signal sent by $\mathcal{P}_1$ at time t_0. More precisely, the following theorem holds.

Theorem 4.6.1

Consider the two systems $\mathcal{S}$ and $\widetilde{\mathcal{S}}$ as explained above. Let $i \in \{2, \ldots, \ell\}$ be arbitrary. It holds that

$$(\mathsf{CT}_{\text{in}})^i[t] = \left(\widetilde{\mathsf{CT}_{\text{in}}}\right)^i[t] \quad \forall t < t_0 + \delta_{1,i}\,.$$

Proof. For $t \leq t_0$, the claim holds by construction of $\widetilde{\mathcal{S}}$. Assume now that the claim does not hold, i. e., $\exists i \in \{2, \ldots, \ell\}$ and $t_0 < t < t_0 + \delta_{1,i}$ such that $(\mathsf{CT}_{\text{in}})^i[t] \neq \left(\widetilde{\mathsf{CT}_{\text{in}}}\right)^i[t]$.

Let t^* be the smallest value for which the claim does not hold, i. e., we have an index i^* and a point in time $t^* < t_0 + \delta_{1,i^*}$ such that

$$(\mathsf{CT}_{\text{in}})^{i^*}[t^*] \neq \left(\widetilde{\mathsf{CT}_{\text{in}}}\right)^{i^*}[t^*] \tag{4.12}$$

but for all $j \neq i^*$ and $t \leq t^*$, it holds that

$$(\mathsf{CT}_{\text{in}})^j[t] = \left(\widetilde{\mathsf{CT}_{\text{in}}}\right)^j[t]\,.$$

Observe that content on communication-in tapes results from content previously written on communication-out tapes. Thus, Equation (4.12) implies the existence of some other party with index $j^* \neq i^*$ such that

$$(\mathsf{CT}_{\text{out}})^{j^*}[t^* - \delta_{j^*,i^*}] \neq \left(\widetilde{\mathsf{CT}_{\text{out}}}\right)^{j^*}[t^* - \delta_{j^*,i^*}]\,.$$

As the behavior of $\mathcal{P}_{j^*}$ is completely determined by the content of the random tape (which is equal in both systems $\mathcal{S}$ and $\widetilde{\mathcal{S}}$) and the content of its input tape, it follows that the content of the input tape of $\mathcal{P}_{j^*}$ must have been different at some point in time before $t^* - \delta_{j^*,i^*}$. That is, there exists a point in time $t^{**} < t^* - \delta_{j^*,i^*}$ such that

$$(\mathsf{CT}_{\text{in}})^{j^*}[t^{**}] \neq \left(\widetilde{\mathsf{CT}_{\text{in}}}\right)^{j^*}[t^{**}]\,.$$

Because of $t^{**} < t^* - \delta_{j^*,i^*} \leq t^*$ and the minimality of t^*, it must follow that $t^{**} \geq t_0 + \delta_{1,j^*}$. Thus, we get

$$t_0 + \delta_{1,j^*} \leq t^{**} < t^* - \delta_{j^*,i^*} < t_0 + \delta_{1,i^*} - \delta_{j^*,i^*} .$$

We focus on the left- and right-most terms and slightly transform the inequality, i. e.,

$$\begin{aligned} & t_0 + \delta_{1,j^*} < t_0 + \delta_{1,i^*} - \delta_{j^*,i^*} \\ \Leftrightarrow \quad & t_0 + \delta_{1,j^*} + \delta_{j^*,i^*} < t_0 + \delta_{1,i^*} . \end{aligned}$$

Additionally considering the triangle inequality (cf. Definition 4.4.8 and Equation (4.9)), this implies

$$t_0 + \delta_{1,i^*} \leq t_0 + \delta_{1,j^*} + \delta_{j^*,i^*} < t_0 + \delta_{1,i^*} . \tag{4.13}$$

Equation (4.13) is obviously a contradiction, thus proving the claim. □

Theorem 4.6.1 implies the security of a DB protocol as follows. Assume that a verifier V and a prover P are running a *secure* challenge-response protocol. That is, V sends a challenge c at some point in time t_0 and receives a response r at point in time $t_1 < t_0 + 2 \cdot \delta$ where δ is the delay between V and P.

If the response has not been generated by the prover P, it follows from the security of the underlying challenge-response protocol that r is accepted only with negligible probability. Hence, we can restrict to the situation that r is coming directly from P as shown in Figure 4.3.[8]

As the response r reaches V at time t_1, it must have been sent by P at time $t^* := t_1 - \delta < t_0 + \delta$. According to Theorem 4.6.1, it holds that $(\mathsf{CT}_{\text{in}})^P[t]$ is independent of c for all $t \leq t^*$. As the value r is uniquely determined by the values $(\mathsf{CT}_{\text{in}})^P[t]$ with $t \leq t^*$ and the random tape (which is independent of c as well), it follows that the response r is independent of the challenge c and security follows from the challenge-response protocol as well.

4.6.3 Friendly Jamming for Confidentiality

Complementing the typically disruptive form of jamming (cf. Sections 4.2.2 and 4.5.2), *friendly jamming* uses deliberate jamming as a defensive measure to ensure authorization or confidentiality (cf. Section 4.2.3). An example of this is the protection of communication to and from an IMD that only supports plain (unencrypted) communication, as described by Gollakota et al. [GHR+11]. They propose a wearable external device, called the *shield*, which acts as a gateway and is located close to the IMD. Using a two-radio design, the shield friendly-jams to enforce authorization and confidentiality for the communication with the IMD. Due to its design, the shield is capable of recovering the original data signals, while other parties receive only jammed signals. However, the authors have not provided any formal representation of the security goals or analysis.

[8] Due to the triangle inequality (Definition 4.4.8), the following statements also hold if r has been forwarded by other parties.

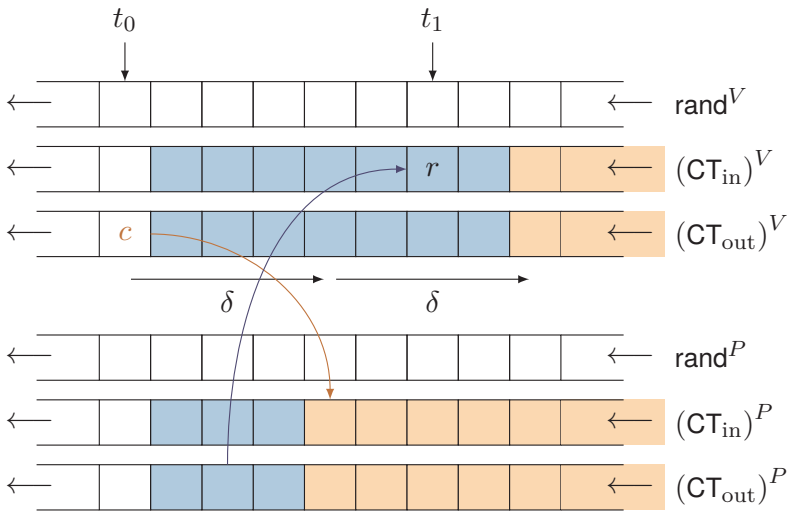

Figure 4.3: A challenge-response protocol run between verifier V sending challenge c at time t_0 and prover P sending response r at time $t_1 - \delta < t_0 + \delta$. According to Theorem 4.6.1, the content of the blue slots are independent of the concrete value of c. This holds in particular for r as well in this case.

4.6.3.1 Formalization of Friendly Jamming

The typical scenario for friendly jamming involves two collaborating parties: the jamming party J and the receiving party R, who have an additional private communication channel (cf. Section 4.4.1). The primary goal of friendly jamming is to protect a signal σ. For example, in the use case presented by Gollakota et al. [GHR+11], σ could either be a signal from the IMD whose content needs to be concealed, or a non-genuine signal sent *to* the IMD that should be rendered illegible. To achieve this, J generates a jamming signal γ and transmits it such that it collides with σ. As a result, any potential attacker will only see $\sigma+\gamma$. To reverse the jamming, J shares all the necessary information with receiver R so that R can generate an appropriate antidote signal α. When α collides with $\sigma + \gamma$, the original signal σ is recovered.

Note that this approach requires J to send γ in time, which can be challenging in practice (cf. Section 4.5.2). As this depends on various aspects such as the relative positions of parties, the reaction time of J, etc., we consider this a separate question and focus on the generation of the jamming signal and its antidote. Consequently, one can formalize[9] such a scheme as follows.

Definition 4.6.1: Friendly Jamming Scheme

A *friendly jamming scheme* $\Gamma = (\mathsf{Gen}, \mathsf{Jam}, \mathsf{Antijam}, \Sigma_f)$ consists of three algorithms and a signal space restricted to some frequency f. $\mathsf{Gen}(\chi)$ takes as input a security parameter χ and outputs a seed κ. $\mathsf{Jam}(\kappa)$ takes as input some seed κ and creates a jamming signal $\gamma \in \Sigma_f$. $\mathsf{Antijam}(\kappa)$ takes as input some seed κ and creates an antidote

[9] For the sake of simplicity, here, we consider one signal only. The formalization can be extended easily to a sequence of signals.

signal $\alpha \in \Sigma_f$.

The scheme is correct if for all seeds $\kappa \leftarrow \mathsf{Gen}(\chi)$ and for all security parameters χ, it holds that

$$\mathsf{Jam}(\kappa) + \mathsf{Antijam}(\kappa) = 0\,.$$

Correctness essentially means that $\mathsf{Antijam}(\kappa) = -\mathsf{Jam}(\kappa)$ for all seeds (cf. Section 4.4.2.4 and Equation (4.5)). In addition to correctness, a friendly jamming scheme should also be sound. Depending on the application, various security goals may be considered, such as authenticity or confidentiality. Here, we focus on confidentiality.

Intuitively, confidentiality ensures that the original message cannot be recovered from the jammed signals. To formalize this security goal, we adopt the well-established security definition of IND-CPA (Indistinguishability under Chosen Plaintext Attack) and introduce IND-CSA-n. The parameter n highlights the main difference from IND-CPA, i. e., the attacker may use n antennas, formally represented by n parties under the attacker's control.

Definition 4.6.2: IND-CSA-n Game

The IND-CSA-n game with respect to some friendly jamming scheme Γ according to Definition 4.6.1, i. e., $\Gamma = (\mathsf{Gen}, \mathsf{Jam}, \mathsf{Antijam}, \Sigma_f)$, considers an attacker $\mathcal{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ and an oracle $\mathcal{O} = \{\mathcal{O}_1, \mathcal{O}_2\}$. Each oracle has access to one communication-out tape and each attacker to one communication-in tape. The tapes controlled by $\mathcal{O}$ are linked to all tapes controlled by $\mathcal{A}$. For $i \in \{1, 2\}$ and $j \in \{1, \ldots, n\}$, we denote by $\lambda_{i,j}$ the linkage procedure of the linkage between the communication-out tape of $\mathcal{O}_i$ and the communication-in tape of $\mathcal{A}_j$.

$\mathcal{A}$ chooses two signals $\sigma_0 \neq \sigma_1 \in \Sigma_f$ and sends these to $\mathcal{O}$. Then, $\mathcal{O}$ generates a (secret) seed $\kappa \leftarrow \mathsf{Gen}(\chi)$ and a jamming signal $\gamma \leftarrow \mathsf{Jam}(\kappa)$. It samples uniformly $b \xleftarrow{\$} \{0, 1\}$ and, at time t, instructs $\mathcal{O}_1$ to write σ_b on $(\mathsf{CT}_{\text{out}})^{\mathcal{O}_1}$ and $\mathcal{O}_2$ to write γ on $(\mathsf{CT}_{\text{out}})^{\mathcal{O}_2}$ (simultaneously). Consequently, $\mathcal{A}_j$ receives the collided signal

$$\sigma_j^{\text{col}} = \lambda_{1,j}(\sigma_b) + \lambda_{2,j}(\gamma)\,.$$

$\mathcal{A}$ outputs $b^* \in \{0, 1\}$ and wins if $b^* = b$. A friendly jamming scheme is called IND-CSA-n-secure with respect to the linkage procedures $\{\lambda_{i,j}\}_{i=1,2;j=1,\ldots,n}$ if the attacker's advantage $\mathsf{Adv}(\mathcal{A}) = \left|\Pr[\mathcal{A} \text{ wins}] - \frac{1}{2}\right|$ is negligible in χ.

4.6.3.2 IND-CSA-2 Attacker

Given the formalization, the logical next step is to ask if IND-CSA-n security can be achieved for any realistic choice of $\{\lambda_{i,j}\}$. While this might be possible for the IND-CSA-1 case, Tippenhauer et al. [TMRC13] definitively described an IND-CSA-2 attacker. In the following, we use the notation from Definition 4.6.2 with $n = 2$ to describe this attack. The attack assumes a set of linkage procedures $\{\lambda_{1,1}, \lambda_{1,2}, \lambda_{2,1}, \lambda_{2,2}\}$ with the Shift operation (cf. Equation (4.2)) such that

$$\lambda_{1,1} = \lambda_{2,2} = \mathrm{id} \quad \text{and} \quad \lambda_{1,2} = \lambda_{2,1} = \mathsf{Shift}_{\frac{\pi}{2}}\,. \tag{4.14}$$

Tippenhauer et al. [TMRC13] demonstrated that this may indeed be possible in practice if the locations for $\mathcal{A}_1$ and $\mathcal{A}_2$ are chosen such that the four parties $\mathcal{O}_1, \mathcal{O}_2, \mathcal{A}_1, \mathcal{A}_2$ form an isosceles trapezoid, while the following properties must be satisfied.

1. The distance d between $\mathcal{O}_1$ and $\mathcal{O}_2$ is smaller than half of the carrier frequency's wavelength, while the distance between $\mathcal{A}_1$ and $\mathcal{A}_2$ is greater than half this wavelength (and d).

2. The delays between $\mathcal{O}_1$ and $\mathcal{A}_1$, and between $\mathcal{O}_2$ and $\mathcal{A}_2$, respectively, are equal. Likewise, the delays between $\mathcal{O}_2$ and $\mathcal{A}_1$, and between $\mathcal{O}_1$ and $\mathcal{A}_2$, respectively, are also equal.

3. $\mathcal{A}_2$ (or $\mathcal{A}_1$) perceives signals from $\mathcal{O}_1$ (or $\mathcal{O}_2$) with a channel-induced phase shift of $\frac{\pi}{2}$ relative to $\mathcal{A}_1$ (or $\mathcal{A}_2$).

Due to the differing phase offsets (cf. Equation (4.14)), the signals received by $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively, are

$$\begin{aligned}\sigma_1^{\text{col}} &= \lambda_{1,1}(\sigma) + \lambda_{2,1}(\gamma) = \sigma + \mathsf{Shift}_{\frac{\pi}{2}}(\gamma)\,,\\ \sigma_2^{\text{col}} &= \lambda_{1,2}(\sigma) + \lambda_{2,2}(\gamma) = \mathsf{Shift}_{\frac{\pi}{2}}(\sigma) + \gamma\,.\end{aligned}$$

The attacker then shifts and collides these two signals as follows:

$$\begin{aligned}&\mathsf{Shift}_{\frac{3\pi}{2}}\left(\sigma_1^{\text{col}}\right) + \mathsf{Shift}_{\pi}\left(\sigma_2^{\text{col}}\right)\\ &= \mathsf{Shift}_{\frac{3\pi}{2}}(\sigma) + \mathsf{Shift}_{2\pi}(\gamma) + \mathsf{Shift}_{\frac{3\pi}{2}}(\sigma) + \mathsf{Shift}_{\pi}(\gamma)\\ &= \mathsf{Shift}_{\frac{3\pi}{2}}(\sigma) + \mathsf{Shift}_{\frac{3\pi}{2}}(\sigma)\,.\end{aligned}$$

In other words, the attacker recovers an amplified variant of σ, which immediately allows to successfully recover σ.

4.7 Related Work

While several works discuss possible extensions of the Dolev-Yao model, e. g., see Herzog [Her05] or Mao [Mao04], we focus on those that address the connection of the MitMM and wireless attackers.

Pöpper et al. [PTDC11] question the applicability of the MitMM in the context of wireless communication (see also Section 4.5.2). They show the difficulties of symbol flipping and signal annihilation and that these succeed with a certain probability only. Contrary to our work, the proposal of an appropriate model is out of scope.

Schaller et al. [SSBC09] (see also [BCSS09; BCSS11]) developed a formal model for wireless networks that addresses physical properties using inductive, trace-based, symbolic approaches, designed for use with a theorem prover. Some of the aspects we identified for the WCM are also present in their work, e. g., the properties of communication (transmission delays based on distance and propagation speed), location (of network nodes), and time (for temporal dependencies). However, there are several key differences. For instance, messages are treated as events, and transmission time is independent of message length. Modulation

schemes, the resulting signals, and different available communication channels are not considered in their model. As a result, schemes like FHSS and the one proposed by Gollakota et al. [GHR+11] (Section 4.6.3) cannot be fully represented within their model.

Rocchetto and Tippenhauer [RT16] proposed CPDY, an extended Dolev-Yao model for cyber-physical systems, to allow for covering physical interactions between components and the notion of distance. However, they propose and implement new rules in a formal security specification language while focusing on physical interactions in the literal sense, such as an attacker physically manipulating a device, while our focus lies on the communication between parties.

Over the years, different DB protocols evolved, e. g., [AH13; AMSP17; DM07; RC10; RTŠ+12; SP07; TČ09]. In addition, *provable* DB protocols emerged, e. g., a formalism for DB protocols developed by Boureanu et al. [BMV15a; BMV15b] to prove security of their `SKI` family against distance fraud among others. On the basis of ITMs, they developed a communication and attacker model while including the concept of time-of-flight of messages based on the distance between parties. Aspects like modulation schemes or the properties of signals are not covered by the proposed formalism.

Avoine et al. [ABG+21] review relay attacks and the threat model for DB protocols in general. They also consider practical effects, e. g., provers' processing delays, and relate these to theoretical approaches for proving the security of DB protocols. They conclude that formal models for DB protocols are inaccurate, as these make unrealistic assumptions, such as assuming that processing delays during the fast phase are nonexistent or disallowing colluding attackers from communicating during this phase. However, they reach this conclusion without proposing any formal model of their own.

We agree with Avoine et al. [ABG+21] that more realistic models are needed. Our model imposes no additional restrictions on communication between colluding attackers or other parties, neither during the fast phase nor any other phase. The only requirement is that all (wireless) communication must adhere to the physical rules defined by our model, e. g., signal properties and delays. Since our model focuses on communication between parties, processing delays are not included either.

Dürholz et al. [DFKO11] provide a formal (simulator-assisted) analysis of DB RFID protocols. They propose formal models for Mafia, terrorist, and distance fraud, applying them to the RFID DB scheme by Kim and Avoine [KA09] as a case study. While Dürholz et al. also incorporate a global clock, their "time" is message-driven, i. e., a unit of time corresponds to a complete protocol message and is therefore independent of the message's length. On the one hand, they account for noisy communication based on the constraints of RFID and wireless communication; on the other hand, modulation schemes, signals, and channels are deemed irrelevant for the RFID scenario.

Boureanu et al. [BDD+21] developed a parameterized cryptographic model for DB. They highlight potential attacks on existing DB schemes and implement their (parameterized) model in an interactive cryptographic prover, applying it to a contactless payment scheme to prove man-in-the-middle security. In their model, Boureanu et al. define notions of time, location, and distances, and introduce oracles to capture an attacker's capabilities. Some of these oracles bear some similarity to our defined operations. However, their model does not account for modulation schemes, signal properties, or available channels. Additionally, their `replace` oracle allows for targeted replacement of specific message bits for any arbitrary subset of parties, *regardless* of their relative positions. This ignores the fact that different parties may receive message bits at different times depending on their positions. In contrast,

our model requires that an attacker actively sends signals that overlap with the original signals, as this is the only way to influence a received signal. Furthermore, the attacker cannot control the timing at which the original signals reach the receivers.

4.8 Conclusion

Our proposed cryptographic model allows to represent existing mechanisms that build on the peculiarities of wireless communication. We hope that our model will be useful for further analysis and design of cryptographic schemes in the wireless area. We formalized existing schemes (culminating in a DB security proof, the IND-CSA-n game, and the IND-CSA-2 attacker) that can only be represented using the three aspects identified in Section 4.3, i. e., communication channels, signals, and locality. In this sense, these three aspects are necessary. However, they may not be sufficient, as further aspects may need integration into our model in the long run. Finally, we leave it as an open task to extend our model to also include non-wireless communication, e. g., as required to fully represent the IMD shield [GHR+11], or processing time as suggested by Avoine et al. [ABG+21] or as observed in practical jamming applications by O'Flynn [OFl11].

Sir! The radar, sir! It appears to be…jammed!

Radar Technician (Spaceballs)

5 One Antenna to Recover Them All

Abstract

In the previous chapter, we highlighted shortcomings of the established MitMM in a wireless setting and proposed the WCM and WAM tailored specifically to wireless communication. An important observation is that the effect of jamming depends on the modulation in use. As detailed, under certain conditions, it is even possible to reliably recover jammed signals, thus, questioning the efficacy of jamming.

Generally, adversarial jamming aims to make messages incomprehensible for legitimate parties, whereas friendly jamming aims to protect communication from illegitimate eavesdroppers. In both cases, it is crucial to understand if recovering jammed signals may be possible, either as a defensive or an offensive measure. The prevailing assumption is that effective signal recovery requires at least as many antennas as the data sender and the jammer, who use separate antennas, have in total.

This chapter challenges this convention by showing that the effects of jamming can be effectively neutralized using *only one antenna* when BPSK modulation is used. BPSK is a mandatory modulation as per the IEEE 802.15.4 specification which is the basis for Zigbee or WirelessHART among others.

Our work presented in this chapter builds on detailed simulation and practical experiments. Using BRR as performance metric, the BRR is mostly near 1 in simulations; in OTA tests with a 15 dB stronger jamming signal, the average BRR is 0.70 (maximum 0.92). When the jamming signal is 20 dB stronger, the average BRR decreases to 0.34, still allowing for significant signal recovery even with substantial jamming. Promising results of a follow-up experiment with a communication distance of 10 m underscore our method's feasibility in the real-world. Our code and data are publicly available to enable replicability and transparency.

Declaration of Origin: This chapter is based on the paper *One Antenna is Enough for Neutralizing Jamming in IEEE 802.15.4 BPSK* [MMSA25], written together with *Vasily Mikhalev, Yves T. Staudenmaier, and Frederik Armknecht*, and presented at *IWSEC 2025*. First published in *Advances in Information and Computer Security*, Lecture Notes in Computer Science, volume 16208, pp. 489–509, 2025 by Springer Nature. Reproduced with permission from Springer Nature.

5.1 Introduction

With the continuing proliferation of wireless communication protocols and devices implementing them, an increasing amount of communication uses some form of wireless communication [Rap24; SABY23]. Likewise, the possibilities for deliberate disturbances, e. g., in the form of adversarial jamming (see also Section 4.2.2), become more relevant and easier to mount [JZJ22; NN23; PZ22; ZCLY17]. Depending on the specific type of jamming, e. g., constant or reactive, reliable communication might even be completely impossible, the detection of a jammer's presence easy or difficult. On the other hand, utilizing jamming as a defensive mechanism, i. e., *friendly jamming* (see also Section 4.2.3), is a measure that might be able to guard a potentially unsecured network from outsiders [BGF+14; BGF+16; GHR+11; JZJ22; JZZ21; LDW+20; SNHD13]. Here, the recipient does not only see the jammed message but also knows what jamming signals have been used, allowing to ultimately recover the original message. An outsider, however, who lacks this knowledge should not be able to recover the message.

As motivated in Section 4.2.3, friendly jamming may help in scenarios where traditional encryption methods fail, e. g., IMDs cannot easily be upgraded to support a secure encryption scheme as they are implanted and thus hard to physically access. In Sections 4.6.3 and 4.6.3.1, we excerpted and formalized the work originally presented by Gollakota et al. [GHR+11] who developed a friendly jamming device, the *shield*, to protect an IMD. Recall that their goal was to achieve authorization and confidentiality for communication with the IMD (cf. Section 4.6.3), i. e., their *shield* jams communication from and to the IMD, and ultimately recovers the original information for further processing. Evidently, their system is vulnerable to (at least) a two-antenna attacker as Tippenhauer et al. [TMRC13] demonstrated approximately two years after Gollakota et al.'s work was published. With their attacker, Tippenhauer et al. were able to illegitimately recover the (supposedly) protected information. Using our model introduced in Chapter 4, we formalized this type of attacker as IND-CSA-2 attacker in Section 4.6.3.2.

While both, jamming and encryption, try to prevent information from being disclosed without knowing some secret, e. g., a cryptographic key in the encryption case, they operate on different principles.[1] Whereas traditional encryption mechanisms transform the information *directly*, jamming introduces deliberate interference to obscure the transmission encoding the sensitive information (cf. Section 4.2.3). Depending on the specific scenario and mechanisms used, this may lead to information being successfully protected or not. While the construction of secure encryption schemes is well-understood, the situation for (friendly) jamming is different. Besides the used modulation scheme, the actual choice of jamming signals, the ratio of sending to jamming power, i. e., Signal-to-Jamming Ratio (SJR), the relative geographic positions of all involved parties, and the number and distribution of antennas involved heavily affect the success of jamming. Over the years, some specific values, e. g., the range of usable SJR values of $-15\,\mathrm{dB}$ to $-30\,\mathrm{dB}$, have been commonly accepted.

Interestingly, just as there are a multitude of friendly jamming schemes claimed to be secure [BGF+16; GHR+11; JWKD19; JZJ22; SNHD13], there also exist attacks on such schemes [PZ22], ultimately, questioning the security of friendly jamming schemes in general. Most of these attacks focus on the widely-used 2.4 GHz band with QPSK modulation, others

[1] Note that, in the case of *adversarial* jamming, the information is hidden from *legitimate* parties. Interestingly, this is similar to so-called ransomware, e. g., WannaCry [KHT19] or NotPetya [Log17], which use encryption to prevent legitimate parties from reading their data.

target the 403 MHz band and FSK. In addition, techniques like machine learning are also applied to recover jammed communication opening up another angle of attacks [LLJL23; NN23; VNT+20].

Most existing attacks rely on multi-antenna setups (see Table 5.1 for details), and the lack of publicly available code and data makes it difficult to re-evaluate or extend these. This leaves the question whether a weaker attacker, not relying on multiple antennas or machine learning, could still succeed. Showing that such a simpler attacker can recover jammed bits would significantly lower the bar for real-world attacks. We address this question while focusing on transmissions using BPSK modulation.

5.1.1 Contribution

We propose a novel method for recovering jammed bits from real OTA, jammed transmissions, achieving BRRs of up to 0.92 (average 0.70) under 15 dB stronger jamming. In comparison to related work (see Section 5.3), our key contributions are:

- *Single Antenna:* In contrast to existing methods that rely on multiple antennas or machine learning with dedicated training data, our approach uses only a single antenna. Through experiments, we demonstrate that high BRRs are achievable without elaborate setups, thereby challenging prevailing assumptions necessitating more complex setups, e. g., multiple antennas.

- *868 MHz Band:* While most prior work targets the 2.4 GHz or 403 MHz band with QPSK or FSK modulation, respectively, we address the less-studied 868 MHz band, for which BPSK modulation is mandatory in IEEE 802.15.4 [IEE11] which is the basis for protocols like Zigbee or WirelessHART, among others. As this band is common in Europe, is energy-efficient, offers better indoor propagation, and is less congested than 2.4 GHz, it is ideal for WPANs. Our method, however, is *not* frequency-dependent.

- *Simplicity:* Unlike machine learning-based recovery techniques, which require training and careful generalization, our method relies solely on PHY properties of the transmission medium and modulation. This simplicity lowers the bar and allows for more practical, real-world deployments.

- *Reproducibility:* Unfortunately, lack of publicly available code and data hinders reproduction, replication, or extending previous work. To foster transparency and future research, we provide all code and data[2] from both, our simulation and real-world experiments.

5.1.2 Organization of This Chapter

The remainder of this chapter is organized as follows. In Section 5.2, we present the considered system and attacker model, a concise introduction to sub-GHz IEEE 802.15.4, and evaluation metrics used. The state-of-the-art and related work are discussed in Section 5.3, also detailing the differences to our work. Section 5.4 presents our method enabling the recovery of jammed bits. In Section 5.5, we elaborate on the implementation of both, simulation and real-world experiments. Section 5.6 presents the corresponding results. We

[2] https://github.com/TheChrisse/OneAntNeutralizeJamming

discuss implications and limitations of our method in Section 5.7. Section 5.8 concludes this chapter and highlights open questions.

5.2 Background

In this section, we first present the considered system and attacker model in Section 5.2.1, covering the involved parties and their respective roles as well as the type of jamming employed. Then, Section 5.2.2 describes technical details regarding IEEE 802.15.4 in its sub-GHz variant including the relevant modulation and the effect of simultaneous transmissions. The considered evaluation metrics are defined in Section 5.2.3.

5.2.1 System and Attacker Model

Throughout this chapter, we consider (at least) two different parties, i. e., a data sender and a receiver, where the former is sending signals intended for the latter. The data sender and receiver aim to securely exchange messages OTA. We assume all communicating parties share a common (wireless) communication medium, e. g., as specified by IEEE 802.15.4 [IEE11], and that parameters like chosen channel or modulation are known to all parties.

In addition, we assume a jammer jammer. She transmits noise or other electromagnetic signals on the frequencies used by target devices for deliberate interference with communication signals (see also Section 5.2.2.2).

While data sender and jammer are active parties, i. e., they are allowed to transmit at some point in time, the receiver is purely passive, i. e., only listens for transmissions and is thus an eavesdropper. As we motivate in Section 4.2.2.1, we assume the jammer is implemented as a *reactive jammer*, restricting her to reacting to the data sender's activity. This strategy also addresses coordination challenges in legacy systems, where synchronizing a jamming signal with the data sender can be impractical, e. g., the protected data sender might not have a predefined sending schedule, and it might not be upgradable, ultimately restricting the jammer to reacting to the data sender's activity. This is also prevalent in previous works on friendly jamming [BGF+14; JZZ21]. Additionally, we consider the jammer to send correctly BPSK-modulated signals as this is the most effective strategy (cf. Section 4.2.4).

5.2.2 Sub-GHz IEEE 802.15.4

In this chapter, we focus on the 868 MHz band because it is commonly used in Europe, offers lower power consumption, and provides better propagation characteristics for indoor scenarios, while also being less crowded than 2.4 GHz, all advantageous for WPANs. However, our approach itself does not rely on any frequency-specific aspects. Recall that, for the 868 MHz band, BPSK is mandatory according to the specification [IEE11] (cf. Section 4.2.1).

In IEEE 802.15.4, when considering a transmission on the PHY using the 868 MHz band, each bit of a PPDU is first differentially encoded[3], then mapped to 15-bit DSSS *chips*[4], and lastly, modulated[5] using BPSK.

[3] Differential encoding combines the previous and current bit using XOR, i. e., a bit change is encoded as a 1 bit, sequences of the same bit result in sequences of 0 bits.

[4] Each PPDU is encoded into one of two codewords with 15 bits each. For clear distinction from PPDU bits, a single DSSS codeword is referred to by the established term *chip*, whereas a single bit of such a codeword is referred to as a *chip bit*.

[5] Modulation refers to changing characteristics like amplitude, frequency, or phase of a carrier wave in order

5.2.2.1 Binary Phase-Shift Keying

Similar to QPSK, BPSK modulates information onto the carrier wave by varying the carrier's phase (cf. Section 4.4.2.2). The main difference is that BPSK only considers one bit at a time, whereas QPSK considers two bits simultaneously (cf. Definition 4.4.4). Recall the formalized signal representation as presented in Equation (4.1) and Definition 4.4.1 (with scalars as before). For BPSK, there are only two phase shift values available, $\varphi = 0$ and $\varphi = \pi$. This results in a constellation with two real-valued points, i. e., -1 and $+1$.

5.2.2.2 Simultaneous Transmissions

Recall that when at least two transmissions are overlapping, they interfere with each other (cf. Section 4.3.3.1). If they use the same carrier frequency[6] and modulation, they affect each other such that the resulting transmission might not represent any of the original transmissions anymore, e. g., they could completely annihilate each other by destructive interference (cf. Sections 4.3.3 and 4.4.2.4). In essence, considering two signals ψ_0, ψ_1 (represented as in Equation (4.1)) having the same fixed frequency, and fixed amplitudes A_0, A_1 and phase shifts φ_0, φ_1, the resulting signal is another signal with the same frequency, i. e.,

$$\begin{aligned}\psi(t) = \psi_0(t) + \psi_1(t) &= A_0 \cdot \big(\cos(2\pi f t + \varphi_0) + i \cdot \sin(2\pi f t + \varphi_0)\big) \\ &\quad + A_1 \cdot \big(\cos(2\pi f t + \varphi_1) + i \cdot \sin(2\pi f t + \varphi_1)\big) \\ &= A \cdot \big(\cos(2\pi f t + \varphi) + i \cdot \sin(2\pi f t + \varphi)\big)\end{aligned}$$

with

$$A = \sqrt{\big(A_0 \cos(\varphi_0) + A_1 \cos(\varphi_1)\big)^2 + \big(A_0 \sin(\varphi_0) + A_1 \sin(\varphi_1)\big)^2}\,, \tag{5.1}$$

$$\varphi = \arctan\left(\frac{A_0 \sin(\varphi_0) + A_1 \sin(\varphi_1)}{A_0 \cos(\varphi_0) + A_1 \cos(\varphi_1)}\right). \tag{5.2}$$

5.2.3 Evaluation Metrics

We shortly introduce evaluation metrics and abbreviations which are prevalent in previous works and, consequently, used throughout this chapter.

To assess our success in recovering jammed transmissions, we use the widely accepted metric of BRR [JZJ22; JZZ21]. BRR is based on the Bit Error Rate (BER) and defined as

$$\mathrm{BRR} = \frac{|0.5 - \mathrm{BER}|}{0.5} \quad \text{with} \quad \mathrm{BER} = \frac{\#\mathrm{WrongBits}}{\#\mathrm{allBits}}\,. \tag{5.3}$$

When BRR $= 0$, a recovery attempt to decode a message is equivalent to random guessing. Conversely, when BRR $= 1$, the original message can be deterministically deduced, ultimately recovering all (correct) bit(s) [JZZ21, p. 42]. As a rule of thumb, the higher the BRR, the better the method's performance for recovering jammed bits. We also calculate the *average BRR* as unweighted (arithmetic) mean, i. e.,

$$\overline{\mathrm{BRR}} = \frac{1}{n} \sum_{i=1}^{n} \mathrm{BRR}_i\,. \tag{5.4}$$

to transmit information.

[6] If they are using different carrier frequencies, they might not affect each other at all.

Packet Delivery Ratio (PDR) is defined as a ratio of delivered packets to all packets sent [YZJ+16]. In previous works, authors chose this metric as an evaluation metric since it provides insights on a packet level and, thus, on successful communication between two parties in general [YZJ+16]. Conversely, Packet Loss Rate (PLR) describes the fraction of packets which are lost during transmission, i. e., PLR $= 1 -$ PDR [BGF+16].

Furthermore, SJR is a variation of Signal-to-Noise Ratio (SNR) which is defined as SNR $= P_s - P_n$, where P_s and P_n are the power level of signal and noise in dB, respectively [Gol05, p. 30]. Thus, SJR represents the ratio between the power of the transmitted signal to the power of the jamming signal, while a negative value indicates that the jamming signal is stronger.

5.3 Related Work

This section summarizes the state-of-the-art in recovery of information subjected to adversarial and friendly jamming. First, Section 5.3.1 provides a brief overview on the foundations and practical implementations of friendly jamming. Section 5.3.2 presents related work in detail, thereby highlighting the differences to our approach.

5.3.1 Foundations and Implementations of Friendly Jamming

Friendly jamming (or cooperative jamming) dates back to Wyner's [Wyn75], and Csiszár and Korner's [CK78] 1970s studies, which laid the theoretical groundwork. Later, emphasizing the potential of friendly jamming in secure communications, Vilela et al. [VBBM10; VBBM11] expanded on these foundations. The practical deployment of friendly jamming and its implementation challenges in real-world scenarios have been explored in several key studies. Papers by Berger et al. [BGF+14; BGF+16], Gollakota et al. [GHR+11], Li et al. [LDW+20], and Martinovic et al. [MPS09] discuss these practical aspects, focusing on how friendly jamming can be implemented effectively to enhance communication security.

In the field of friendly jamming, several advanced techniques have been proposed to enhance both, practicality and security. Shen et al. [SNHD13] developed Ally Friendly Jamming, a strategy that disables unauthorized wireless communications using intelligently controlled jamming signals. Rahbari and Krunz [RK14] proposed Friendly CryptoJam, an approach that combines analog-domain friendly jamming with modulation-level encryption to protect PHY fields from adversaries. Ma et al. [MDL+22] introduced an Active Anti-Jamming (AAJ) scheme that utilizes the jamming signal as a medium for communication by employing a programmable-gain amplifier to re-modulate jammed signals, thus enhancing channel capacity under high jamming-to-noise ratios. Concurrently, Ma et al. [MLL+23] explored a similar AAJ technique, providing detailed analyses on BERs and optimal detection thresholds for various jamming types, demonstrating superior channel capacity over direct transmissions in environments with strong jamming. Yaman et al. [YAE23] introduced a lightweight friendly jamming scheme tailored to IoT devices, which stands out for its simplicity and effectiveness. It maintains performance levels comparable to more complex existing schemes while requiring minimal computational and energy resources.

5.3.2 Comparing Our Method with Others

Although the focus in this chapter is particularly on attacks targeting friendly jamming, we consider the following two scenarios:

1. An adversary attempts to disrupt the communication, but the receiver tries to cancel out the jammer's signal in order to recover the original signal.

2. An attacker is trying to eavesdrop on communication that is being friendly-jammed.

Both scenarios are relevant for comparing our work with prior studies, as they primarily involve recovering original signals from jammed signals. For a concise overview of existing papers and to allow for a detailed comparison to ours, we refer to Table 5.1 which outlines various studies addressing these specific scenarios. To enable comparison across different works, we homogenized the evaluation metrics shown in Table 5.1. Specifically, we converted BER to BRR and PLR to PDR (cf. Section 5.2.3). PDR cannot be directly converted to BRR because PDR does not provide information about the number of erroneous bits in packets considered undelivered.

A significant point of comparison in our research is the ability of various systems to recover signals using a different number of receiving antennas. According to the rule of thumb discussed by Jin et al. [JZJ22], a Multiple-Input and Multiple-Output receiver with N antennas can separate and decode N signals transmitted concurrently on the wireless medium. This principle is well-supported by several works shown in our comparison table which demonstrate effective results using two antennas to separate the original and jammed signals.

One of the first papers to address the limitations of friendly jamming was by Tippenhauer et al. [TMRC13], which specifically discusses the constraints and vulnerabilities of such techniques. Investigating the friendly jamming scheme originally proposed by Gollakota et al. [GHR+11], they demonstrate that with two antennas, it is feasible to recover the original (supposedly protected) signals with a relatively high BRR even when the jamming signal is 30 dB stronger than the communication signal.

A more sophisticated technique was investigated in the study by Nguyen and Noubir [NN23] utilizing advanced tools like convolutional neural networks. However, instead of using actual *jammed signals* in practice, they synthesize data by simulating how the jammed signal merges with the original one. Since they employ a different performance metric — specifically, PLR, which we have converted to PDR in Table 5.1 — our results are not directly comparable. Nonetheless, it is evident that their method's effectiveness significantly declines at an SJR of −20 dB, which aligns with our findings in a similar scenario but using only one antenna.

Vo-Huu et al. [VBN13] present a mechanical beam-forming approach combined with a fast auto-configuration algorithm and a digital interference cancellation technique. They achieve reliable communication in presence of a 48 dB stronger jammer.

In contrast, this chapter presents findings that demonstrate the ability to recover a significant portion of the original signal using just *one antenna*, challenging the conventional belief that more antennas necessarily provide better separation and decoding of signals.

Another study by Jin et al. [JZZ21] shows that, using a single antenna, one can recover the original signal from a jammed one. However, as indicated in Table 5.1, our method exhibits considerably enhanced performance in scenarios with stronger jammers.

Table 5.1: Comparison of research studies on data recovery under adversarial and friendly jamming scenarios, values homogenized, SJR in dB. If a paper presents several performance results under different scenarios, we consider the best one. '—' indicates not available, Freq. the frequency, BW bandwidth (in MHz), #A number of antennas used, and DCA the data and code availability with '✓' indicating availability.

Work	Modulation	SJR	Freq.	BW	#A	Performance	DCA
[TMRC13]	FSK	$\{-20, -25, -30, -35\}$	403 MHz	0.3	2	BRR > 0.6 (SJR $= -30$ dB)	—
[VBN13]	BPSK, QPSK	$[-50; 0]$	2.4 GHz	—	2	BRR > 0.8 (SJR $= -48$ dB)	—
[YZJ+14]	OFDM + BPSK	0	2.45 GHz	0.5 to 1	2	PDR > 0.6 (500 kHz BW)	—
[ZCLY17]	OFDM + QPSK	$\{0, -20\}$	2.4 GHz	—	≥ 1	BRR > 0.98 (2, 3, 4 ant.*, SJR $\in [-20; -17]$)	—
[JZZ21]	OQPSK	-10 (average)	2.4 GHz	2	1	$\overline{\text{BRR}} = 0.21$	—
[NN23]	{B, Q, 8}-PSK, 16-QAM	$\{-10, -20\}$	—	—	2	PDR > 0.9 (SJR $= -10$ dB), PDR ≈ 0.5 (SJR $= -20$ dB, BPSK)	—
Ours	BPSK	$\{-15, -20, -25, -30\}$	868.3 MHz	0.6	1	$\overline{\text{BRR}} = 0.7$ (SJR $= -15$ dB)	✓

* To the best of our knowledge, the method in [ZCLY17] does not yield any significant results when using only one antenna.

To the best of our knowledge, none of the studies exploring the recovery of the original signal from jammed signals in the scenarios mentioned earlier publish their code or data, which would facilitate result replication. In contrast, we provide access to our code and data in order to support verification and with the hope of stimulating further research.

5.4 Attack Description

In the following, we introduce a new method that allows for the recovery of jammed bits. For clarity and as we consider this the more security-critical use case, we present it in the context of friendly jamming (cf. Section 4.2.3). Accordingly, we frame the method as an *attack* by an outsider attempting to recover friendly-jammed bits. Notably, the same method can serve as a defense against adversarial jamming.

First, we provide an overview of the attack scenario and objectives in Section 5.4.1. We then give a high-level, illustrative sketch of our attack in Section 5.4.2, followed by a more detailed and precise description in Section 5.4.3. The concept of the attack was initially explored and validated through simulation. This theoretical examination demonstrated the general feasibility of our approach. In an exploratory practical setting, we encountered and addressed real-world challenges like background noise, interference, and alignment issues. We detail both the simulated and practical attacks in Section 5.5. The simulation also serves as a baseline for replicability.

5.4.1 Attack Scenario and Objectives

The *objective* of our attack is to recover the original communication obscured by jamming using only *one* antenna, even though data sender and jammer use separate antennas.

The transmission process begins with the transmission of the original signal by the data sender. The receiver, as outlined in Section 5.2.1, also acts as a jammer and monitors for the signal initiation. Once the transmission is detected, the receiver starts emitting jamming signals. Due to the detection and response time, the initial part of the signal is transmitted without any jamming interference. Jamming is then performed via concurrent transmission of (random) chip bits to obscure the signal.

The attacker is passive, i. e., she records waveforms without knowing transmission start (or end) times, resulting in a waveform W comprising three parts:

1. pre-transmission background noise,
2. an unjammed preamble due to the jammer's delay, and
3. the remainder of the transmission including a jammed part.

Mathematically, $W = (W_1, W_2, \ldots)$ is a sequence of complex samples representing the signal ψ (cf. Equation (4.1)). Ideally, BPSK values are ± 1, but real-world effects like noise, interference, and misalignments cause deviations, producing scattered complex values.

5.4.2 Attack Procedure

Before we detail our attack, we first provide a brief overview of it, followed by illustrative examples from the simulation and real-world scenarios.

5.4.2.1 Our Attack in a Nutshell

Using BPSK modulation, only two defined phase values are available (cf. Section 5.2.2.1). Consequently, the receiver needs to identify the phase shift to recognize the bit being transmitted. Now, if a transmission is jammed, the stronger jamming signal dominates the resulting signal's phase shift. The overall goal is to estimate the data sender's and the jammer's signal strengths, and to use this knowledge to infer the original signal.

More precisely, recall that the captured waveform W can be divided into three parts. In the attack description below, we use $N_{\text{data}}^{\text{begin}}$ and $N_{\text{data}}^{\text{end}}$ to refer to the beginning and end of a transmission, respectively. Similarly, we use $N_{\text{jam}}^{\text{begin}}$ and $N_{\text{jam}}^{\text{end}}$ to refer to the jammed region within the transmission. That is, we have

$$W = \Big(\underbrace{W_1, \ldots, W_{N_{\text{data}}^{\text{begin}}-1}}_{\text{Part 1}}, \underbrace{W_{N_{\text{data}}^{\text{begin}}}, \ldots, \overbrace{W_{N_{\text{jam}}^{\text{begin}}}, \ldots, W_{N_{\text{jam}}^{\text{end}}}}^{\text{Part 3}}, \ldots, W_{N_{\text{data}}^{\text{end}}}}_{\text{Part 2}} \Big) . \tag{5.5}$$

Note that, a priori, the attacker does not know when the individual parts start and end, so these need to be found during the attack. Once accomplished, the attacker uses a small portion of Part 2, $\left(W_{N_{\text{data}}^{\text{begin}}}, \ldots, W_{N_{\text{data}}^{\text{end}}}\right)$, to estimate the data sender's signal strength. Together with the observed maximum amplitude of the remaining Part 3, $\left(W_{N_{\text{jam}}^{\text{begin}}}, \ldots, W_{N_{\text{jam}}^{\text{end}}}\right)$, she can also estimate the jammer's signal strength. For each jammed bit, the attacker uses these estimations to create two candidates for the data sender's value. Ultimately, she compares these two candidates with the observed value to decide which bit the data sender actually used.

5.4.2.2 Illustration

As an example, Figure 5.1 shows the first 10,000 signal samples of a simulated, jammed transmission (without preceding noise-only part). In the beginning of this plot, the signal is visible while unimpaired from jamming. Then, the jamming starts, yielding a signal with generally higher amplitudes. If the data sender's and the jammer's chip bits are identical, the jammer's amplitude is increased by the data sender's amplitude. If the chip bits differ, the jammer's amplitude is decreased by approximately the data sender's.

Figure 5.2 shows a partial plot containing three consecutively transmitted packets obtained during the real-world evaluation. Contrary to Figure 5.1, the plot contains regions with only background noise, e. g., in the beginning. Before the actual bit recovery process, the different parts' positions need to be identified, i. e., parts containing only background noise, unjammed, and jammed parts of any transmission.

Figure 5.3 displays the second message visible in Figure 5.2 in more detail. Note the pronounced imaginary parts in the beginning of both, unjammed parts and throughout the jammed part, indicating the phase offsets between the involved parties.

5.4.3 Detailed Attack Description

The steps of our proposed attack are visually represented as a flowchart in Figure 5.4 and described in the following.

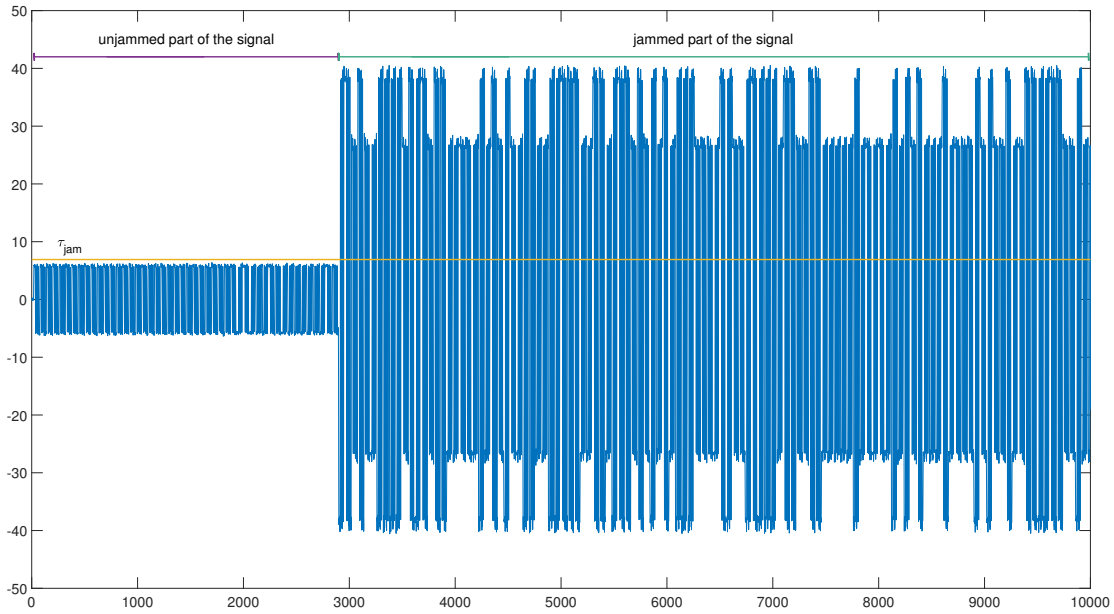

Figure 5.1: Jamming plot from the simulation experiment, showing the unimpaired PHR, the jamming detection threshold τ_{jam} (horizontal yellow line), and the combined/jammed waveform.

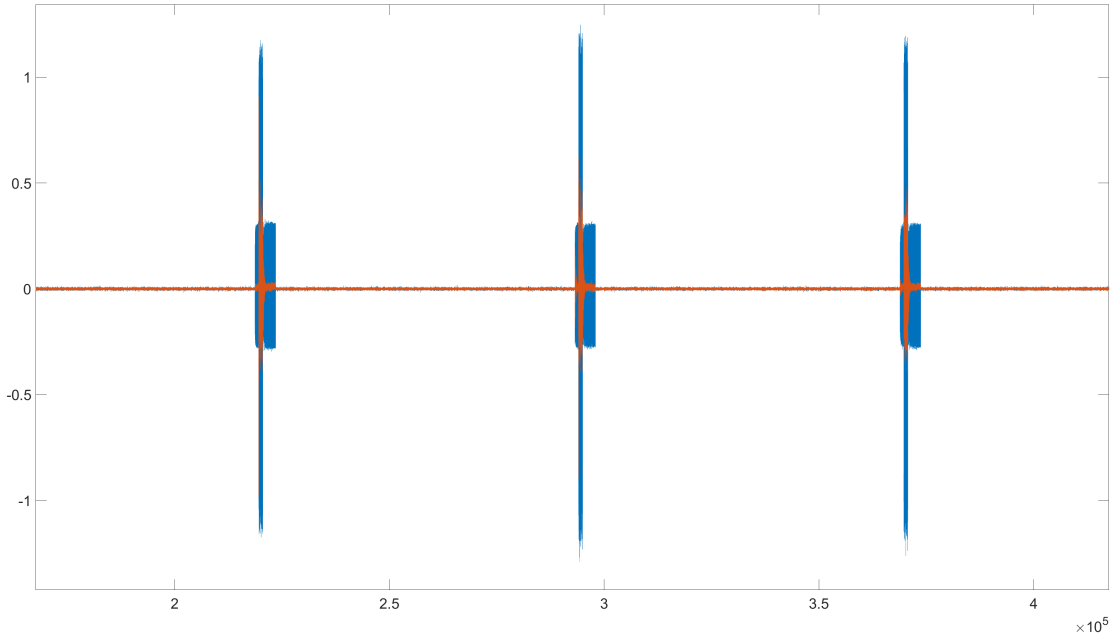

Figure 5.2: Waveform plot of three jammed messages from the real-world experiment with blue representing real and orange the imaginary part.

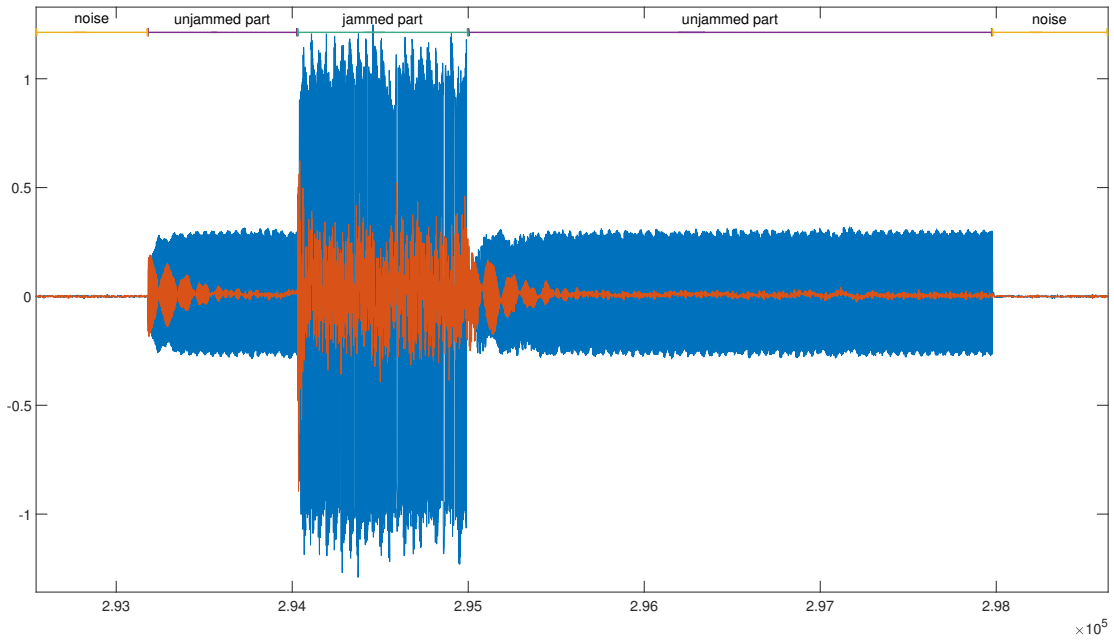

Figure 5.3: Detailed waveform plot of the jammed second message visible in Figure 5.2 with preceding and trailing noise obtained during the real-world experiment, with blue representing real and orange the imaginary part.

1. The first step of the attack involves determining $N_{\text{data}}^{\text{begin}}$, i. e., the beginning of a transmission in the captured waveform W (cf. Equation (5.5)). To this end, the attacker assumes that the first few samples of the waveform only represent background noise (Part 1 in Equation (5.5); see also beginning of Figures 5.2 and 5.3). Consequently, she uses these to calculate the background noise level as the maximum observed at the very beginning of the waveform W, i. e.,

$$A_{\text{noise}} = \max_{i=1..N} \left\{ \left| W_i \right| \right\}, \tag{5.6}$$

where N is chosen such that $N < N_{\text{data}}^{\text{begin}}$ holds with high probability. Subsequently, a data threshold, τ_{data}, is calculated as

$$\tau_{\text{data}} = \alpha_{\text{noise}} \cdot A_{\text{noise}}\,, \tag{5.7}$$

where α_{noise} is a scaling factor[7] to ensure that the threshold represents an amplitude significantly exceeding the background noise. If N consecutive samples have amplitudes exceeding τ_{data}, then the attacker assumes the presence of a transmission, and uses the first such sample's position for $N_{\text{data}}^{\text{begin}}$ (the beginning of Part 2 in Equation (5.5); see also beginning of Figure 5.1). The end of a transmission, $N_{\text{data}}^{\text{end}}$ (the end of Part 2 in Equation (5.5); see also end of Figure 5.3), is found similarly, i. e., finding N consecutive samples with amplitudes below τ_{data}.

[7] We elaborate on the choice of scaling factors in Section 5.5.3.

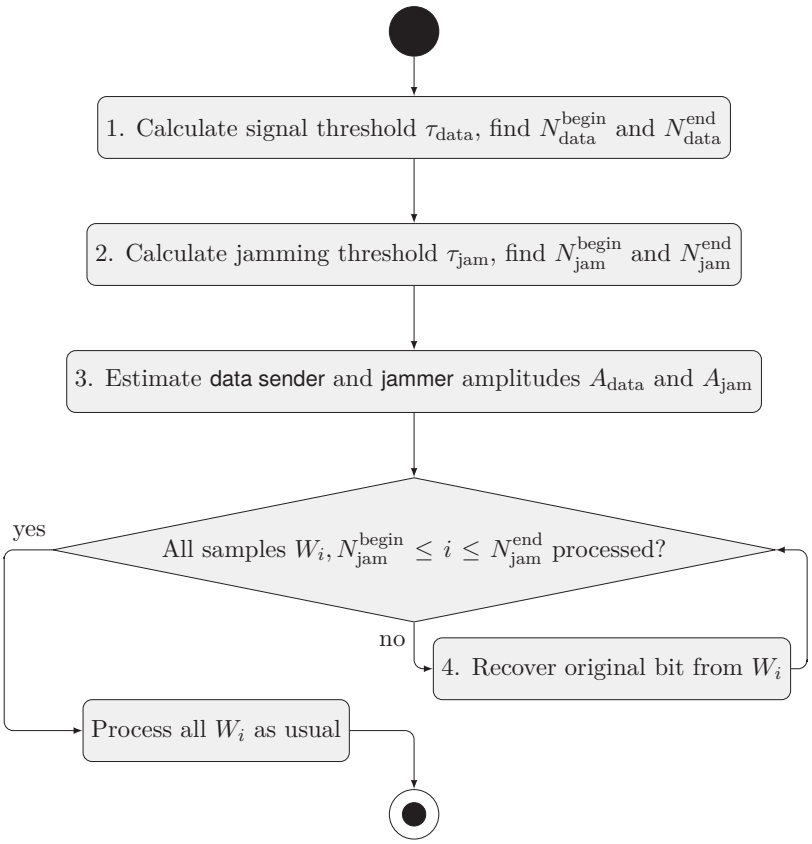

Figure 5.4: Flowchart illustrating the processing steps for recovering jammed bits.

2. Analogously to Step 1, the beginning of a detected transmission (unaffected by the jammer) is used to obtain the maximal amplitude of the data sender's signal, i. e.,

$$A_{\text{data}}^{\max} = \max\left\{\left|W_{N_{\text{data}}^{\text{begin}}}\right|, \ldots, \left|W_{N_{\text{data}}^{\text{begin}}+M-1}\right|\right\}, \tag{5.8}$$

where M denotes the subsequent number of (unjammed transmission) samples considered. M is chosen to *not* overlap with the jammed part in W, i. e., $N_{\text{data}}^{\text{begin}} + M - 1 < N_{\text{jam}}^{\text{begin}}$. With this, a jamming threshold, τ_{jam} (see horizontal yellow line in Figure 5.1), is defined to identify sections with considerably higher amplitude than $A_{\text{data}}^{\max}$, indicating ongoing jamming to obfuscate the original signal, i. e.,

$$\tau_{\text{jam}} = \alpha_{\text{data}}^{\max} \cdot A_{\text{data}}^{\max}, \tag{5.9}$$

where $\alpha_{\text{data}}^{\max}$ is a scaling factor chosen to ensure that the threshold represents an amplitude indicating a jammed signal. Once M consecutive samples exceed τ_{jam}, the attacker considers this the beginning of the jammed region, $N_{\text{jam}}^{\text{begin}}$ (the beginning of Part 3 in Equation (5.5); see also high-amplitude part starting close to sample 2.94×10^5 in Figure 5.3). As soon as M consecutive samples are below τ_{jam}, the attacker found the end of the jammed region, $N_{\text{jam}}^{\text{end}}$ (the end of Part 3 in Equation (5.5); see also high-amplitude part ending close to sample 2.95×10^5 in Figure 5.3).

3. The attacker uses $A_{\text{data}}^{\max}$ (cf. Equation (5.8)) and the maximal signal amplitude, $A_{\text{jam}}^{\max}$, to estimate the data sender's and the jammer's signal amplitudes, A_{data} and A_{jam}, as

$$A_{\text{jam}}^{\max} = \max\left\{ \left| W_{N_{\text{jam}}^{\text{begin}}} \right|, \ldots, \left| W_{N_{\text{jam}}^{\text{begin}}+M-1} \right| \right\}, \tag{5.10}$$

$$A_{\text{data}} = 0.97 \cdot A_{\text{data}}^{\max}, \tag{5.11}$$

$$A_{\text{jam}} = \alpha_{\text{jam}}^{\max} \cdot A_{\text{jam}}^{\max} - \frac{A_{\text{data}}}{2}, \tag{5.12}$$

where $\alpha_{\text{jam}}^{\max}$ is a scaling factor.[8] These estimations use the unjammed preamble to benchmark normal data sender signal levels, and the jammed region to assess the strength and impact of the jammer's signal.

4. For each jammed sample W_i, $N_{\text{jam}}^{\text{begin}} \leq i \leq N_{\text{jam}}^{\text{end}}$, the attacker tries to recover the original value, denoted W_i^*. To this end, she constructs two potential phase shift values, one for each possible data sender phase shift value $\varphi_{\text{data}} \in \{0, \pi\}$ (cf. Section 5.2.2.1), to estimate the jammer's phase shift value. Based on Equation (5.1) and the previously estimated amplitudes of data sender and jammer (cf. Step 3, Equations (5.11) and (5.12)), she calculates for the two possible phase shift values 0 and π, the expected phase shift values (of the jammed sample) φ_0 and φ_π, respectively. First, we rewrite and transform Equation (5.1) as follows by squaring and resolving the parentheses.

$$\begin{aligned}
A &= \Big(\big(A_{\text{data}} \cos(\varphi_{\text{data}}) + A_{\text{jam}} \cos(\varphi_{\text{jam}}) \big)^2 \\
&\quad + \big(A_{\text{data}} \sin(\varphi_{\text{data}}) + A_{\text{jam}} \sin(\varphi_{\text{jam}}) \big)^2 \Big)^{1/2} \\
\Leftrightarrow \quad A^2 &= \big(A_{\text{data}} \cos(\varphi_{\text{data}}) + A_{\text{jam}} \cos(\varphi_{\text{jam}}) \big)^2 \\
&\quad + \big(A_{\text{data}} \sin(\varphi_{\text{data}}) + A_{\text{jam}} \sin(\varphi_{\text{jam}}) \big)^2 \\
\Leftrightarrow \quad A^2 &= A_{\text{data}}^2 \cos^2(\varphi_{\text{data}}) + 2A_{\text{data}}A_{\text{jam}} \cos(\varphi_{\text{data}}) \cos(\varphi_{\text{jam}}) + A_{\text{jam}}^2 \cos^2(\varphi_{\text{jam}}) \\
&\quad + A_{\text{data}}^2 \sin^2(\varphi_{\text{data}}) + 2A_{\text{data}}A_{\text{jam}} \sin(\varphi_{\text{data}}) \sin(\varphi_{\text{jam}}) + A_{\text{jam}}^2 \sin^2(\varphi_{\text{jam}})
\end{aligned}$$

With the Pythagorean trigonometric identity $\sin^2\theta + \cos^2\theta = 1$, we can further simplify the intermediate result:

$$\begin{aligned}
\Leftrightarrow \quad A^2 &= A_{\text{data}}^2 \left(\cos^2(\varphi_{\text{data}}) + \sin^2(\varphi_{\text{data}}) \right) + A_{\text{jam}}^2 \left(\cos^2(\varphi_{\text{jam}}) + \sin^2(\varphi_{\text{jam}}) \right) \\
&\quad + 2A_{\text{data}}A_{\text{jam}} \big(\cos(\varphi_{\text{data}}) \cos(\varphi_{\text{jam}}) + \sin(\varphi_{\text{data}}) \sin(\varphi_{\text{jam}}) \big) \\
\Leftrightarrow \quad A^2 &= A_{\text{data}}^2 + A_{\text{jam}}^2 + 2A_{\text{data}}A_{\text{jam}} \big(\cos(\varphi_{\text{data}}) \cos(\varphi_{\text{jam}}) + \sin(\varphi_{\text{data}}) \sin(\varphi_{\text{jam}}) \big)
\end{aligned}$$

We rearrange and multiply to collect all sine and cosine terms on one side, all terms involving amplitudes on the other side:

$$\begin{aligned}
\Leftrightarrow \quad A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2 &= 2A_{\text{data}}A_{\text{jam}} \big(\cos(\varphi_{\text{data}}) \cos(\varphi_{\text{jam}}) + \sin(\varphi_{\text{data}}) \sin(\varphi_{\text{jam}}) \big) \\
\Leftrightarrow \quad \frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2A_{\text{data}}A_{\text{jam}}} &= \cos(\varphi_{\text{data}}) \cos(\varphi_{\text{jam}}) + \sin(\varphi_{\text{data}}) \sin(\varphi_{\text{jam}})
\end{aligned}$$

[8] The scaling factor 0.97 was empirically determined to compensate for coincidental cases in which the background noise peaks, thus, increasing the data sender's signal amplitude more strongly than usual.

Now, considering that there are exactly two defined phase values (0 and π) and using these, e. g., for φ_{data}, we end up with two equations, in which the sine product can be omitted as it results in zero in both cases. For the cosine product, we know that $\cos(0) = -\cos(\pi) = 1$, i. e., we can simplify even further in order to calculate an expected phase, e. g., the jammer's phase, φ_{jam}, under the assumption that the data sender's phase, φ_{data}, is either 0 or π.

$$
\begin{aligned}
& \pm\cos(\varphi_{\text{jam}}) = \frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2A_{\text{data}}A_{\text{jam}}} \\
\Leftrightarrow \quad & \cos(\varphi_{\text{jam}}) = \pm\frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2A_{\text{data}}A_{\text{jam}}} \\
\Leftrightarrow \quad & \varphi_{\text{jam}} = \arccos\left(\pm\frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2A_{\text{data}}A_{\text{jam}}}\right) \\
\Leftrightarrow \quad & \varphi_{\text{jam}} = \begin{cases} \arccos\left(\frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2A_{\text{data}}A_{\text{jam}}}\right) & \text{if } \varphi_{\text{data}} = 0 \\ \pi - \arccos\left(\frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2A_{\text{data}}A_{\text{jam}}}\right) & \text{if } \varphi_{\text{data}} = \pi \end{cases}
\end{aligned}
$$

Note that $\arccos(-\theta) = \pi - \arccos(\theta)$, and recall that any two phase values, $\varphi_1 = \theta$, $\varphi_2 = k \cdot 2\pi + \varphi_1, k \in \mathbb{Z}$, i. e., separated by (a multiple of) 2π, can be considered equivalent (cf. Section 4.4.2.1). Also note that if $\arccos(\theta) = 0$ or $\arccos(\theta) = \pi$, respectively, then $\arccos(\theta)$ and $-\arccos(\theta)$ are equivalent. In other words, we can describe the jammer's phase depending on the (hypothesized) data sender's phase as

$$
\varphi_0 = 0 - \arccos\left(\frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2 \cdot A_{\text{data}} \cdot A_{\text{jam}}}\right), \quad \varphi_\pi = \pi - \arccos\left(\frac{A^2 - A_{\text{data}}^2 - A_{\text{jam}}^2}{2 \cdot A_{\text{data}} \cdot A_{\text{jam}}}\right), \tag{5.13}
$$

with φ_0 representing the case that the data sender's phase is 0, and φ_π that it is π. In general, we can assume that the result of two concurrent signals is dominated by the stronger (jamming) signal (cf. Equations (4.3), (4.4), (5.1) and (5.2)), i. e., the resulting signal's *observed* phase, φ, should be closer to the dominant jammer's phase, φ_{jam}. Consequently, we can derive the hypothesis that $\varphi \approx \varphi_{\text{jam}}$. In combination with the observed phase, we can then derive our decision criterion based on the (absolute) differences between these calculated values (cf. Equation (5.13)) and the observed phase in order to derive which hypothesis, i. e., $\varphi_{\text{data}} = 0$ or $\varphi_{\text{data}} = \pi$ holds. Thus, W_i^* is estimated as[9]

$$
W_i^* = \begin{cases} 0, & \text{if } |\varphi_0 - \varphi| > |\varphi_\pi - \varphi| \text{ ,} \\ 1, & \text{else.} \end{cases} \tag{5.14}
$$

5.5 Implementation

This section explores the implementation of our attack from Section 5.4 through both, simulation-based (Section 5.5.1) and real-world experiments (Section 5.5.2), which are de-

[9] Technically, the waveform samples are shaped as pulses of $\{-1; +1\}$ which directly represent chip bits. We omit this technicality for clarity and brevity.

signed to investigate the effectiveness of the proposed methods for countering jamming in wireless communication. Furthermore, Section 5.5.3 reports chosen values like scaling factors required for our attack (cf. Section 5.4.3) used in the simulation and real-world experiments.

5.5.1 Simulation Experiment

To assess the feasibility of our approach for recovering jammed bits (cf. Section 5.4.2), we implemented an end-to-end simulation for IEEE 802.15.4 in Matlab. In a nutshell, standard functions and modules are used where possible, ensuring generated packets and waveforms adhere to the specification [IEE11].

More precisely, we use the `lrwpan.PHYGeneratorBPSK` and `lrwpan.PHYDecoderBPSK` modules which are available in Matlab R2024a. The data sender constructs PPDUs as required by the standard (cf. Figure 4.1 and Sections 4.2.1 and 5.2.2), i. e., including the required spreading and filtering as defined in the standard [IEE11]. For simulating the payload (or PHY Service Data Unit (PSDU)), a fixed amount of random bits is generated. Then, the resulting PPDU is processed by going through the differential encoding, the bit-to-chip mapping, i. e., spreading [IEE11, Table 78], and the modulator resulting in a waveform that is standard-conforming.

As motivated in Section 5.2.1, the jammer is simulated as a reactive jammer (cf. Section 4.2.2.1), i. e., the jammer's activity is delayed until after the PHR. Then, the data sender's waveform and the jammer's waveform are combined into a single waveform which is subjected to additive white Gaussian noise (using the integrated Matlab function `awgn`) to simulate real-world background noise. Additionally, we consider that data sender and jammer might have a small offset in the transmission process denoted *Chip Shift*. This is due to the fact that a chip bit is not directly transmitted, instead, *oversampling* is used to generate a number of consecutive samples (usually 4) constituting the chip bit to be transmitted. Hence, the offset between the data sender's and the jammer's (over-)sampling process is captured by the Chip Shift parameter. The following SNR, SJR, and Chip Shift values are considered for the simulation:

- The data sender's SNR is approximately 30 dB.
- SJR $\in \{n \mid n = -(15 + 5k),\ k \in \mathbb{Z},\ 0 \leq k \leq 3\}$.
- Chip Shift $\in \{0, 1, 2, 3\}$.

Instead of generating valid PPDUs, the jammer generates a fixed amount of random bits. However, the jammer can choose to either generate random *message* or random *chip* bits, which we refer to as *Message Bit Jamming* or *Chip Bit Jamming*, respectively.

5.5.1.1 Message Bit Jamming

In this scenario, the jammer generates random message bits following the same processing steps as the data sender. That is, the jammer chooses random message bits and the waveform is generated by going through the differential encoder, spreading (thus, the chips conform to the specification), and modulation (cf. Section 5.2.2.1). The example shown in Figure 5.1 demonstrates the first 10,000 signal samples of a jammed transmission (using this form of jamming with an SJR of −15 dB) as discussed in Section 5.4.2.

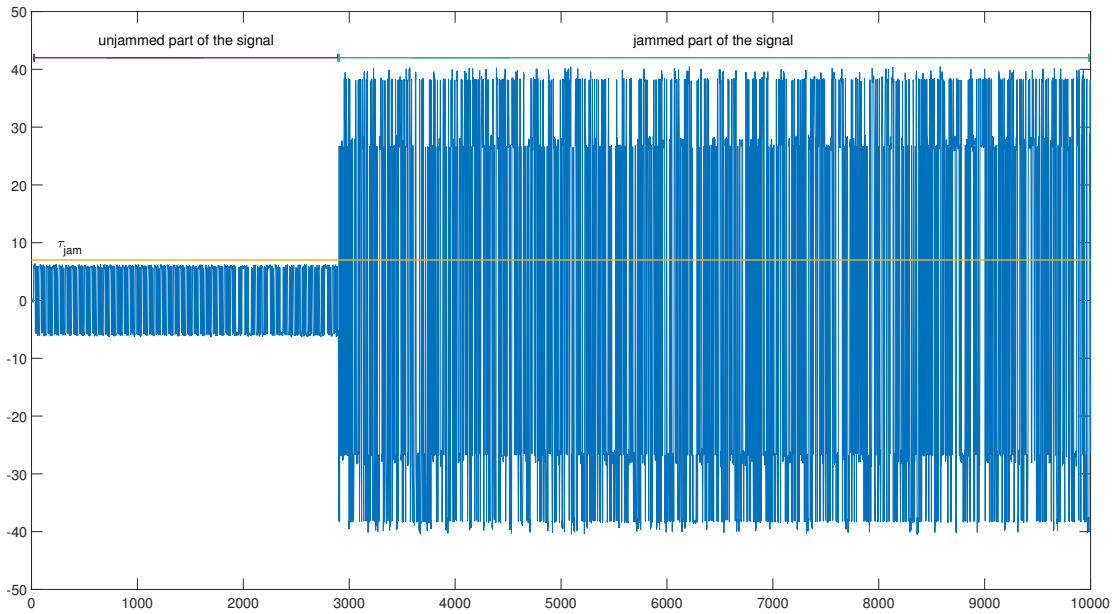

Figure 5.5: Chip Bit Jamming plot showing the unimpaired PHR, the jamming detection threshold τ_{jam} (horizontal yellow line), and the combined/jammed waveform.

5.5.1.2 Chip Bit Jamming

In the Chip Bit Jamming scenario, the jammer selects and transmits random *chip bits* instead of message bits, i. e., BPSK-modulates and sends a fixed number of random *chip* bits that are *not* standard-compliant, as they do not represent 0 or 1 message bits. This also incurs a variation in the waveform, the impact of which is visible in Figure 5.5, especially in comparison to Figure 5.1.

Although the process of generating the jamming signal differs (focusing on chip bits rather than message bits), the attack procedure remains the same. By maintaining this methodological consistency, we ensure that the efficacy of our approach to recover jammed bits is robust across different types of jamming scenarios.

5.5.2 Real-World Experiment

To validate our simulation-based findings, we conducted a real-world experiment utilizing SDRs and GNU Radio. We implemented the IEEE 802.15.4 868 MHz PHY [IEE11, Section 11] for a data sender, a (reactive) jammer, and an attacker partially from scratch while using available blocks where possible.

5.5.2.1 Test Setup

The test setup comprises a desktop PC (running Ubuntu Linux 22.04 LTS) with an Ettus B210[10] USRP [Ett18a] and two Ettus LP0965 log-periodic antennas[11] [Ett12], and a laptop (running Kali Linux 2024.1) with an Ettus B200mini[10] USRP [Ett18b] with one LP0965 and

[10] https://kb.ettus.com/B200/B210/B200mini/B205mini
[11] https://kb.ettus.com/Antennas#LP0965

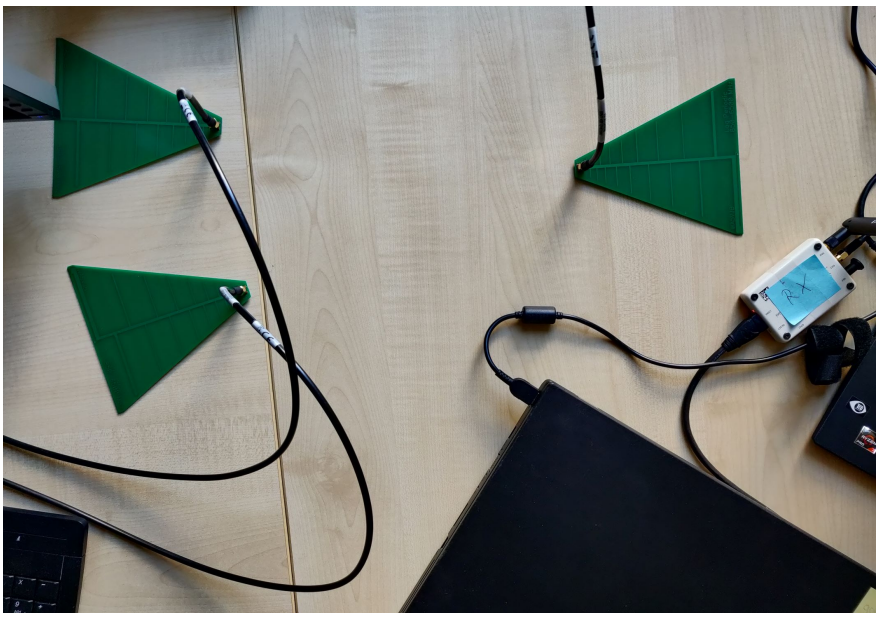

Figure 5.6: Lab setup for the real-world experiment showing three LP0965 and one omnidirectional antenna. The left antennas are connected to the B210; the upper one is used by the jammer, the lower one by the data sender. For receiving, the attacker uses either the right LP0965 or the omnidirectional antenna which are connected to the B200mini.

one Delock 89769 omnidirectional antenna[12] [Del23], where only one of which is used during the experimental runs. The software implementation is based on GNU Radio version 3.10[13], with USRP Hardware Driver support[14], the gr-ieee802-15-4[15] [BLDS13] and the required gr-foo[16] modules. Data sender and jammer are implemented on the desktop PC, the attacker on the laptop. Figure 5.6 shows the lab setup with a measured distance of 31 cm from the tips of the jammer's (upper left) and data sender's (lower left) antennas to the log-periodic attacker antenna (right side), resulting in a distance of 59 cm between the far side of these antennas. The distance from the tips of the data sender's and the jammer's antennas to the attacker's omnidirectional antenna is approximately 54 cm, the distance between the tips of the data sender's and jammer's antenna is 14 cm. Thus, data sender and jammer are placed less than half a wavelength from each other, making their signals practically inseparable from a third party's perspective as they are effectively using the same channel [ESWM12; ZZP+15].

Ethical Considerations. To reduce unintentional effects of our experiments on others, we considered the following precautions. First, we checked for general communication activity which was very rarely and only barely noticeable, i. e., if we detected some communication activity, it had very low amplitude. Second, we started with the lowest transmit gain setting to establish the minimum gain required for reliable communication and only exceeded this

[12] https://www.delock.com/produkt/89769/merkmale.html
[13] https://github.com/gnuradio/gnuradio/tree/maint-3.10
[14] https://github.com/EttusResearch/uhd/tree/UHD-4.5
[15] https://github.com/bastibl/gr-ieee802-15-4/tree/maint-3.10
[16] https://github.com/bastibl/gr-foo/tree/maint-3.10

if necessary, e. g., due to increased distance. Furthermore, our implementation ensures that the jammer only targets transmissions from our data sender and only jams a limited portion of the transmission, i. e., the jammer starts jamming after the data sender and stops before the data sender, thereby limiting the time the channel is occupied to the duration of the data sender's transmission.

5.5.2.2 GNU Radio Implementation

The gr-ieee802-15-4 module is *not directly* applicable as it only implements the 2.4 GHz OQPSK and CSS PHYs [BLDS13]. Nonetheless, besides other more general-purpose blocks, it provides useful blocks for the MAC. Thus, when appropriate, our implementation reuses existing modules and blocks to implement the IEEE 802.15.4 868 MHz PHY. Ultimately, our implementation consists of three different parts, resulting in two different flowgraphs, one of which comprises the data sender and jammer while the other one represents the attacker. These two flowgraphs are depicted in Figures 5.7 and 5.8, respectively.

Data Sender Implementation. In a nutshell, the data sender flowgraph works as follows. The data sender generates a valid PPDU by using available blocks from gr-ieee802-15-4 and GNU Radio itself. This PPDU is processed as required by the standard [IEE11], i. e., it is differentially encoded, mapped to the corresponding DSSS chips, modulated using BPSK, and filtered. Lastly, the connected SDR is instructed to send the resulting waveform.

More specifically, the data sender uses the gr-ieee802-15-4 MAC block to generate a PSDU. Then, the SHR and PHR are prepended with the appropriate values (as defined in the standard [IEE11]) using the `Access Code Prefixer` and `PDU to Tagged Stream` block, respectively. The resulting PPDU bits are then differentially encoded, mapped to DSSS chips, and ultimately modulated using BPSK. Before forwarding the modulated symbols to the connected SDR, they are filtered according to specification [IEE11, Section 11.2.5.1], i. e., using a (root) raised cosine filter. In our implementation, we use a `Constellation Modulator` block that modulates and, at the same time, filters with an overall oversampling rate of 8. A `Tagged Stream Multiply Length Tag` block ensures that the connected SDR sends out the correct amount of samples, i. e., each single message bit results in $15 \cdot 8 \cdot 8 = 960$ samples due to spreading (factor 15), oversampling (factor 8), and sample rate (which is 8 times higher than the symbol rate).

Note that, for all significant bits to be flushed out of the filter, we pad the PPDU with (irrelevant) zero bits as the filter introduces a delay and would otherwise suppress some PPDU bits. This is achieved by the `Pad Tagged Stream` block. Also, at some stages, additional type conversion or (un-)packing blocks are necessary to ensure the output of one block is a compatible input to a subsequent block.

Jammer Implementation. The jammer is added to the data sender implementation as a second path, enabling a reliable triggering of sending the jamming signal while the data sender is active. Jamming is realized as Chip Bit Jamming (cf. Section 5.5.1).

That is, after the differentially encoded bits are mapped to DSSS chips, we use a custom embedded Python block (denoted `Gen Jamming Bits` in Figure 5.7) for generating jamming bits, effectively realizing random Chip Bit Jamming (cf. Section 5.5.1). Whenever a PPDU is propagated to this block, it discards the incoming PPDU and outputs 960 random chip bits (an equivalent of 64 data bits). These are modulated and filtered exactly like the data

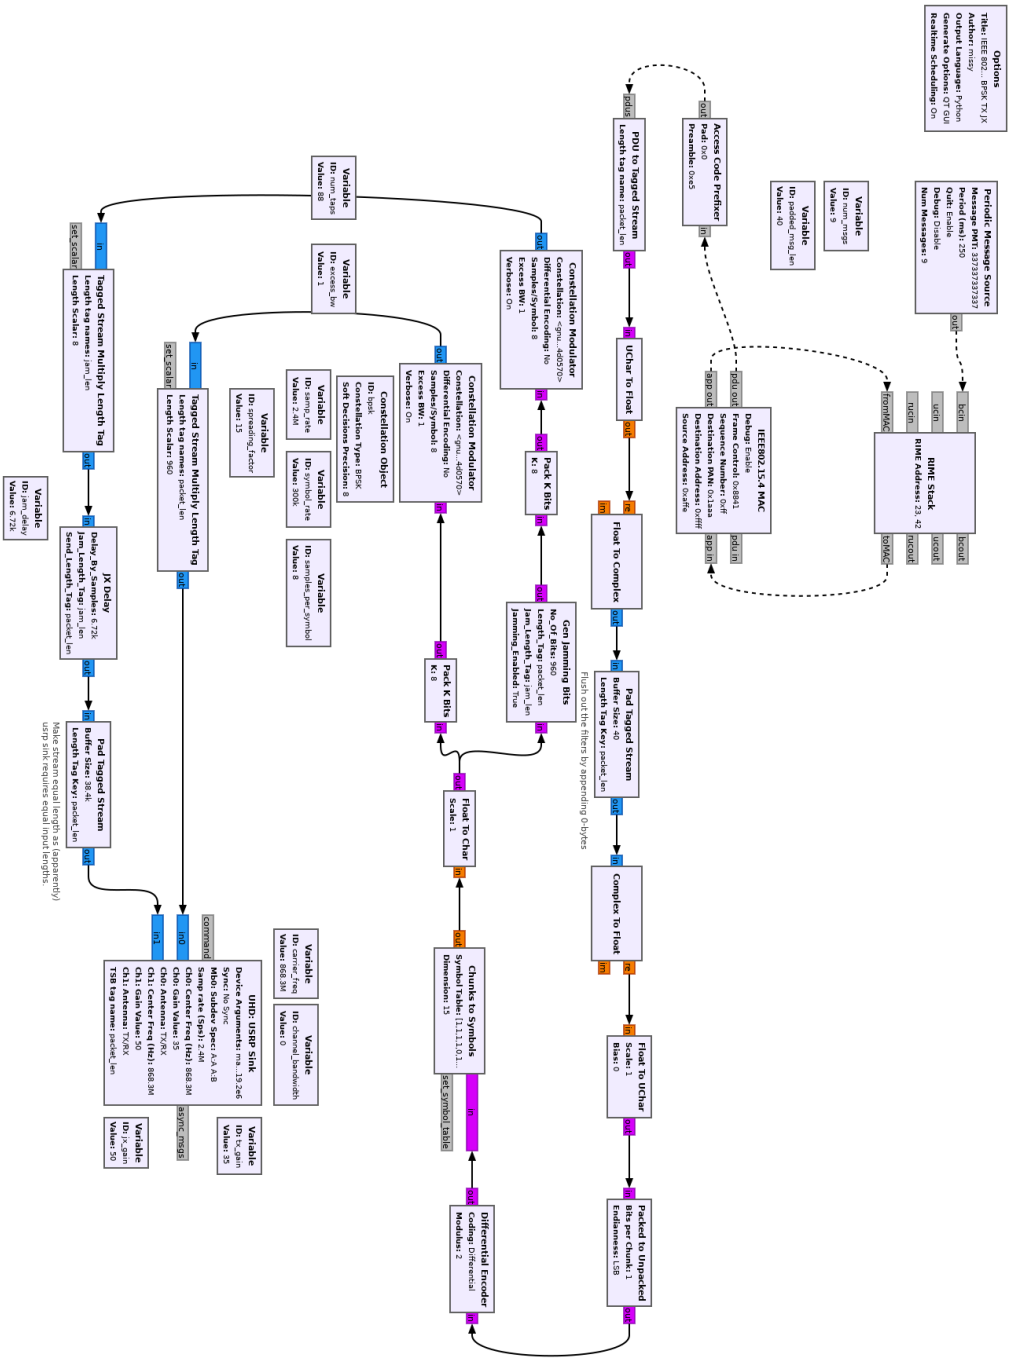

Figure 5.7: The developed GNU Radio flowgraph comprising data sender and jammer.

sender's. Then, this modulated jamming waveform is prepended with zeros such that it is delayed until the data sender has sent SHR and PHR. For technical reasons, both, the data sender's and the jammer's path to the SDR, have to have equal size. Therefore, the jamming waveform is also padded with zeros to match the data sender's signal length. The complete flowgraph comprising data sender and jammer is depicted in Figure 5.7.

Attacker Implementation. Our attacker implementation complements the data sender implementation. However, it includes additional logic required for aligning and decoding incoming packets from the received waveforms. Furthermore, frequency and phase synchronization is carried out in order to synchronize with the data sender. Received PPDUs are processed as mandated by IEEE 802.15.4, hence, the implementation can also serve as a regular receiver. We use a `File Sink` after downsampling, filtering, aligning, and synchronizing, but before demodulation, to store the received jammed waveforms on the local file system.

In more detail, first, the waveform received by the attacker's SDR is passed to an `FLL Band-Edge` block for frequency correction, and is then filtered by a root raised cosine filter (complementary to the data sender side). Subsequently, our implementation uses a `Symbol Sync` block to recover the phase and timing alignment, a `Costas Loop` to synchronize to the BPSK-modulated symbols (after which a `File Sink` is used to store received waveforms), and, for demodulation, a `Constellation Decoder`. After BPSK demodulation, the bit stream is scanned for the beginning of the expected SHR (which, at this point, is differentially encoded and spread) using two `Correlate Access Code` blocks and a custom embedded Python block (denoted `Detect preamble - Align output` in Figure 5.8). Based on the correlator blocks' results, the custom block matches the (now tagged) bit stream against the full differentially encoded and spread SHR and aligns it, such that the sequence of received chips starts at a 15-bit offset. These (DSSS) chips are then decoded (or de-mapped) according to the specification [IEE11, Table 78] by the `Codeword Demapper` block to yield the differentially encoded message bits. After this, a `Differential Decoder` is used and the result is then processed by another custom embedded Python block (dubbed `Detect and write PDU`) as follows. First, the incoming bits are matched against the expected values of the SHR. If there is a positive match, the PHR (cf. Figure 4.1) is extracted and used to determine the number of bits pertaining to the PSDU, i. e., the payload. Once the indicated number of bits is processed, the custom block switches back to scanning continuously for any new SHR, while everything after the end of a received PPDU and before the beginning of another is simply discarded. Since SHR and PHR are not further required, only the resulting PSDU is fed to an IEEE 802.15.4 (or Zigbee) `Wireshark Connector` (provided by gr-foo). The resulting output is written to the local file system as a wireshark-compatible capture file for persistence and possible later inspection.

Ultimately, the attacker stores two separate files on her local file system. The waveform before demodulation and a wireshark-compliant capture file containing received, jammed messages including an FCS respectively a CRC that is mismatching with high probability. To recover jammed bits, the stored waveforms are further processed in a custom-built Matlab script that carries out each of the attack steps as outlined in Section 5.4.3.

5.5.2.3 Recovery Evaluation

Using the jammer, data sender, and attacker implementations, we send 9 packets with a 250 ms interval and a PSDU of 27 bytes each. While the data sender transmits with either 35 dB or

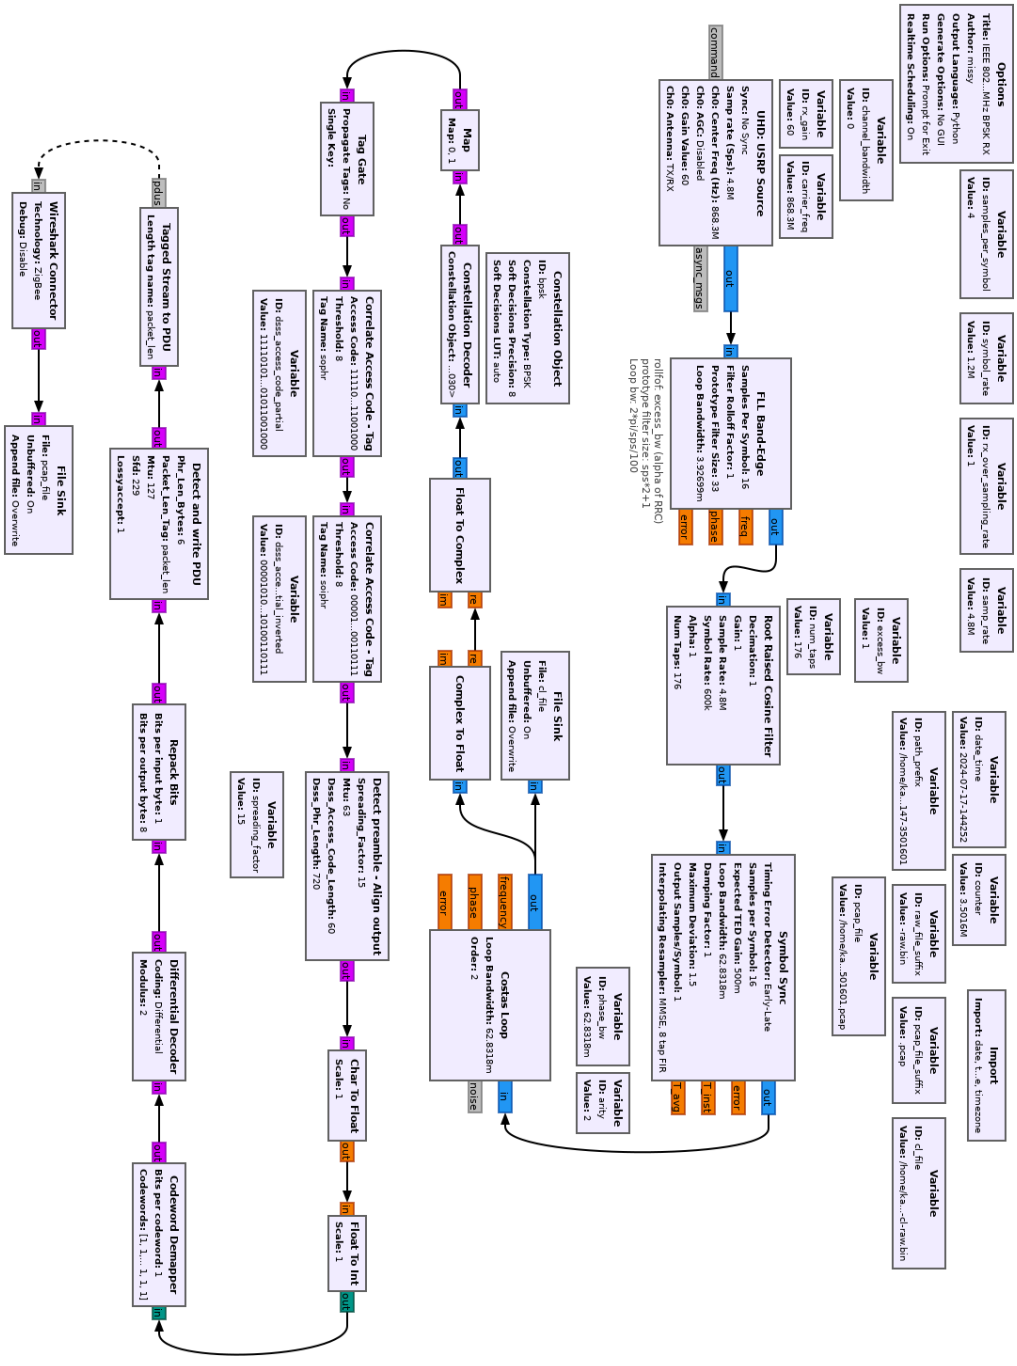

Figure 5.8: The developed GNU Radio flowgraph for the attacker.

40 dB, the jammer uses 15 dB, 20 dB, 25 dB, or 30 dB *more* power than the data sender. This leads to a total of eight different scenarios covering four different SJR settings per attacker antenna.

The transmit power was chosen after preliminary testing and to satisfy two major constraints. Firstly, reliable communication should be possible. Secondly, the data sender's power should be low enough to allow for all chosen SJR values as the maximum sending power is hardware-limited, and to prevent jamming with unnecessarily high power.

5.5.3 Attack Specificities

Here, we elaborate on the different scaling factors used in our attack.

In the simulation-based evaluation, the data sender immediately starts to send the signal. Thus, the waveform does not contain any region with only background noise, and we can omit Attack Step 1 as we do not need to locate the point in time at which the actual transmission starts.

In contrast, in the real-world evaluation, neither do we know when the data sender nor when the jammer starts sending. Thus, first, the beginning and end offsets, $N_{\text{data}}^{\text{begin}}$ and $N_{\text{data}}^{\text{end}}$, of each contained packet need to be found. For this, the general background noise level, A_{noise}, is determined by considering the first $N = 200$ samples' maximum amplitude to compute the data threshold τ_{data}, using scaling factor $scaleNoise = 4$ (cf. Equations (5.6) and (5.7)). Packet beginnings $\left(N_{\text{data}}^{\text{begin}}\right)$ are found by searching for $M = 250$ consecutive samples having a higher amplitude than τ_{data}, packet ends $\left(N_{\text{data}}^{\text{end}}\right)$ are detected by considering $M = 250$ consecutive samples having an amplitude less or equal to τ_{data} (cf. Attack Step 1).

To identify the jammed region as described in Attack Step 2, our attack determines the jamming threshold, τ_{jam}, based on $A_{\text{data}}^{\text{max}}$. For this, the simulation considers the first $M = 1000$ data samples' maximum aplitude while the real-world case uses $M = 250$ (cf. Equation (5.8)). As for τ_{jam}, in the simulation, we use the scaling factor $\alpha_{\text{data}}^{\text{max}} = 1.3$, in the real-world, it is set to $\alpha_{\text{data}}^{\text{max}} = 1.1$ (cf. Equation (5.9)).

In order to estimate the jammer's amplitude, $A_{\text{jam}}^{\text{max}}$ (cf. Equation (5.10)) is determined similarly to $A_{\text{data}}^{\text{max}}$, i. e., by considering the first $M = 1000$ samples of the found jammed region in the simulation case and $M = 250$ samples in the real-world case. Then, A_{jam} is estimated by using the scaling factor $\alpha_{\text{jam}}^{\text{max}} = 0.96$ in the simulation respectively $\alpha_{\text{jam}}^{\text{max}} = 1$ in the real-world case (cf. Equation (5.12)).

All scaling factors are collectively shown in Table 5.2. We identify potential root causes for the scaling factors differing between simulation and real-world in Section 5.7.4.

5.6 Results

This section presents the evaluation results of our attack (cf. Section 5.4) for the simulation and the real-world case (cf. Sections 5.5.1 and 5.5.2) in Sections 5.6.1 and 5.6.2, respectively.

5.6.1 Simulation Attack Results

In Tables 5.3 and 5.4, we present the results of our simulation experiment for Message Bit Jamming and Chip Bit Jamming, respectively.

Table 5.2: Overview of the scaling factors chosen in the respective attack step for simulation and practical evaluation, where "—" means not applicable. N samples are used for noise level estimation, M samples for data threshold and jamming threshold estimation.

Step	Factor	Simulation	Practical
1.	α_{noise}	—	4 ($N = 200$)
2.	$\alpha_{\text{data}}^{\max}$	1.1 ($M = 1000$)	1.3 ($M = 250$)
3.	$\alpha_{\text{jam}}^{\max}$	0.96 ($M = 1000$)	1 ($M = 1000$)

We simulated the attack for 1,000 packets per parameter combination, i. e., for SJR $\in \{n \mid n = -(15 + 5k),\ k = 0, 1, 2, 3\}$, and Chip Shift $\in \{0, 1, 2, 3\}$. To allow for comparison between simulation and practical attack results, we calculated the mean, minimum, and maximum BRR (cf. Equations (5.3) and (5.4)) for each set of parameters.

Attack BRR refers to the attack case in which we apply our recovery action (cf. Step 4 in Section 5.4.2). On the other hand, *Original BRR* pertains to the case without applying our attack, i. e., a standard processing pipeline *without* any corrective actions. For our suggested attack (cf. Section 5.4), most BRRs are very high, and often close to 1.

However, two exceptions stand out. First, the BRR for the Chip Shift case of 2 is the worst in almost all respective settings, probably because two consecutive attacker chip bits affect one data sender bit, leading to ambiguity in our recovery process. Second, the only deviations from this occur in the Message Bit Jamming scenario with SJR $\in \{-25, -30\}$, where a Chip Shift of 3 results in the respective worst BRR (see Table 5.3). Interestingly, in the Chip Bit Jamming scenario, for a Chip Shift of 2, the Original BRR is very high, suggesting that, in these cases, no recovery action is necessary after all (see Table 5.4).

5.6.2 Real-World Attack Results

Table 5.5 provides the practical experiment results. Analyzing the $\overline{\text{BRR}}$ column, it is evident that our method demonstrates robust cancellation of friendly jamming with an SJR of −15 dB and continues to produce encouraging outcomes even at an SJR of −20 dB. For SJR $\in \{-25, -30\}$, our approach results in very low average BRR values of 0.07 and 0.08, respectively.

More precisely, for SJR = −15 dB, the BRRs range from 0.44 to 0.92 ($\overline{\text{BRR}} = 0.70$). Regarding SJR = −20 dB, the range is from 0.19 to 0.76 ($\overline{\text{BRR}} = 0.34$). For SJR = −25 dB, the range is from 0 to 0.16 ($\overline{\text{BRR}} = 0.07$), and for SJR = −30 dB from 0.04 to 0.14 ($\overline{\text{BRR}} = 0.08$).

In Table 5.5, the real-world (Chip Bit Jamming) results are grouped by SJR and antenna type (AT), i. e., O for the omnidirectional and LP for the log-periodic directional antenna (cf. Section 5.5.2). The results show that for SJR = −15 dB the variation of BRR for the omnidirectional antenna is smaller than for the directional antenna. Furthermore, the minimum BRR is higher than the minimum of the directional antenna setup. The described trend of higher BRRs can only be recognized for the maximum BRR in the SJR = −20 dB case. For SJR $\in \{-25, -30\}$, the BRR approaches zero indicating performance similar to random guessing.

Table 5.3: Results of our simulation-based attack experiment for Message Bit Jamming, SJR in dB. 'Attack BRR' refers to a processing pipeline using our attack, 'Original BRR' to a standard processing without any corrective measures. We report the average BRR as $\overline{\text{BRR}}$, as well as minimum and maximum for both, a processing pipeline including our attack and one without it.

SJR	Chip Shift	Attack BRR			Original BRR		
		$\overline{\text{BRR}}$	min	max	$\overline{\text{BRR}}$	min	max
−15	0	0.99	0.94	0.99	0.04	0	0.16
−15	1	0.98	0.94	0.99	0.04	0	0.13
−15	2	0.50	0.02	0.89	0.25	0.09	0.43
−15	3	0.92	0.78	0.99	0.04	0	0.17
−20	0	0.99	0.96	0.99	0.04	0	0.13
−20	1	0.99	0.92	0.99	0.03	0	0.15
−20	2	0.51	0.07	0.89	0.25	0.09	0.40
−20	3	0.91	0.70	0.99	0.03	0	0.14
−25	0	0.99	0.95	0.99	0.04	0	0.13
−25	1	0.83	0.68	0.98	0.04	0	0.15
−25	2	0.80	0.68	0.92	0.25	0.08	0.40
−25	3	0.52	0.32	0.88	0.03	0	0.15
−30	0	0.97	0.66	0.99	0.04	0	0.16
−30	1	0.80	0.68	0.91	0.04	0	0.16
−30	2	0.79	0.65	0.89	0.25	0.06	0.41
−30	3	0.53	0.39	0.70	0.04	0	0.17

Table 5.5 clearly shows an imbalance in the number of measurements (#M)[17] for the used antenna types at a given SJR. To clarify, in our effort to make the data publicly available, we excluded measurements that contained other transmissions. This exclusion was necessary due to data privacy and regulatory concerns. Additionally, to ensure a fully replicable methodology, we filtered out these random, unknown transmissions during preprocessing, as they can neither be reproduced nor published. One might suggest filtering out other transmissions to balance measurements per antenna type and SJR class. However, deciding which transmissions to exclude raises concerns about potentially skewing results.

5.7 Discussion

In this section, we discuss issues and limitations of our attack including its generalizability and practicality (Section 5.7.1). Regarding a possible influence of communication distance on our attack's feasibility, Section 5.7.2 reports on a follow-up experiment where we applied our attack in a setting with increased communication distance. Obviously, the evaluation results (cf. Section 5.6) for the simulation and the real-world differ, e. g., for SJR $= -15\,\text{dB}$, the simulation achieves $\overline{\text{BRR}} \approx 0.8$ (see Table 5.4), the real-world only $\overline{\text{BRR}} = 0.70$ (see

[17] To be precise, this also implies an imbalance in the number of messages (#msg).

Table 5.4: Results of our simulation-based attack experiment for Chip Bit Jamming, SJR in dB. 'Attack BRR' refers to a processing pipeline using our attack, 'Original BRR' to a standard processing without any corrective measures. We report the average BRR as $\overline{\text{BRR}}$, as well as minimum and maximum for both, a processing pipeline including our attack and one without it.

SJR	Chip Shift	Attack BRR			Original BRR		
		$\overline{\text{BRR}}$	min	max	$\overline{\text{BRR}}$	min	max
−15	0	0.93	0.83	0.99	0.04	0	0.16
−15	1	0.85	0.70	0.98	0.03	0	0.14
−15	2	0.58	0.42	0.75	0.93	0.87	1.00
−15	3	0.87	0.71	0.98	0.03	0	0.15
−20	0	0.93	0.86	0.99	0.04	0	0.14
−20	1	0.87	0.70	0.98	0.03	0	0.17
−20	2	0.60	0.41	0.74	0.93	0.85	0.99
−20	3	0.85	0.68	0.98	0.04	0	0.14
−25	0	0.93	0.84	0.99	0.03	0	0.13
−25	1	0.73	0.57	0.88	0.04	0	0.16
−25	2	0.61	0.47	0.75	0.93	0.87	0.99
−25	3	0.72	0.54	0.89	0.04	0	0.16
−30	0	0.83	0.41	0.96	0.04	0	0.14
−30	1	0.68	0.44	0.82	0.04	0	0.14
−30	2	0.54	0.13	0.75	0.93	0.86	0.99
−30	3	0.65	0.46	0.78	0.03	0	0.16

Table 5.5). This trend is rooted in different reasons present in the real-world setting, e. g., chip, phase, and frequency misalignments. We elaborate on these (Section 5.7.3) as well as their influence on threshold computations (Section 5.7.4). Lastly, we highlight possible countermeasures for defending against our attack in Section 5.7.5.

5.7.1 Generalizability and Practicality of Our Attack

We want to stress that all experiments were conducted in an uncontrolled university environment, offering more realistic conditions than a strictly controlled lab setting.

Our attack demonstrates that friendly jamming offers only limited confidentiality by exploiting vulnerabilities at the most fundamental level of data transmission, the PHY. This focused approach allows us to uncover vulnerabilities at this fundamental level of data transmission. The efficacy of our attack is independent of the carrier frequency, as we work with a signal that is already downconverted to baseband containing the modulated information. We therefore conclude that, assuming sufficient available bandwidth, our attack is agnostic to the carrier frequency.[18] As we exploit properties of the modulation scheme, our attack

[18] We tested our attack on the 433 MHz band using the exact same settings and equipment as described in Section 5.7.2 for SJR = −15. While the equipment is not suited for this frequency band, the average BRR is 0.47. The data is similarly available at https://github.com/TheChrisse/OneAntNeutralizeJamming.

Table 5.5: Comparison of mean, minimum, and maximum BRR per antenna type (AT with 'O' denoting the omnidirectional and 'LP' the log-periodic antenna) and SJR level (in dB) in the real-world experiment using Chip Bit Jamming. #M is the number of measurements and #msg the number of messages for each setting. All decimal figures are rounded to the 2nd decimal place.

SJR	AT	#M	#msg	BRR			$\overline{\text{BRR}}$
				mean	min	max	
−15	O	3	27	0.74	0.59	0.86	0.70
−15	LP	4	36	0.68	0.44	0.92	
−20	O	4	36	0.63	0.57	0.76	0.34
−20	LP	2	18	0.25	0.19	0.31	
−25	O	3	27	0.09	0.04	0.16	0.07
−25	LP	3	27	0.05	0.00	0.08	
−30	O	2	18	0.07	0.05	0.08	0.08
−30	LP	4	36	0.09	0.04	0.14	

remains effective across protocols and frequencies that use the same modulation. In principle, our attack might be adapted to other PSK variants. However, more than two possible discrete phase states would require the recovery process to be different. As substantial modifications, parameter optimizations, and new experiments would be required, forming the basis of a separate study, we leave this as an open direction for future work. Furthermore, our method's phase-based design inherently precludes its use with non-PSK modulations. Still, our findings are applicable to a diverse array of higher-layer protocols.

Moreover, our method is not limited to friendly jamming only, it is likewise applicable to *adversarial* jamming. Its fundamental goal is to recover jammed signals, irrespective of the nature of the jammer, i. e., being a friendly or adversarial jammer. Consequently, our approach can also be used as a *defense* against adversarial jamming, ultimately allowing for successful communication under ongoing adversarial jamming, thus defeating the jammer.

Demonstrating a successful attack using simple tools, i. e., a single antenna without machine learning, underscores a stronger potential threat. This approach has dual relevance. First, *practicality*, i. e., a single antenna, easier to conceal and deploy, increases the feasibility of attacks in real-world scenarios. Second, *accessibility*, i. e., avoiding machine learning, which requires dedicated training data, shows that low-resource attackers can still pose significant threats. Additionally, depending on the architecture of such a trained and tuned model, small changes in the setup, e. g., changed distance between data sender and attacker, may decrease its performance and accuracy significantly.

5.7.2 Influence of Communication Distance

As described in Section 5.5.2, the setup is established with short communication distances. However, before the experimental runs with the jammer, we evaluated the power (more precisely, gain) required by the data sender to allow for reliable communication in general, i. e., without active interference by our jammer. This determined gain was then used during

all experimental runs. To further assess the feasibility of our attack, we conducted a follow-up experimental run[19] with the following conditions. We established our test setup in a hallway of the same building in order to increase the communication distance to approximately 10 m (line-of-sight). To account for the larger distance, the data sender's (and likewise the jammer's) gain had to be increased. As the results from the real-world experiment (cf. Sections 5.5.2 and 5.6.2) exhibit a significant drop in BRR from SJR = −20 to SJR = −25, and since the (simpler) omnidirectional antenna generally performs better (cf. Table 5.5), we focus on the SJR ∈ {−15, −20} cases with the omnidirectional antenna.

Table 5.6 presents the results of the 10 m setup. For the evaluation, we reused the *exact same* processing pipeline and settings as in the real-world experiment described before. Interestingly, the BRR is consistently high, suggesting that our attack strategy also works when communication distance is increased. Compared to Table 5.5, the BRRs are even higher. We believe these improvements potentially stem from the fact that the chosen gain is higher than the bare minimum, and, thus, enhances the communication quality in general, allowing for faster, more reliable signal acquisition and synchronization during reception. These results confirm our attack's feasibility and applicability for increased distances.

Table 5.6: Comparison of mean, minimum, and maximum BRR per SJR level (in dB) with communication distance of 10 m and the omnidirectional antenna. #M denotes the number of measurements, #msg the number of messages for the respective setting. All decimal figures are rounded to the 2nd decimal place.

SJR	#M	#msg	BRR		
			mean	min	max
−15	4	36	0.91	0.88	0.92
−20	3	27	0.89	0.89	0.90

5.7.3 Misalignments

Due to possible asynchronicity between jammer and data sender, their signals might not be perfectly aligned, e. g., the signals might exhibit chip, phase, and frequency misalignments.

5.7.3.1 Chip Misalignment

In terms of chip misalignment, the jammer's chips could start mid-way through the transmission of the data sender's chips, resulting in partial overlap and potential misalignment. We incorporate this in our simulation, and allow it to be programmatically set to simulate different synchronization scenarios.

The pulse shaping process is designed with an oversampling rate of 4, meaning that each chip bit is represented by four samples. Depending on the relative offset (denoted Chip Shift, cf. Section 5.5.1) between data sender and jammer, four received samples might represent two complete chip bits (one of each, data sender's and jammer's) or one complete (e. g., the data sender's) and parts of two adjacent chip bits (e. g., the jammer's). Note that any Chip Shift $c' \in \mathbb{N}, c' > 3$ can be represented by $c = c' \bmod 4$.

[19] Data also available at: https://github.com/TheChrisse/OneAntNeutralizeJamming.

As described in Tables 5.3 and 5.4 and Section 5.6.1, when the offset does not align perfectly, some degree of recovery is still possible. If the offset is 2, recovery is not feasible due to the overlap dividing two chips equally, making it impossible to discern the boundary between two consecutive bits. However, evidently in some settings no recovery procedure is necessary since standard processing achieves high BRRs. In all other cases, the attacker can potentially recover the bits by using our attack.

One can observe that the Chip Bit Jamming (and Message Bit Jamming) results displayed in Table 5.4 (and Table 5.3) from the simulation differ from those in Table 5.5 obtained in the practical experiments. In contrast to the simulation, in which the Chip Shift is perfectly consistent for each transmission, in the real-world experiment, the Chip Shift is not directly under our control, i. e., it may fluctuate between jammed transmissions, and it may even vary *within* a single transmission. These effects may stem from fluctuations in the physical environment, and that the receiving hardware is constantly trying to improve the synchronization with the currently incoming signal.

5.7.3.2 Phase and Frequency Misalignment

Due to production inaccuracies and the exact distance between jammer, data sender, and attacker, received signals might be slightly misaligned, e. g., the attacker observes an additional relative phase shift between the two signals or slightly different frequencies of all parties involved. Whereas the former, e. g., caused by clock drift or jitter, has a relatively low impact, the actual distance may induce a significant phase shift, introducing uncertainties that further reduce the attacker's estimation accuracy regarding the recovery of samples' phase values (cf. Property 4 and Equations (5.13) and (5.14)). As Figure 5.3 evidently shows, the signal is distributed over both, real *and* imaginary axis[20] (in Figure 5.3, blue color represents real and orange the imaginary part of the complex-valued signal).

5.7.4 Differing Thresholds

A major difference between the simulation and the real-world is that the observed signal levels are completely different, e. g., in the practical evaluation the peaks are approximately between -1.3 and $+1.3$, whereas the same settings in our simulation result in peaks approximately between -40 and $+40$ (cf. Figures 5.1 and 5.5). We attribute this difference to the different processing pipelines as we expect that, e. g., the real-world waveform is (nonlinearly) normalized by the hardware and by the used GNU Radio blocks as well. Therefore, the calculations for τ_{jam} (cf. Equation (5.9)) use a different scaling factor for the simulation and practical evaluation (cf. Table 5.2).

5.7.5 Countermeasures

As our attack requires reliable estimations of both, the data sender's and the jammer's amplitude, varying these during transmission represents a promising countermeasure. However, the jammer needs to choose amplitude variations that are unpredictable for the attacker as, otherwise, she could adapt to and compensate for the varying values and still carry out the

[20] Received signals are generally complex-valued. However, as data sender and attacker are not perfectly synchronized, the received signal exhibits a phase offset, thus, parts of it are spread over both, real and imaginary part, although for BPSK, the signal's phase should be 0 or π, i. e., the received signal should be real-valued (cf. Section 5.2.2.1 and Equation (4.1)).

attack successfully. That is, the jammer has to use some *unpredictably* (or pseudorandomly) varying amplitudes to protect against our attack. At the same time, the jammer needs to respect the limited range of acceptable SJRs for it to be effective.

Additionally, inhibiting any unjammed transmissions from a friendly jamming-protected device would prevent the attacker from estimating the data sender's amplitude. This would require the data sender be aware of the jammer and a synchronization mechanism, which could be challenging by itself. However, depending on the actual setup, it might happen that, due to constructive and destructive interference, the data sender's amplitude could still be estimated by comparing the highest with the lowest peaks (cf. Figures 5.1 and 5.5). Nonetheless, in practice, this might be challenging for the attacker.

5.8 Conclusion

In this chapter, we presented a novel one-antenna attack on jamming to successfully recover jammed bits in both, a simulation and a practical evaluation. Following established research, we used the commonly accepted metric BRR to evaluate our recovery method under four different SJR values. The simulation-based evaluation is implemented in Matlab using standard methods. Here, the attacker is able to recover a significant fraction of bits leading to average BRRs from 0.54 to 0.93. For the practical evaluation, we implemented GNU Radio flowgraphs that allow for OTA transmissions being jammed *during transmission*. Our results show that the attacker is able to recover a significant number of bits with $\overline{\text{BRR}}$ of 0.70 and 0.34 in the SJR $= -15\,\text{dB}$ and SJR $= -20\,\text{dB}$ settings, respectively. The performance in the SJR $= -25\,\text{dB}$ and SJR $= -30\,\text{dB}$ cases with $\overline{\text{BRR}}$ of 0.07 and 0.08, respectively, still leaves room for improvement. In contrast to this, our follow-up experiment over a distance of 10 m shows that our attack is applicable in such a setting with $\overline{\text{BRR}}$ of 0.91 and 0.89 in the SJR $= -15\,\text{dB}$ and SJR $= -20\,\text{dB}$ settings, respectively. However, our findings clearly narrow down the space of possible SJR values for friendly and adversarial jamming. All developed code and data are made available in the hope of stimulating further research.

Based on our work and results presented in this chapter, several future research directions emerge, e. g., expanding our study by developing single-antenna attacks for further modulation schemes. Another direction would be to create a more realistic experimental environment with additional noise and complexity, e. g., simulating a typical home setting with multiple legitimate users and wireless standards. This would allow for a better understanding of current attack models' jamming cancellation capabilities. Lastly, we propose developing open datasets specifically for jamming and anti-jamming research to facilitate standardized benchmarking and to promote collaborative advancements within the community.

Ah, I always rip out the last page of a book.
That way it doesn't have to end. I hate endings.

Eleventh Doctor (Doctor Who)

6 Conclusion

On the one hand, the IoT has established itself as an integral part of everyday life and offers new applications and convenience features like smart lighting or autonomous heating control, all combined in smart homes. On the other hand, it is not completely obvious what kind of potential privacy risks arise from collecting seemingly innocuous information such as room climate data. However, providing convincing incentives for why this kind of data should be protected in the first place is not straightforward. As long as it can be ensured that such collected information is protected throughout its lifecycle, the mere collection might not necessarily be problematic at all. How exactly this protection should be implemented is a different question.

Summarizing, pertinent to the aforementioned aspects of the IoT, the following three main questions came to mind:

1. Is it necessary to protect basic data such as temperature and (relative) humidity emanating from a smart home?

2. If such data require protection, are there any problems when using established communication protocols? And if so, what are the root causes of these problems?

3. To what extent can insights into the root causes of these fundamental issues be exploited to derive novel attack vectors or to advance analytical frameworks?

In this thesis, we addressed each one of these identified questions. First, in Chapter 2, we studied the potential privacy risks of room climate measurements, e. g., to detect the presence of persons using the two physical quantities temperature and (relative) humidity. More alarmingly, we showed that the four defined activities, Read, Stand, Walk, and Work, can reliably be distinguished from each other, raising the question of what other activities in what circumstances could be identified. These results provide a strong point for the necessity of protecting basic smart home data like room climate data from unwarranted access throughout their entire lifecycle.

In Chapter 3, we then analyzed the popular wireless communication protocol Zigbee. More precisely, we discovered several vulnerabilities in Zigbee's touchlink commissioning procedure, allowing exploitation by an attacker, e. g., to take over devices or to eavesdrop on and decrypt encrypted transmissions. As some of these vulnerabilities are rooted in the specification itself, they cannot easily be fixed without breaking compatibility. Already a single touchlink-enabled device puts the complete network at risk, therefore, support for the touchlink commissioning procedure should be completely removed. Furthermore, we believe that some of the discovered weaknesses stem from wrong assumptions about wireless communication which, however, might be correct for wired communication.

Consequently, in Chapter 4, we introduced a new communication and attacker model tailored to the peculiarities of wireless communication, the WCM and WAM. There, we identified the three key concepts — communication channels, signals, and locality — that are necessary to represent the unique properties of wireless communication. We showed our model's versatility by modeling FHSS, DB protocols, friendly jamming, and a two-antenna attacker (according to Tippenhauer et al. [TMRC13]) able to recover friendly-jammed signals, thereby observing that the effect of jamming depends on the chosen modulation among others. As our design is highly influenced by real-world aspects, we hope that our model stimulates the development of new solutions that are secure in a wireless setting, regardless of whether these would be considered secure in a wired setting.

Motivated by the two-antenna attacker, in Chapter 5, we developed a novel attack method targeting (friendly) jammed IEEE 802.15.4 communication. More precisely, specifically targeting the 868 MHz BPSK variant, our attack reliably neutralizes the effects of (friendly) jamming while using a low-tech approach, i. e., the attacker has only *one* receiving antenna without any machine learning. We examined our attack in both, a theoretical and a real-world setting, and recover approximately 70% of jammed data bits on average under 15 dB stronger jamming. Under very strong jamming, i. e., the jamming signal is 25 dB or 30 dB stronger than the data signal, our recovery attack performs close to random guessing. Additionally, we performed a proof-of-concept with an indoor communication distance of 10 m to estimate the influence of the increased distance on our recovery approach as, in our first experimental setting, the communication distance (approximately 0.5 m) and the chosen gain values were minimal. Our results underline the attack's real-world feasibility as we can recover almost 90% of jammed data bits on average under more challenging circumstances. Ultimately, these results demonstrate the threat potential of low-resource one-antenna attackers against friendly-jamming techniques.

Bibliography

[3GP16] 3GPP. *3GPP TS 36.300: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall Description; Stage 2 (Release 13)*. Dec. 2016. URL: https://www.3gpp.org/ftp/Specs/archive/36_series/36.300/36300-d60.zip.

[3GP22a] 3GPP. *3GPP TS 36.300: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2 (Release 17)*. Dec. 2022. URL: https://www.3gpp.org/ftp/Specs/archive/36_series/36.300/36300-h30.zip.

[3GP22b] 3GPP. *3GPP TS 38.300: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; NR; NR and NG-RAN Overall Description; Stage 2 (Release 17)*. Dec. 2022. URL: https://www.3gpp.org/ftp/Specs/archive/38_series/38.300/38300-h30.zip.

[AB14] SaiDhiraj Amuru and R. Michael Buehrer. "Optimal jamming strategies in digital communications — Impact of modulation." In: *2014 IEEE Global Communications Conference*. 2014, pp. 1619–1624. DOI: 10.1109/GLOCOM.2014.7037040.

[ABG+21] Gildas Avoine, Ioana Boureanu, David Gérault, Gerhard P. Hancke, Pascal Lafourcade, and Cristina Onete. "From Relay Attacks to Distance-Bounding Protocols." In: *Security of Ubiquitous Computing Systems: Selected Topics*. Ed. by Gildas Avoine and Julio Hernandez-Castro. Cham: Springer, 2021, pp. 113–130. ISBN: 978-3-030-10591-4. DOI: 10.1007/978-3-030-10591-4_7.

[ABM+19] Frederik Armknecht, Zinaida Benenson, Philipp Morgner, Christian Müller, and Christian Riess. "Privacy implications of room climate data." In: *Journal of Computer Security* 27.1 (2019), pp. 113–136. ISSN: 1875-8924. DOI: 10.3233/JCS-181133.

[AFG14] Bing Ai, Zhaoyan Fan, and Robert X. Gao. "Occupancy estimation for smart buildings by an auto-regressive hidden Markov model." In: *American Control Conference, ACC 2014, Portland, OR, USA, June 4-6, 2014*. IEEE, 2014, pp. 2234–2239. ISBN: 978-1-4799-3272-6. DOI: 10.1109/ACC.2014.6859372.

[AH13] Adnan Abu-Mahfouz and Gerhard P. Hancke. "Distance Bounding: A Practical Security Solution for Real-Time Location Systems." In: *IEEE Transactions on Industrial Informatics* 9.1 (2013), pp. 16–27. DOI: 10.1109/TII.2012.2218252.

[ALAM19] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. "SoK: Security Evaluation of Home-Based IoT Deployments." In: *2019 IEEE Symposium on Security and Privacy*. San Francisco, CA, USA: IEEE Computer Society Press, 2019, pp. 1362–1380. DOI: 10.1109/SP.2019.00013.

[AM24] Frederik Armknecht and Christian Müller. "Towards a Cryptographic Model for Wireless Communication." In: *Proceedings of the 21st International Conference on Security and Cryptography — SECRYPT*. Ed. by Sabrina De Capitani Di Vimercati and Pierangela Samarati. INSTICC. SciTePress, 2024, pp. 249–261. ISBN: 978-989-758-709-2. DOI: 10.5220/0012809300003767.

[AM25] Frederik Armknecht and Christian Müller. "Towards a Cryptographic Model for Wireless Communication: Extended Version." In: *Security and Cryptography. 20th International Conference, SECRYPT 2023, Rome, Italy, July 10–12, 2023, and 21st International Conference, SECRYPT 2024, Dijon, France, July 8–10, 2024, Revised Selected Papers*. Ed. by Pierangela Samarati and Sabrina De Capitani di Vimercati. Vol. 2588. Communications in Computer and Information Science. Cham: Springer Nature Switzerland, Nov. 16, 2025, pp. 215–242. ISBN: 978-3-032-09598-5. DOI: 10.1007/978-3-032-09598-5_10.

[AMSP17] Aysajan Abidin, Eduard Marin, Dave Singelée, and Bart Preneel. "Towards Quantum Distance Bounding Protocols." In: *Radio Frequency Identification and IoT Security: 12th International Workshop, RFIDSec 2016, Hong Kong, China, November 30 – December 2, 2016, Revised Selected Papers*. Ed. by Gerhard P. Hancke and Konstantinos Markantonakis. Vol. 10155. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2017, pp. 151–162. ISBN: 978-3-319-62024-4. DOI: 10.1007/978-3-319-62024-4_11.

[Ash09] Kevin Ashton. *That 'Internet of Things' Thing*. RFID Journal, Emerald X, LLC, 2009. URL: https://web.archive.org/web/20231013205225/https://www.rfidjournal.com/that-internet-of-things-thing.

[BC94] Stefan Brands and David Chaum. "Distance-Bounding Protocols." In: *Advances in Cryptology — EUROCRYPT '93: Workshop on the Theory and Application of Cryptographic Techniques Lofthus, Norway, May 23–27, 1993 Proceedings*. Ed. by Tor Helleseth. Berlin, Heidelberg: Springer Berlin Heidelberg, 1994, pp. 344–359. ISBN: 978-3-540-48285-7. DOI: 10.1007/3-540-48285-7_30.

[BCSS09] David Basin, Srdjan Capkun, Patrick Schaller, and Benedikt Schmidt. "Let's Get Physical: Models and Methods for Real-World Security Protocols." In: *Theorem Proving in Higher Order Logics*. Ed. by Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1–22. ISBN: 978-3-642-03359-9. DOI: 10.1007/978-3-642-03359-9_1.

[BCSS11] David Basin, Srdjan Capkun, Patrick Schaller, and Benedikt Schmidt. "Formal Reasoning about Physical Properties of Security Protocols." In: *ACM Transactions on Information and System Security* 14.2 (Sept. 2011). ISSN: 1094–9224. DOI: 10.1145/2019599.2019601.

[BDD+21] Ioana Boureanu, Constantin Cătălin Drăgan, François Dupressoir, David Gérault, and Pascal Lafourcade. "Mechanised Models and Proofs for Distance-Bounding." In: *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*. June 2021, pp. 1–16. DOI: 10.1109/CSF51468.2021.00049.

[BGF+14] Daniel S. Berger, Francesco Gringoli, Nicolò Facchi, Ivan Martinovic, and Jens Schmitt. "Gaining Insight on Friendly Jamming in a Real-World IEEE 802.11 Network." In: *Proceedings of the 2014 ACM Conference on Security and Privacy in Wireless & Mobile Networks*. WiSec '14. Oxford, United Kingdom: Association for Computing Machinery, 2014, pp. 105–116. ISBN: 978-1-4503-2972-9. DOI: 10.1145/2627393.2627403.

[BGF+16] Daniel S. Berger, Francesco Gringoli, Nicolò Facchi, Ivan Martinovic, and Jens B. Schmitt. "Friendly Jamming on Access Points: Analysis and Real-World Measurements." In: *IEEE Transactions on Wireless Communications* 15.9 (2016), pp. 6189–6202. DOI: 10.1109/TWC.2016.2581165.

[BLDS13] Bastian Bloessl, Christoph Leitner, Falko Dressler, and Christoph Sommer. "A GNU Radio-based IEEE 802.15.4 Testbed." In: *12. GI/ITG KuVS Fachgespräch Drahtlose Sensornetze (FGSN 2013)*. Cottbus, Germany, Sept. 2013, pp. 37–40. URL: https://nbn-resolving.org/urn:nbn:de:kobv:co1-opus-28986.

[Blu16] Bluetooth SIG. *Bluetooth Core Specification 5.0*. Dec. 6, 2016. URL: https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=587173.

[Blu24] Bluetooth SIG. *Bluetooth Core Specification 6.0*. Aug. 27, 2024. URL: https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=588232.

[BMV15a] Ioana Boureanu, Aikaterini Mitrokotsa, and Serge Vaudenay. "Practical and Provably Secure Distance-Bounding." In: *Information Security*. Ed. by Yvo Desmedt. Cham: Springer International Publishing, 2015, pp. 248–258. ISBN: 978-3-319-27659-5. DOI: 10.1007/978-3-319-27659-5_18.

[BMV15b] Ioana Boureanu, Aikaterini Mitrokotsa, and Serge Vaudenay. "Practical and provably secure distance-bounding." In: *Journal of Computer Security* 23.2 (2015), pp. 229–257. ISSN: 1875-8924. DOI: 10.3233/JCS-140518.

[BSI14] BSI. *Protection Profile for the Gateway of a Smart Metering System (Smart Meter Gateway PP)*. Mar. 2014. URL: https://www.commoncriteriaportal.org/files/ppfiles/pp0073b_pdf.pdf.

[CAF19] Gordon Chu, Noah Apthorpe, and Nick Feamster. "Security and Privacy Analyses of Internet of Things Children's Toys." In: *IEEE Internet of Things Journal* 6.1 (2019), pp. 978–985. DOI: 10.1109/JIOT.2018.2866423.

[CCC16] CCC. *Chaos Computer Club: Guidelines for Smart Home Solutions*. Feb. 2016. URL: https://www.ccc.de/en/updates/2016/smarthome.

[CF16] Luis M. Candanedo and Véronique Feldheim. "Accurate occupancy detection of an office room from light, temperature, humidity and CO2 measurements using statistical learning models." In: *Energy and Buildings* 112 (2016), pp. 28–39. ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2015.11.071.

[CFG+16] Franco Cicirelli, Giancarlo Fortino, Andrea Giordano, Antonio Guerrieri, Giandomenico Spezzano, and Andrea Vinci. "On the Design of Smart Homes: A Framework for Activity Recognition in Home Environment." In: *Journal of Medical Systems* 40.9 (July 2016), p. 200. ISSN: 1573-689X. DOI: 10.1007/s10916-016-0549-7.

[Cha14] Alex Chapman. "Hacking into Internet Connected Light Bulbs." July 2014. URL: https://web.archive.org/web/20141226123510/https://www.contextis.com/resources/blog/hacking-internet-connected-light-bulbs/.

[Cho18] Niraj Chokshi. *Amazon Alexa Recorded and Shared a Conversation*. May 2018. URL: https://www.nytimes.com/2018/05/25/business/amazon-alexa-conversation-shared-echo.html.

[CK78] Imre Csiszár and Janos Korner. "Broadcast channels with confidential messages." In: *IEEE Transactions on Information Theory* 24.3 (1978), pp. 339–348. DOI: 10.1109/TIT.1978.1055892.

[CPW10] Ann Cavoukian, Jules Polonetsky, and Christopher Wolf. "SmartPrivacy for the Smart Grid: embedding privacy into the design of electricity conservation." In: *Identity in the Information Society* 3.2 (2010), pp. 275–294. ISSN: 1876-0678. DOI: 10.1007/s12394-010-0046-y.

[DAL+10] Bing Dong, Burton Andrews, Khee Poh Lam, Michael Höynck, Rui Zhang, Yun-Shang Chiou, and Diego Benitez. "An information technology enabled sustainability test-bed (ITEST) for occupancy detection through an environmental sensing network." In: *Energy and Buildings* 42.7 (2010), pp. 1038–1046. ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2010.01.016.

[Del15] Deloitte. *Ready for Takeoff?* Consumer Survey (in German). July 2015. URL: https://web.archive.org/web/20151021100322/https://www.deloitte.com/de/de/pages/technology-media-and-telecommunications/articles/smart-home-consumer-survey.html.

[Del23] Delock. *Data Sheet Delock LoRa 868 MHz Antenna SMA plug 3 dBi omnidirectional*. Delock. 2023. URL: https://www.delock.com/produkt/89769/pdf.html?sprache=en.

[DFKO11] Ulrich Dürholz, Marc Fischlin, Michael Kasper, and Cristina Onete. "A Formal Approach to Distance-Bounding RFID Protocols." In: *Information Security*. Ed. by Xuejia Lai, Jianying Zhou, and Hui Li. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 47–62. ISBN: 978-3-642-24861-0. DOI: 10.1007/978-3-642-24861-0_4.

[DGB88] Yvo Desmedt, Claude Goutier, and Samy Bengio. "Special Uses and Abuses of the Fiat-Shamir Passport Protocol (extended abstract)." In: *Advances in Cryptology — CRYPTO '87: Proceedings*. Ed. by Carl Pomerance. Berlin, Heidelberg: Springer Berlin Heidelberg, 1988, pp. 21–39. ISBN: 978-3-540-48184-3. DOI: 10.1007/3-540-48184-2_3.

[DGV04] Adam Dunkels, Björn Grönvall, and Thiemo Voigt. "Contiki – a lightweight and flexible operating system for tiny networked sensors." In: *29th Annual IEEE International Conference on Local Computer Networks, 2004*. IEEE. 2004, pp. 455–462. DOI: 10.1109/LCN.2004.38.

[Dha13] Nitesh Dhanjani. *Hacking Lightbulbs: Security Evaluation of the Philips Hue Personal Wireless Lighting System*. Aug. 2013. URL: https://web.archive.org/web/20140213012545/https://www.dhanjani.com/blog/2013/08/hacking-lightbulbs.html.

[DM07] Saar Drimer and Steven J. Murdoch. "Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks." In: *16th USENIX Security Symposium (USENIX Security 07)*. Boston, MA: USENIX Association, Aug. 2007. URL: https://www.usenix.org/conference/16th-usenix-security-symposium/keep-your-enemies-close-distance-bounding-against.

[DY83] Danny Dolev and Andrew Chi-Chih Yao. "On the security of public key protocols." In: *IEEE Trans. Information Theory* 29.2 (1983), pp. 198–207. DOI: 10.1109/TIT.1983.1056650.

[EBPF13] Tobore Ekwevugbe, Neil Brown, Vijay Pakka, and Denis Fan. "Real-time building occupancy sensing using neural-network based sensor network." In: *7th IEEE International Conference on Digital Ecosystems and Technologies (DEST), 2013*. July 2013, pp. 114–119. DOI: 10.1109/DEST.2013.6611339.

[EBV+13] Afrooz Ebadat, Giulio Bottegal, Damiano Varagnolo, Bo Wahlberg, and Karl Henrik Johansson. "Estimation of building occupancy levels through environmental signals deconvolution." In: *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. BuildSys '13. Roma, Italy: Association for Computing Machinery, 2013, pp. 1–8. ISBN: 978-1-4503-2431-1. DOI: 10.1145/2528282.2528290.

[EBV+15] Afrooz Ebadat, Giulio Bottegal, Damiano Varagnolo, Bo Wahlberg, and Karl Henrik Johansson. "Regularized Deconvolution-Based Approaches for Estimating Room Occupancies." In: *IEEE Transactions on Automation Science and Engineering* 12.4 (2015), pp. 1157–1168. DOI: 10.1109/TASE.2015.2471305.

[Eco15] Ecobee. *Privacy Policy & Terms of Use*. Apr. 2015. URL: https://web.archive.org/web/20150324182340/https://www.ecobee.com/legal/use/.

[Elb08] Bruce R. Elbert. *Introduction to Satellite Communication*. 3rd ed. The Artech House Telecommunications Library. Boston: Artech House, Inc, 2008. ISBN: 978-1-59693-210-4.

[Eri23] Ericsson. *Number of cellular IoT connections worldwide from 2022 to 2028, by connection technology (in millions)*. Statista, Feb. 2023. URL: https://www.statista.com/statistics/1388055/cellular-iot-connections-worldwide/.

[ESWM12] Simon Eberz, Martin Strohmeier, Matthias Wilhelm, and Ivan Martinovic. "A Practical Man-In-The-Middle Attack on Signal-Based Key Generation Protocols." In: *Computer Security – ESORICS 2012: 17th European Symposium on Research in Computer Security, Pisa, Italy, September 10-12, 2012. Proceedings*. Ed. by Sara Foresti, Moti Yung, and Fabio Martinelli. Berlin, Heidelberg:

Springer Berlin Heidelberg, 2012, pp. 235–252. ISBN: 978-3-642-33167-1. DOI: 10.1007/978-3-642-33167-1_14.

[Ett12] Ettus Research. *LP0965 Data Sheet*. Ettus Research. 2012. URL: https://kb.ettus.com/images/0/06/ettus_research_lp0965_datasheet.pdf.

[Ett18a] Ettus Research. *USRP B200-B210 – Spec Sheet*. Ettus Research. 2018. URL: https://www.ettus.com/wp-content/uploads/2019/01/b200-b210_spec_sheet.pdf.

[Ett18b] Ettus Research. *USRP B200mini Series – Data Sheet*. Ettus Research. 2018. URL: https://www.ettus.com/wp-content/uploads/2019/01/USRP_B200mini_Data_Sheet.pdf.

[Eur15] European Union Agency For Network And Information Security. *Security and Resilience of Smart Home Environments – Good practices and recommendations*. Dec. 2015. URL: https://www.enisa.europa.eu/publications/security-resilience-good-practices.

[Eur16] European Parliament and the Council of the European Union. "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)." In: *Official Journal of the European Union* (May 2016), pp. 1–88. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L:2016:119:TOC.

[FXL+17] Xiaohu Fan, Qubo Xie, Xuebin Li, Hoa Huang, Jian Wang, Si Chen, Changsheng Xie, and Jiajing Chen. "Activity Recognition as a Service for Smart Home: Ambient Assisted Living Application via Sensing Home." In: *2017 IEEE International Conference on AI Mobile Services (AIMS)*. June 2017, pp. 54–61. DOI: 10.1109/AIMS.2017.29.

[GBM+12] Travis Goodspeed, Sergey Bratus, Ricky Melgares, Ryan Speers, and Sean W. Smith. "Api-do: Tools for Exploring the Wireless Attack Surface in Smart Meters." In: *45th Hawaii International Conference on Systems Science (HICSS-45 2012), Proceedings, 4-7 January 2012, Grand Wailea, Maui, HI, USA*. IEEE Computer Society, 2012, pp. 2133–2140. ISBN: 978-0-7695-4525-7. DOI: 10.1109/HICSS.2012.115.

[GGJL12] Ulrich Greveler, Peter Glösekötterz, Benjamin Justusy, and Dennis Loehr. "Multimedia content identification through smart meter power usage profiles." In: *Proceedings of the International Conference on Information and Knowledge Engineering (IKE)*. Jan. 2012. URL: https://www.researchgate.net/publication/266461208_Multimedia_Content_Identification_Through_Smart_Meter_Power_Usage_Profiles.

[GHR+11] Shyamnath Gollakota, Haitham Hassanieh, Benjamin Ransford, Dina Katabi, and Kevin Fu. "They Can Hear Your Heartbeats: Non-invasive Security for Implantable Medical Devices." In: *Proceedings of the ACM SIGCOMM 2011 Conference*. SIGCOMM '11. Toronto, Ontario, Canada: Association for Computing Machinery, 2011, pp. 2–13. ISBN: 978-1-4503-0797-0. DOI: 10.1145/2018436.2018438.

[GLO15] Sébastien Gambs, Carlos Eduardo Rosar Kós Lassance, and Cristina Onete. "The Not-so-Distant Future: Distance-Bounding Protocols on Smartphones." In: *Smart Card Research and Advanced Applications - 14th International Conference, CARDIS 2015, Bochum, Germany, November 4-6, 2015. Revised Selected Papers*. Ed. by Naofumi Homma and Marcel Medwed. Vol. 9514. Lecture Notes in Computer Science. Springer, 2015, pp. 209–224. DOI: 10.1007/978-3-319-31271-2_13.

[GLY14] Kanika Grover, Alvin Lim, and Qing Yang. "Jamming and Anti-Jamming Techniques in Wireless Networks: A Survey." In: *International Journal of Ad Hoc and Ubiquitous Computing (IJAHUC)* 17.4 (Dec. 2014), pp. 197–215. ISSN: 1743-8225. DOI: 10.1504/IJAHUC.2014.066419.

[Gol05] Andrea Goldsmith. *Wireless Communications*. 1st ed. Cambridge University Press, Aug. 2005. ISBN: 978-0-511-84122-4. DOI: 10.1017/CBO9780511841224.

[Goo26] Google. *HTTPS encryption on the web – Google Transparency Report*. Mar. 8, 2026. URL: https://transparencyreport.google.com/https/overview.

[GRR+17] Shabnam Ghaffarzadegan, Attila Reiss, Mirko Ruhs, Robert Duerichen, and Zhe Feng. "Occupancy Detection in Commercial and Residential Environments Using Audio Signal." In: *Interspeech 2017. 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*. Ed. by Francisco Lacerda. ISCA, 2017, pp. 3802–3806. DOI: 10.21437/Interspeech.2017-524.

[GSM23] GSMA. *Worldwide Internet of Things (IoT) connections from 2020 to 2030, by application (in billions)*. Statista, Mar. 2023. URL: https://www.statista.com/statistics/1403256/global-iot-connections/.

[Hal15] John E. Hall. *Guyton and Hall Textbook of Medical Physiology*. 13th edition. Philadelphia: Elsevier, 2015. ISBN: 978-1-4557-7005-2.

[Har89] George W. Hart. "Residential energy monitoring and computerized surveillance via utility power flows." In: *IEEE Technology and Society Magazine* 8.2 (June 1989), pp. 12–16. DOI: 10.1109/44.31557.

[Has22] Mohammad Hasan. *State of IoT 2022: Number of connected IoT devices growing 18% to 14.4 billion globally*. May 2022. URL: https://web.archive.org/web/20220602160047/https://iot-analytics.com/number-connected-iot-devices/.

[HC21] Yao Han and Meng Chen. "A Jamming Method for Costas Loop Carrier Synchronization Based BPSK Signals." In: *2021 4th International Conference on Information Communication and Signal Processing (ICICSP)*. 2021, pp. 11–15. DOI: 10.1109/ICICSP54369.2021.9611843.

[Hei16] Deral Heiland. "R7-2016-10: Multiple Osram Sylvania Osram Lightify Vulnerabilities (CVE-2016-5051 through 5059)." July 2016. URL: https://www.rapid7.com/blog/post/2016/07/26/r7-2016-10-multiple-osram-sylvania-osram-lightify-vulnerabilities-cve-2016-5051-through-5059/.

[Her05] Jonathan Herzog. "A computational interpretation of Dolev-Yao adversaries." In: *Theoretical Computer Science* 340.1 (2005). Theoretical Foundations of Security Analysis and Design II, pp. 57–81. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2005.03.003.

[HFH+09] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. "The WEKA Data Mining Software: An Update." In: *SIGKDD Explor. Newsl.* 11.1 (Nov. 2009), pp. 10–18. ISSN: 1931-0145. DOI: 10.1145/1656274.1656278.

[HGAK11] Ebenezer Hailemariam, Rhys Goldstein, Ramtin Attar, and Azam Khan. "Real-time occupancy detection using decision trees with multiple sensor types." In: *2011 Spring Simulation Multi-conference, SpringSim '11, Boston, MA, USA, April 03-07, 2011.* 2011, pp. 141–148. URL: https://dl.acm.org/doi/10.5555/2048536.2048555.

[HGF12] Zhenyu Han, Robert X. Gao, and Zhaoyan Fan. "Occupancy and indoor environment quality sensing for smart buildings." In: *2012 IEEE International Instrumentation and Measurement Technology Conference (I2MTC).* May 2012, pp. 882–887. DOI: 10.1109/I2MTC.2012.6229557.

[HHY95] John E. Hershey, Amer A. Hassan, and Rao Yarlagadda. "Unconventional cryptographic keying variable management." In: *IEEE Transactions on Communications* 43.1 (1995), pp. 3–6. DOI: 10.1109/26.385951.

[HK05] Gerhard P. Hancke and Markus G. Kuhn. "An RFID distance bounding protocol." In: *First International Conference on Security and Privacy for Emerging Areas in Communications Networks (SECURECOMM'05).* IEEE. 2005, pp. 67–73. DOI: 10.1109/SECURECOMM.2005.56.

[Hon15] Honeywell. *HONEYWELL CONNECTED HOME PRIVACY STATEMENT.* Dec. 2015. URL: https://web.archive.org/web/20151206210204/https://yourhome.honeywell.com/privacy-statement.

[HPP+16] Verena Huppert, Jan Paulus, Ute Paulsen, Martin Burkart, Bernd Wullich, and Bjoern Eskofier. "Quantification of Nighttime Micturition With an Ambulatory Sensor-Based System." In: *IEEE Journal of Biomedical and Health Informatics* 20.3 (May 2016), pp. 865–872. DOI: 10.1109/JBHI.2015.2421487.

[HTF09] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* 2nd ed. Springer Series in Statistics. New York, NY, USA: Springer, 2009. ISBN: 978-0-387-84857-0.

[ico15] icontrol Networks. *2015 State of the Smart Home Report.* June 2015. URL: https://web.archive.org/web/20151218231614/http://www.icontrol.com/wp-content/uploads/2015/06/Smart_Home_Report_2015.pdf.

[IDC24] IDC. *Wearables unit shipments worldwide from 2014 to 2028 (in millions).* Statista, Apr. 2024. URL: https://www.statista.com/statistics/437871/wearables-worldwide-shipments/.

[IEC16] IEC. *IEC 62591:2016 – Industrial networks – Wireless communication network and communication profiles – WirelessHART.* IEC. Mar. 2016.

[IEE11] IEEE. "IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)." In: *IEEE Std 802.15.4-2011 (Revision of IEEE Std 802.15.4-2006)* (Sept. 2011), pp. 1–314. DOI: 10.1109/IEEESTD.2011.6012487.

[IEE21] IEEE. "IEEE Standard for Information Technology—Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks—Specific Requirements – Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications." In: *IEEE Std 802.11-2020 (Revision of IEEE Std 802.11-2016)* (Feb. 2021), pp. 1–4379. DOI: 10.1109/IEEESTD.2021.9363693.

[IEE97] IEEE. "802.11-1997 - IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications." In: *IEEE Std 802.11-1997* (Nov. 1997), pp. 1–445. DOI: 10.1109/IEEESTD.1997.85951.

[Int16] Intel Security. *Intel Security's International Internet of Things Smart Home Survey Shows Many Respondents Sharing Personal Data for Money*. Intel, Mar. 2016. URL: https://web.archive.org/web/20171230165918/https://newsroom.intel.com/news-releases/intel-securitys-international-internet-of-things-smart-home-survey/.

[ITU23] ITU. *Number of fixed broadband internet subscriptions worldwide from 2005 to 2023 (in millions)*. Statista, Nov. 2023. URL: https://www.statista.com/statistics/268673/number-of-broadband-internet-subscriptions/.

[JBKE16] Ulf Jensen, Peter Blank, Patrick Kugler, and Bjoern Eskofier. "Unobtrusive and Energy-Efficient Swimming Exercise Tracking Using On-Node Processing." In: *IEEE Sensors Journal* 16.10 (May 2016), pp. 3972–3980. DOI: 10.1109/JSEN.2016.2530019.

[JJK11] Marek Jawurek, Martin Johns, and Florian Kerschbaum. "Plug-In Privacy for Smart Metering Billing." In: *Privacy Enhancing Technologies - 11th International Symposium, PETS 2011, Waterloo, ON, Canada, July 27-29, 2011. Proceedings*. Ed. by Simone Fischer-Hübner and Nicholas Hopper. Vol. 6794. Lecture Notes in Computer Science. Springer, 2011, pp. 192–210. ISBN: 978-3-642-22262-7. DOI: 10.1007/978-3-642-22263-4_11.

[JKD12] Marek Jawurek, Florian Kerschbaum, and George Danezis. *SoK: Privacy Technologies for Smart Grids – A Survey of Options*. Tech. rep. MSR-TR-2012-119. Microsoft, Nov. 2012. URL: https://www.microsoft.com/en-us/research/publication/privacy-technologies-for-smart-grids-a-survey-of-options/.

[JLS+22] Ga-Hye Jeon, Ji-Hyun Lee, Yeon-Su Sung, Hyun-Ju Park, You-Jin Lee, Sun-Woo Yun, and Il-Gu Lee. "Cooperative Friendly Jamming Techniques for Drone-Based Mobile Secure Zone." In: *Sensors* 22.3 (2022). ISSN: 1424-8220. DOI: 10.3390/s22030865.

[JSJ+25] Chenghua Jin, Yuxin Song, Yan Jia, Qingyin Tan, Rui Yang, and Zheli Liu. *Security and Privacy Measurement on Chinese Consumer IoT Traffic based on Device Lifecycle*. Sept. 28, 2025. arXiv: 2505.09929v5 [cs.CR].

[JWKD19] Furqan Jameel, Shurjeel Wyne, Georges Kaddoum, and Trung Q. Duong. "A Comprehensive Survey on Cooperative Relaying and Jamming Strategies for Physical Layer Security." In: *IEEE Communications Surveys & Tutorials* 21.3 (2019), pp. 2734–2771. ISSN: 1553-877X. DOI: 10.1109/COMST.2018.2865607.

[JZJ22] Rong Jin, Kai Zeng, and Chuan Jiang. "Friendly spectrum jamming against MIMO eavesdropping." In: *Wireless Networks* 28.6 (Aug. 2022), pp. 2437–2453. ISSN: 1572-8196. DOI: 10.1007/s11276-022-02967-1.

[JZZ21] Rong Jin, Kai Zeng, and Kai Zhang. "A Reassessment on Friendly Jamming Efficiency." In: *IEEE Transactions on Mobile Computing* 20.1 (2021), pp. 32–47. DOI: 10.1109/TMC.2019.2940941.

[KA09] Chong Hee Kim and Gildas Avoine. "RFID Distance Bounding Protocol with Mixed Challenges to Prevent Relay Attacks." In: *Cryptology and Network Security*. Ed. by Juan A. Garay, Atsuko Miyaji, and Akira Otsuka. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 119–133. ISBN: 978-3-642-10433-6. DOI: 10.1007/978-3-642-10433-6_9.

[Kat02] Jonathan Katz. "Efficient Cryptographic Protocols Preventing 'Man-in-the-Middle' Attacks." PhD thesis. Columbia University, 2002.

[KDK11] Klaus Kursawe, George Danezis, and Markulf Kohlweiss. "Privacy-Friendly Aggregation for the Smart-Grid." In: *Privacy Enhancing Technologies - 11th International Symposium, PETS 2011, Waterloo, ON, Canada, July 27-29, 2011. Proceedings*. Ed. by Simone Fischer-Hübner and Nicholas Hopper. Vol. 6794. Lecture Notes in Computer Science. Springer, 2011, pp. 175–191. ISBN: 978-3-642-22262-7. DOI: 10.1007/978-3-642-22263-4_10.

[KHT19] Da-Yu Kao, Shou-Ching Hsia, and Raylin Tso. "Analyzing WannaCry Ransomware Considering the Weapons and Exploits." In: *2019 21st International Conference on Advanced Communication Technology (ICACT)*. Feb. 2019, pp. 1098–1107. DOI: 10.23919/ICACT.2019.8702049.

[KNEK08] Tim van Kasteren, Athanasios Noulas, Gwenn Englebienne, and Ben Kröse. "Accurate activity recognition in a home setting." In: *Proceedings of the 10th International Conference on Ubiquitous Computing*. Association for Computing Machinery. 2008. DOI: 10.1145/1409635.1409637.

[LDW+20] Xuran Li, Hong-Ning Dai, Qubeijian Wang, Muhammad Imran, Dengwang Li, and Muhammad Ali Imran. "Securing Internet of Medical Things with Friendly-jamming schemes." In: *Computer Communications* 160 (2020), pp. 431–442. ISSN: 0140-3664. DOI: 10.1016/j.comcom.2020.06.026.

[Lee14] Franciscus Wilhelmus Adrianus Alphonsus Van Leeuwen. *Network discovery with touchlink option*. WO Patent App. PCT/IB2013/056663. Feb. 2014. URL: https://www.google.com/patents/WO2014030103A2.

[LHD+09] Khee Poh Lam, Michael Höynck, Bing Dong, Burton Andrews, Yun-shang Chiou, Diego Benitez, and Joonho Choi. "Occupancy detection through an extensive environmental sensor network in an open-plan office building." In: *Building Simulation 2009*. IBPSA. 2009, pp. 1452–1459. URL: https://www.aivc.org/resource/occupancy-detection-through-extensive-environmental-sensor-network-open-plan-office.

[LLD+22] Jianglong Li, Xianfu Lei, Panagiotis D. Diamantoulakis, Lisheng Fan, and George K. Karagiannidis. "Security Optimization of Cooperative NOMA Networks With Friendly Jamming." In: *IEEE Transactions on Vehicular Technology* 71.12 (2022), pp. 13422–13428. DOI: 10.1109/TVT.2022.3200253.

[LLJL23] Sun-Jin Lee, Yu-Rim Lee, So-Eun Jeon, and Il-Gu Lee. "Machine learning-based jamming attack classification and effective defense technique." In: *Computers & Security* 128 (May 2023), p. 103169. ISSN: 0167-4048. DOI: 10.1016/j.cose.2023.103169.

[Log17] LogRhythm Labs. *NotPetya Technical Analysis*. Archived from: https://gallery.logrhythm.com/threat-intelligence-reports/notpetya-technical-analysis-logrhythm-labs-threat-intelligence-report.pdf. 2017. URL: https://web.archive.org/web/20241001182116/https://gallery.logrhythm.com/threat-intelligence-reports/notpetya-technical-analysis-logrhythm-labs-threat-intelligence-report.pdf.

[LoR20] LoRa Alliance Technical Committee. *LoRaWAN L2 1.0.4 Specification (TS001-1.0.4)*. Ed. by T. Kramp and O. Seller. Oct. 2020. URL: https://resources.lora-alliance.org/technical-specifications/ts001-1-0-4-lorawan-l2-1-0-4-specification.

[LoR22] LoRa Alliance Technical Committee Regional Parameters Workgroup. *RP002-1.0.4 LoRaWAN Regional Parameters*. Ed. by D. Kjendal. Sept. 2022. URL: https://resources.lora-alliance.org/technical-specifications/rp002-1-0-4-regional-parameters.

[LSS+10] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. "The smart thermostat: using occupancy sensors to save energy in homes." In: *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*. Association for Computing Machinery. 2010, pp. 211–224. DOI: 10.1145/1869983.1870005.

[MABM16] Christian Müller, Frederik Armknecht, Zinaida Benenson, and Philipp Morgner. "On the security of the ZigBee light link touchlink commissioning procedure." In: *Sicherheit 2016 – Sicherheit, Schutz und Zuverlässigkeit*. Ed. by Michael Meier, Delphine Reinhardt, and Steffen Wendzel. Lecture Notes in Informatics (LNI) – Proceedings, Volume P-256. Bonn: Gesellschaft für Informatik e.V., 2016, pp. 229–240. ISBN: 978-3-88579-650-3. URL: https://dl.gi.de/items/12790b8b-aeb5-40ed-979a-158be5286220.

[MAM17] Vasily Mikhalev, Frederik Armknecht, and Christian Müller. "On Ciphers that Continuously Access the Non-Volatile Key." In: *IACR Transactions on Symmetric Cryptology* 2016.2 (Feb. 2017), pp. 52–79. DOI: 10.13154/tosc.v2016.i2.52-79.

[Mao04] Wenbo Mao. "A Structured Operational Modelling of the Dolev-Yao Threat Model." In: *Security Protocols 2002*. Ed. by Bruce Christianson, Bruno Crispo, James A. Malcolm, and Michael Roe. Vol. 2845. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 34–46. ISBN: 978-3-540-39871-4. DOI: 10.1007/978-3-540-39871-4_5.

[Mat16] Stephan Mattejat. "Traffic analysis and denial-of-service attacks in connected lighting systems." Bachelor's Thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg, June 2016.

[MBMA15] Philipp Morgner, Zinaida Benenson, Christian Müller, and Frederik Armknecht. "Design Space of Smart Home Networks from a Security Perspective." In: *Proceedings of the 14. GI/ITG KuVS Fachgespräch Sensornetze*. Ed. by Jürgen Eckert, Abdalkarim Awad, Kai-Steffen Hielscher, and David Eckhoff. Technical Reports, CS-2015-06. Friedrich-Alexander-Universität Erlangen-Nürnberg, Dept. of Computer Science, Sept. 2015, pp. 229–240. URL: https://open.fau.de/handle/openfau/6617.

[MDL+22] Jianhui Ma, Linsong Du, Chenxing Li, Chao Chen, and Qiang Li. "Jamming Modulation: Make Enemies Become Friends." In: *2022 10th International Workshop on Signal Design and Its Applications in Communications (IWSDA)*. IEEE. 2022, pp. 1–5. DOI: 10.1109/IWSDA50346.2022.9870594.

[MLL+23] Jianhui Ma, Qiang Li, Zilong Liu, Linsong Du, Hongyang Chen, and Nirwan Ansari. "Jamming Modulation: An Active Anti-Jamming Scheme." In: *IEEE Transactions on Wireless Communications* 22.4 (2023), pp. 2730–2743. DOI: 10.1109/TWC.2022.3213572.

[MMB+17] Philipp Morgner, Stephan Mattejat, Zinaida Benenson, Christian Müller, and Frederik Armknecht. "Insecure to the touch: attacking ZigBee 3.0 via touchlink commissioning." In: *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. Ed. by Guevara Noubir, Mauro Conti, and Sneha Kumar Kasera. WiSec '17. Boston, Massachusetts, USA: Association for Computing Machinery, 2017, pp. 230–240. ISBN: 978-1-4503-5084-6. DOI: 10.1145/3098243.3098254.

[MMR+17] Philipp Morgner, Christian Müller, Matthias Ring, Björn Eskofier, Christian Riess, Frederik Armknecht, and Zinaida Benenson. "Privacy Implications of Room Climate Data." In: *Computer Security – ESORICS 2017. 22nd European Symposium on Research in Computer Security, Oslo, Norway, September 11–15, 2017, Proceedings, Part II*. Ed. by Simon N. Foley, Dieter Gollmann, and Einar Snekkenes. Vol. 10493. Lecture Notes in Computer Science. Cham: Springer, Aug. 2017, pp. 324–343. ISBN: 978-3-319-66399-9. DOI: 10.1007/978-3-319-66399-9_18.

[MMSA25] Christian Müller, Vasily Mikhalev, Yves T. Staudenmaier, and Frederik Armknecht. "One Antenna is Enough for Neutralizing Jamming in IEEE 802.15.4 BPSK." In: *Advances in Information and Computer Security. 20th International Workshop on Security, IWSEC 2025, Fukuoka, Japan, November 25–27, 2025, Proceedings*. Ed. by Carlos Cid and Naoto Yanai. Vol. 16208. Lecture Notes in Computer Science. Singapore: Springer Nature Singapore, Nov. 25, 2025, pp. 489–509. ISBN: 978-981-95-4674-9. DOI: 10.1007/978-981-95-4674-9_25.

[Mot06] Moteiv Corporation. *Tmote Sky Datasheet*. Archived from: https://insense.cs.st-andrews.ac.uk/files/2013/04/tmote-sky-datasheet.pdf. 2006. URL: https://web.archive.org/web/20220120072303/https://insense.cs.st-andrews.ac.uk/files/2013/04/tmote-sky-datasheet.pdf.

[MPC16] Homay Danaei Mehr, Huseyin Polat, and Aydin Cetin. "Resident activity recognition in smart homes by using artificial neural networks." In: *2016 4th International Istanbul Smart Grid Congress and Fair (ICSG)*. Apr. 2016, pp. 1–5. DOI: 10.1109/SGCF.2016.7492428.

[MPS09] Ivan Martinovic, Paul Pichota, and Jens B Schmitt. "Jamming for good: a fresh approach to authentic communication in WSNs." In: *Proceedings of the second ACM conference on Wireless network security*. WiSec '09. Zurich, Switzerland: Association for Computing Machinery, 2009, pp. 161–168. ISBN: 978-1-60558-460-7. DOI: 10.1145/1514274.1514298.

[MSC15] Mustafa K. Masood, Yeng Chai Soh, and Victor W.-C. Chang. "Real-time occupancy estimation using environmental parameters." In: *2015 International Joint Conference on Neural Networks, IJCNN 2015, Killarney, Ireland, July 12-17, 2015*. IEEE, 2015, pp. 1–8. ISBN: 978-1-4799-1960-4. DOI: 10.1109/IJCNN.2015.7280781.

[MSF+10] Andrés Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin. "Private Memoirs of a Smart Meter." In: *Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*. BuildSys '10. Zurich, Switzerland: Association for Computing Machinery, 2010, pp. 61–66. ISBN: 978-1-4503-0458-0. DOI: 10.1145/1878431.1878446.

[MSMA25] Christian Müller, Yves T. Staudenmaier, Vasily Mikhalev, and Frederik Armknecht. "On The Recovery Of Jammed Bits." In: *crypto day matters 38*. 38th Crypto Day (Hamburg, Mar. 6–7, 2025). Ed. by Janik Schug and Michael Nüsken. Bonn: Gesellschaft für Informatik e.V. / FG KRYPTO, Nov. 2025. DOI: 10.18420/cdm-2025-38-14.

[Nes16] Nest. *Privacy Statement for Nest Products and Services*. Mar. 2016. URL: https://web.archive.org/web/20160509070702/https://nest.com/legal/privacy-statement-for-nest-products-and-services/.

[NN23] Hai N. Nguyen and Guevara Noubir. "JaX: Detecting and Cancelling High-power Jammers Using Convolutional Neural Network." In: *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. WiSec '23. New York, NY, USA: Association for Computing Machinery, June 2023, pp. 293–304. ISBN: 978-1-4503-9859-6. DOI: 10.1145/3558482.3590178.

[OFl11] Colin P. O'Flynn. "Message Denial and Alteration on IEEE 802.15.4 Low-Power Radio Networks." In: *2011 4th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*. IEEE. 2011, pp. 1–5. DOI: 10.1109/NTMS.2011.5720580.

[OHA+14] Olayemi Olawumi, Keijo Haataja, Mikko Asikainen, Niko Vidgren, and Pekka Toivanen. "Three practical attacks against ZigBee security: Attack scenario definitions, practical experiments, countermeasures, and lessons learned." In: *14th International Conference on Hybrid Intelligent Systems, HIS 2014, Kuwait, December 14-16, 2014*. IEEE, 2014, pp. 199–206. ISBN: 978-1-4799-7633-1. DOI: 10.1109/HIS.2014.7086198.

[Pal20] Palo Alto Networks Unit 42. *2020 Unit 42 IoT Threat Report*. Archived from: https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/reports/Unit_42/2020-unit-42-IoT-threat-report.pdf. Mar. 10, 2020. URL: https://web.archive.org/web/20210922032139/https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/reports/Unit_42/2020-unit-42-IoT-threat-report.pdf.

[Phi15] Philips. *Friends of Hue - Update*. Dec. 2015. URL: https://web.archive.org/web/20151224204750/https://www.developers.meethue.com/documentation/friends-hue-update.

[PKS23] Hans-Christian Pape, Armin Kurtz, and Stefan Silbernagl. *Physiologie*. German. 10th ed. Georg Thieme Verlag, 2023, p. 1088. ISBN: 978-3-13-244608-3.

[PLD14] Jean-Michel Picod, Arnaud Lebrun, and Jonathan-Christofer Demay. "Bringing software defined radio to the penetration testing community." In: *Black Hat USA Conference*. 2014. URL: https://www.blackhat.com/us-14/archives.html#bringing-software-defined-radio-to-the-penetration-testing-community.

[PNP17] Theis Heidmann Pedersen, Kasper Ubbe Nielsen, and Steffen Petersen. "Method for room occupancy detection based on trajectory of indoor climate sensor data." In: *Building and Environment* 115 (2017), pp. 147–156. ISSN: 0360-1323. DOI: 10.1016/j.buildenv.2017.01.023.

[PTDC11] Christina Pöpper, Nils Ole Tippenhauer, Boris Danev, and Srdjan Capkun. "Investigation of Signal and Message Manipulations on the Wireless Channel." In: *Computer Security – ESORICS 2011*. Ed. by Vijay Atluri and Claudia Diaz. Vol. 6879. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer-Verlag Berlin Heidelberg, 2011, pp. 40–59. DOI: 10.1007/978-3-642-23822-2_3.

[PZ22] Hossein Pirayesh and Huacheng Zeng. "Jamming Attacks and Anti-Jamming Strategies in Wireless Networks: A Comprehensive Survey." In: *IEEE Communications Surveys & Tutorials* 24.2 (2022), pp. 767–809. ISSN: 1553-877X. DOI: 10.1109/COMST.2022.3159185.

[Rap24] Theodore S. Rappaport. *Wireless Communications: Principles and Practice*. Updated Second Edition. Cambridge, United Kingdom: Cambridge University Press, Feb. 2024. ISBN: 978-1-009-48984-3. DOI: 10.1017/9781009489843.

[RC10] Kasper Bonne Rasmussen and Srdjan Capkun. "Realization of RF Distance Bounding." In: *19th USENIX Security Symposium (USENIX Security 10)*. Washington, DC: USENIX Association, Aug. 2010. URL: https://www.usenix.org/conference/usenixsecurity10/realization-rf-distance-bounding.

[RCHC09] Kasper Bonne Rasmussen, Claude Castelluccia, Thomas S. Heydt-Benjamin, and Srdjan Capkun. "Proximity-based access control for implantable medical devices." In: *Proceedings of the 16th ACM Conference on Computer and Communications Security*. CCS '09. Chicago, Illinois, USA: Association for Computing Machinery, 2009, pp. 410–419. ISBN: 978-1-60558-894-0. DOI: 10.1145/1653662.1653712.

[RCo14] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria, 2014. URL: https://www.R-project.org.

[RD11] Alfredo Rial and George Danezis. "Privacy-preserving Smart Metering." In: *Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society*. WPES '11. Chicago, Illinois, USA: Association for Computing Machinery, 2011, pp. 49–60. ISBN: 978-1-4503-1002-4. DOI: 10.1145/2046556.2046564.

[RD16] Lee Rainie and Maeve Duggan. *Pew Research: Privacy and Information Sharing*. Jan. 2016. URL: https://www.pewinternet.org/2016/01/14/privacy-and-information-sharing/.

[RDC15] Aanjhan Ranganathan, Boris Danev, and Srdjan Capkun. "Proximity Verification for Contactless Access Control and Authentication Systems." In: *Proceedings of the 31st Annual Computer Security Applications Conference*. ACSAC '15. Los Angeles, CA, USA: Association for Computing Machinery, 2015, pp. 271–280. ISBN: 978-1-4503-3682-6. DOI: 10.1145/2818000.2818004.

[REC15] Andreas Reinhardt, Frank Englert, and Delphine Christin. "Averting the Privacy Risks of Smart Metering by Local Data Preprocessing." In: *Pervasive and Mobile Computing (PMC)* 16 (2015), pp. 171–183. DOI: 10.1016/j.pmcj.2014.10.002.

[RFC4944] Gabriel Montenegro, Jonathan Hui, David Culler, and Nandakishore Kushalnagar. *Transmission of IPv6 Packets over IEEE 802.15.4 Networks*. RFC 4944. Sept. 2007. DOI: 10.17487/RFC4944.

[RJKE12] Matthias Ring, Ulf Jensen, Patrick Kugler, and Bjoern Eskofier. "Software-based performance and complexity analysis for the design of embedded classification systems." In: *Proceedings of the 21st International Conference on Pattern Recognition, ICPR 2012, Tsukuba, Japan, November 11-15, 2012*. IEEE Computer Society, 2012, pp. 2266–2269. ISBN: 978-1-4673-2216-4. URL: https://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=6425799.

[RK14] Hanif Rahbari and Marwan Krunz. "Friendly CryptoJam: A Mechanism for Securing Physical-Layer Attributes." In: *Proceedings of the 2014 ACM conference on Security and privacy in wireless & mobile networks*. WiSec '14. Oxford, United Kingdom: Association for Computing Machinery, 2014, pp. 129–140. ISBN: 978-1-4503-2972-9. DOI: 10.1145/2627393.2627415.

[RS16] Eyal Ronen and Adi Shamir. "Extended Functionality Attacks on IoT Devices: The Case of Smart Lights." In: *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*. IEEE, 2016, pp. 3–12. ISBN: 978-1-5090-1751-5. DOI: 10.1109/EuroSP.2016.13.

[RSWO17] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin O'Flynn. "IoT Goes Nuclear: Creating a ZigBee Chain Reaction." In: *2017 IEEE Symposium on Security and Privacy (SP)*. 2017, pp. 195–212. DOI: 10.1109/SP.2017.14.

[RT16] Marco Rocchetto and Nils Ole Tippenhauer. "CPDY: Extending the Dolev-Yao Attacker with Physical-Layer Interactions." In: *Formal Methods and Software Engineering*. Ed. by Kazuhiro Ogata, Mark Lawford, and Shaoying Liu. Cham: Springer International Publishing, 2016, pp. 175–192. ISBN: 978-3-319-47846-3. DOI: 10.1007/978-3-319-47846-3_12.

[RTŠ+12] Aanjhan Ranganathan, Nils Ole Tippenhauer, Boris Škorić, Dave Singelée, and Srdjan Čapkun. "Design and Implementation of a Terrorist Fraud Resilient Distance Bounding System." In: *Computer Security – ESORICS 2012: 17th European Symposium on Research in Computer Security, Pisa, Italy, September 10-12, 2012. Proceedings*. Ed. by Sara Foresti, Moti Yung, and Fabio Martinelli. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 415–432. ISBN: 978-3-642-33167-1. DOI: 10.1007/978-3-642-33167-1_24.

[SABH15] Sarah Spiekermann, Alessandro Acquisti, Rainer Böhme, and Kai-Lung Hui. "The challenges of personal data markets and privacy." In: *Electronic Markets* 25.2 (Apr. 2015), pp. 161–167. ISSN: 1422-8890. DOI: 10.1007/s12525-015-0191-0.

[SABY23] Ramiz Salama, Fadi Al-Turjman, Dibyahash Bordoloi, and Satya Prakash Yadav. "Wireless Sensor Networks and Green Networking for 6G Communication – An Overview." In: *2023 International Conference on Computational Intelligence, Communication Technology and Networking (CICTN)*. Apr. 2023, pp. 830–834. DOI: 10.1109/CICTN57981.2023.10141262.

[SCFS16] Gina Sprint, Diane Cook, Roschelle Fritz, and Maureen Schmitter-Edgecombe. "Detecting Health and Behavior Change by Analyzing Smart Home Sensor Data." In: *2016 IEEE International Conference on Smart Computing (SMARTCOMP)*. May 2016, pp. 1–3. DOI: 10.1109/SMARTCOMP.2016.7501687.

[Sel14] Markus Selinger. *Test: Smart Home Kits Leave the Door Wide Open – for Everyone*. Apr. 2014. URL: https://web.archive.org/web/20140514064326/https://www.av-test.org/en/news/news-single-view/test-smart-home-kits-leave-the-door-wide-open-for-everyone/.

[Sem11] Sergei Semenov. "Modulation." In: *Modulation and Coding Techniques in Wireless Communications*. Ed. by Sergei Semenov and Evgenii Krouk. 1st ed. Chichester, West Sussex: John Wiley & Sons, Ltd, 2011. Chap. 2, pp. 21–82. ISBN: 978-0-470-97677-7. DOI: 10.1002/9780470976777.ch2.

[Sha16] Chetan Sharma. *Correcting the IoT History*. Chetan Sharma Consulting, LLC, 2016. URL: https://web.archive.org/web/20240112155817/https://chetansharma.com/correcting-the-iot-history/.

[SMA22] Jochen Schäfer, Christian Müller, and Frederik Armknecht. "If You Like Me, Please Don't "Like" Me: Inferring Vendor Bitcoin Addresses From Positive Reviews." In: *Proceedings on Privacy Enhancing Technologies* 2022.1 (Jan. 2022), pp. 440–459. DOI: 10.2478/popets-2022-0022.

[SMA24] Jochen Schäfer, Christian Müller, and Frederik Armknecht. "0 Sterne für die Sicherheit: Wie Kundenbewertungen die Bitcoin-Adressen von Darknet-Händlern verraten." In: *Sicherheit 2024*. Ed. by Steffen Wendzel, Christian Wressnegger, Laura Hartmann, Felix Freiling, Frederik Armknecht, and Lena Reinfelder. Lecture Notes in Informatics (LNI) – Proceedings Volume P-345. Bonn: Gesellschaft für Informatik e.V., 2024, pp. 153–170. ISBN: 978-3-88579-739-5. DOI: 10.18420/sicherheit2024_010.

[SNHD13] Wenbo Shen, Peng Ning, Xiaofan He, and Huaiyu Dai. "Ally Friendly Jamming: How to Jam Your Enemy and Maintain Your Own Wireless Connectivity at the Same Time." In: *2013 IEEE Symposium on Security and Privacy*. 2013, pp. 174–188. DOI: 10.1109/SP.2013.22.

[SP07] Dave Singelée and Bart Preneel. "Distance Bounding in Noisy Environments." In: *Security and Privacy in Ad-hoc and Sensor Networks: 4th European Workshop, ESAS 2007, Cambridge, UK, July 2-3, 2007. Proceedings*. Ed. by Frank Stajano, Catherine Meadows, Srdjan Capkun, and Tyler Moore. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 101–115. ISBN: 978-3-540-73275-4. DOI: 10.1007/978-3-540-73275-4_8.

[SSBC09] Patrick Schaller, Benedikt Schmidt, David Basin, and Srdjan Capkun. "Modeling and Verifying Physical Properties of Security Protocols for Wireless Networks." In: *2009 22nd IEEE Computer Security Foundations Symposium*. 2009, pp. 109–123. DOI: 10.1109/CSF.2009.6.

[Sta25] Statista. *Number of users of smart homes worldwide from 2018 to 2029 (in millions)*. Statista, July 1, 2025. URL: https://www.statista.com/forecasts/887613/number-of-smart-homes-in-the-smart-home-market-in-the-world.

[SW04] Naveen Sastry and David Wagner. "Security considerations for IEEE 802.15.4 networks." In: *Proceedings of the 2004 ACM Workshop on Wireless Security, Philadelphia, PA, USA, October 1, 2004*. Ed. by Markus Jakobsson and Adrian Perrig. Association for Computing Machinery, 2004, pp. 32–42. ISBN: 978-1-58113-925-9. DOI: 10.1145/1023646.1023654.

[TČ09] Nils Ole Tippenhauer and Srdjan Čapkun. "ID-Based Secure Distance Bounding and Localization." In: *Computer Security – ESORICS 2009: 14th European Symposium on Research in Computer Security, Saint-Malo, France, September 21-23, 2009. Proceedings*. Ed. by Michael Backes and Peng Ning. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 621–636. ISBN: 978-3-642-04444-1. DOI: 10.1007/978-3-642-04444-1_38.

[Tex13] Texas Instruments. *2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver. Product data sheet*. 2013. URL: https://www.ti.com/lit/ds/symlink/cc2420.pdf.

[TMRC13] Nils Ole Tippenhauer, Luka Malisa, Aanjhan Ranganathan, and Srdjan Capkun. "On Limitations of Friendly Jamming for Confidentiality." In: *2013 IEEE Symposium on Security and Privacy*. Berkeley, CA, USA: IEEE Computer Society Press, 2013, pp. 160–173. DOI: 10.1109/SP.2013.21.

[VBBM10] Joao P. Vilela, Matthieu Bloch, Joao Barros, and Steven W. McLaughlin. "Friendly jamming for wireless secrecy." In: *2010 IEEE International Conference on Communications*. IEEE. 2010, pp. 1–6. DOI: 10.1109/ICC.2010.5502606.

[VBBM11] Joao P. Vilela, Matthieu Bloch, Joao Barros, and Steven W. McLaughlin. "Wireless Secrecy Regions With Friendly Jamming." In: *IEEE Transactions on Information Forensics and Security* 6.2 (2011), pp. 256–266. DOI: 10.1109/TIFS.2011.2111370.

[VBN13] Triet D. Vo-Huu, Erik-Oliver Blass, and Guevara Noubir. "Counter-jamming using mixed mechanical and software interference cancellation." In: *Proceedings of the sixth ACM conference on Security and privacy in wireless and mobile networks*. WiSec '13. Budapest, Hungary: Association for Computing Machinery, 2013, pp. 31–42. ISBN: 978-1-4503-1998-0. DOI: 10.1145/2462096.2462103.

[VHP+13] Niko Vidgren, Keijo Haataja, Jose Luis Patino-Andres, Juan Jose Ramirez-Sanchis, and Pekka Toivanen. "Security Threats in ZigBee-Enabled Systems: Vulnerability Evaluation, Practical Experiments, Countermeasures, and Lessons Learned." In: *46th Hawaii International Conference on System Sciences, HICSS 2013, Wailea, HI, USA, January 7-10, 2013*. IEEE, 2013, pp. 5132–5138. ISBN: 978-1-4673-5933-7. DOI: 10.1109/HICSS.2013.475.

[VNT+20] Nguyen Van Huynh, Diep N. Nguyen, Dinh Thai Hoang, Eryk Dutkiewicz, and Markus Mueck. "Ambient Backscatter: A Novel Method to Defend Jamming Attacks for Wireless Networks." In: *IEEE Wireless Communications Letters* 9.2 (Feb. 2020), pp. 175–178. ISSN: 2162-2345. DOI: 10.1109/LWC.2019.2947417.

[WAF17] Daniel Wood, Noah Apthorpe, and Nick Feamster. "Cleartext Data Transmissions in Consumer IoT Medical Devices." In: *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy*. IoTS&P '17. Dallas, Texas, USA: Association for Computing Machinery, 2017, pp. 7–12. ISBN: 978-1-4503-5396-0. DOI: 10.1145/3139937.3139939.

[WBRW14] Dominic Wörner, Thomas von Bomhard, Marc Roeschlin, and Felix Wortmann. "Look twice: Uncover hidden information in room climate sensor data." In: *4th International Conference on the Internet of Things, IoT 2014, Cambridge, MA, USA, October 6-8, 2014*. IEEE, 2014, pp. 25–30. ISBN: 978-1-4799-5154-3. DOI: 10.1109/IOT.2014.7030110.

[WFH11] Ian H. Witten, Eibe Frank, and Mark A. Hall. *Data Mining: Practical Machine Learning Tools and Techniques*. 3rd ed. Burlington, MA, USA: Morgan Kaufmann Publishers, 2011. ISBN: 978-0-12-374856-0. DOI: 10.1016/C2009-0-19715-5.

[Wri09] Joshua Wright. "KillerBee: Practical ZigBee Exploitation Framework." ToorCon 11. 2009. URL: https://www.willhackforsushi.com/presentations/toorcon11-wright.pdf.

[Wyn75] Aaron D. Wyner. "The wire-tap channel." In: *The Bell System Technical Journal* 54.8 (1975), pp. 1355–1387. DOI: 10.1002/j.1538-7305.1975.tb02040.x.

[XTZW05] Wenyuan Xu, Wade Trappe, Yanyong Zhang, and Timothy Wood. "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks." In: *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. MobiHoc '05. Urbana-Champaign, IL, USA: ACM, New York, NY, USA, 2005, pp. 46–57. ISBN: 978-1-59593-004-0. DOI: 10.1145/1062689.1062697.

[YAE23] Okan Yaman, Tolga Ayav, and Yusuf Murat Erten. "A Lightweight Self-Organized Friendly Jamming." In: *International Journal of Information Security Science* 12.1 (2023), pp. 13–20. DOI: 10.55859/ijiss.1194643.

[YLBO14] Zheng Yang, Nan Li, Burcin Becerik-Gerber, and Michael D. Orosz. "A systematic approach to occupancy modeling in ambient sensor-rich buildings." In: *Simulation* 90.8 (2014), pp. 960–977. DOI: 10.1177/0037549713489918.

[YLQ+12] Weining Yang, Ninghui Li, Yuan Qi, Wahbeh Qardaji, Stephen McLaughlin, and Patrick McDaniel. "Minimizing private data disclosures in the smart grid." In: *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. CCS '12. Raleigh, North Carolina, USA: Association for Computing Machinery, 2012, pp. 415–427. ISBN: 978-1-4503-1651-4. DOI: 10.1145/2382196.2382242.

[YZJ+14] Qiben Yan, Huacheng Zeng, Tingting Jiang, Ming Li, Wenjing Lou, and Yiwei Thomas Hou. "MIMO-based jamming resilient communication in wireless networks." In: *IEEE INFOCOM 2014 – IEEE Conference on Computer Communications*. 2014, pp. 2697–2706. DOI: 10.1109/INFOCOM.2014.6848218.

[YZJ+16] Qiben Yan, Huacheng Zeng, Tingting Jiang, Ming Li, Wenjing Lou, and Yiwei Thomas Hou. "Jamming Resilient Communication Using MIMO Interference Cancellation." In: *IEEE Transactions on Information Forensics and Security* 11.7 (July 2016), pp. 1486–1499. ISSN: 1556-6021. DOI: 10.1109/TIFS.2016.2535906.

[ZCLY17] Huacheng Zeng, Chen Cao, Hongxiang Li, and Qiben Yan. "Enabling jamming-resistant communications in wireless MIMO networks." In: *2017 IEEE Conference on Communications and Network Security (CNS)*. IEEE. 2017, pp. 1–9. DOI: 10.1109/CNS.2017.8228622.

[Zig12a] ZigBee Alliance. *ZigBee Light Link Standard Version 1.0 – Document 11-0037-10*. Apr. 2012.

[Zig12b] ZigBee Alliance. *ZigBee Smart Energy Standard Version 1.1b – Document 07-5356-18*. Nov. 2012. URL: https://zigbeealliance.org/wp-content/uploads/2021/10/docs-07-5356-18-0zse-zigbee-smart-energy-profile-specification.pdf.

[Zig12c] ZigBee Standards Organization. *ZigBee Specification – Document 053474r20*. Sept. 2012.

[Zig13] ZigBee Alliance. *ZigBee Home Automation Public Application Profile Version 1.2 – Document 05-3520-29*. June 2013.

[Zig16a] ZigBee Alliance. *Base Device Behavior Specification Version 1.0 – Document 13-0402-13*. Feb. 2016. URL: https://zigbeealliance.org/wp-content/uploads/2019/12/docs-13-0402-13-00zi-Base-Device-Behavior-Specification-2-1.pdf.

[Zig16b] ZigBee Alliance. *ZigBee Alliance Accelerates IoT Unification with 20 ZigBee 3.0 Platform Certifications*. Dec. 2016. URL: https://web.archive.org/web/20161218122007/https://www.zigbee.org/zigbee-alliance-accelerates-iot-unification-with-20-zigbee-3-0-platform-certifications/.

[Zig16c] ZigBee Alliance. *ZigBee Cluster Library Specification Revision 6 – Document 07-5123-06*. Jan. 2016. URL: https://zigbeealliance.org/wp-content/uploads/2019/12/07-5123-06-zigbee-cluster-library-specification.pdf.

[Zig17a] ZigBee Alliance. *The zigbee alliance to Unveil Universal Language for the IoT from CES 2017 – Making it Possible for Smart Objects to Work Together on Any Network*. Jan. 2017. URL: https://web.archive.org/web/20170106135954/https://www.zigbee.org/the-zigbee-alliance-to-unveil-universal-language-for-the-iot-from-ces-2017-making-it-possible-for-smart-objects-to-work-together-on-any-network/.

[Zig17b] ZigBee Alliance. *ZigBee Certified Products*. 2017. URL: http://www.zigbee.org/zigbee-products-2/.

[zig17] zigbee alliance. *zigbee Specification Revision 22.1.0 – Document 05-3474-22*. zigbee alliance. Apr. 2017.

[ZLCD12] Rui Zhang, Khee Poh Lam, Yun-Shang Chiou, and Bing Dong. "Information-theoretic environment features selection for occupancy detection in open office spaces." In: *Building Simulation* 5.2 (2012), pp. 179–188. ISSN: 1996-8744. DOI: 10.1007/s12273-012-0075-6.

[ZPP16] Christian T. Zenger, Mario Pietersz, and Christof Paar. "Preventing relay attacks and providing perfect forward secrecy using PHYSEC on 8-bit μC." In: *2016 IEEE International Conference on Communications Workshops (ICC)*. May 2016, pp. 110–115. DOI: 10.1109/ICCW.2016.7503773.

[ZS15] Tobias Zillner and Sebastian Strobl. "ZigBee Exploited – The good, the bad and the ugly." Black Hat USA. 2015. URL: https://www.blackhat.com/us-15/briefings.html#zigbee-exploited-the-good-the-bad-and-the-ugly.

[Zsc19] Zscaler ThreatLabZ. *IoT in the Enterprise: An Analysis of Traffic and Threats*. May 22, 2019. URL: https://www.zscaler.com/press/zscaler-reveals-enterprise-iot-ecosystem-extends-beyond-corporate-devices-56-million-iot-device-transactions-over-30-days.

[ZWF18] Lars Zimmermann, Robert Weigel, and Georg Fischer. "Fusion of Non-Intrusive Environmental Sensors for Occupancy Detection in Smart Homes." In: *IEEE Internet of Things Journal* 5.4 (Aug. 2018), pp. 2343–2352. ISSN: 2327-4662. DOI: 10.1109/JIOT.2017.2752134.

[ZZP+15] Christian T. Zenger, Jan Zimmer, Mario Pietersz, Jan-Felix Posielek, and Christof Paar. "Exploiting the Physical Environment for Securing the Internet of Things." In: *Proceedings of the 2015 New Security Paradigms Workshop*. NSPW '15. Twente, Netherlands: Association for Computing Machinery, 2015, pp. 44–58. ISBN: 978-1-4503-3754-0. DOI: 10.1145/2841113.2841117.