

Diploma Project

**ANALYSIS OF
JAMMING EFFECTS IN A
STATIC LORA
WIRELESS SENSOR
NETWORK**

Ioannis Kyriacou



University of Cyprus
**Department of Computer
Science**

Department of Computer Science

April 2026

UNIVERSITY OF CYPRUS
Faculty of Pure and Applied Sciences
Department of Computer Science

**ANALYSIS OF
JAMMING EFFECTS IN A
STATIC LORA
WIRELESS SENSOR
NETWORK**

Ioannis Kyriacou

Advisor

Professor Vasos Vassiliou

The Thesis was submitted in partial fulfillment of the requirements for obtaining the
Computer Science degree of the Department of Computer Science of the University of
Cyprus

April 2026

Acknowledgements

I would like to express my sincere gratitude to Dr. Michalis Savva, Research Scientist in the Department of Computer Science at the University of Cyprus, for his valuable assistance, technical advice, and the many discussions that contributed to the development of this thesis.

Abstract

Wireless sensor networks and Internet of Things deployments depend on low-power wireless communication, which makes them vulnerable to intentional interference and denial-of-service attacks. This thesis investigates jamming effects in a static indoor LoRa mesh experiment using six measurement sets. `test1` and `test2` represent normal network traffic without a jammer, `test3`, `test4`, and `test5` represent constant proactive jamming experiments, and `test6` represents a reactive-jammer experiment placed in front of node 5. The study analyzes transmit and receive logs from nine devices, estimates directional packet delivery ratio using unique packet sequence numbers, and compares physical-layer indicators including RSSI and SNR. It also introduces expected-versus-received packet graphs for each receiver so directional packet loss can be inspected directly. The results show that link quality is not determined by distance alone. Building layout, directional asymmetry, local interference, receiver position, jammer placement, and log quality all affect packet delivery and signal quality. The thesis also identifies the limits of the current experiment and proposes future work on machine-learning and fuzzy-logic intrusion detection.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
TABLE OF CONTENTS	v
LIST OF FIGURES	xi
LIST OF TABLES	xii
1 Introduction	1
1.1 Context	1
1.2 Problem Statement	1
1.3 Motivation	1
1.4 Aims	1
1.5 Research Questions	2
1.6 Thesis Structure	2
2 Theoretical Background	3
2.1 Long-Range Wide-Area Network (LoRa) Overview	3
2.1.1 LoRa Modulation Fundamentals	3
2.1.2 CSS Modulation Properties	3
2.2 Wireless Signal Propagation in Indoor Environments	4
2.2.1 Path Loss and Propagation Models	4
2.2.2 Multipath Propagation and Fading	5
2.2.3 Wall Attenuation and Shadowing	5
2.3 Jamming and Interference Fundamentals	5
2.3.1 Jammer Classification	5
2.3.1.1 Continuous vs. Pulsed Jamming	5
2.3.1.2 Waveform-Based Classification	6
2.3.2 Jamming Effectiveness Metrics	6
2.3.2.1 Packet Delivery Ratio (PDR)	6
2.3.2.2 Signal-to-Noise+Interference Ratio (SINR)	6
2.3.3 Jamming Detection	7
2.4 Indoor LoRa Propagation Challenges	7
2.4.1 Coherence Bandwidth	7
2.4.2 Doppler Shift	7
2.4.3 Adjacent-Channel Interference	8

2.5	Spread-Spectrum Resilience Limits	8
2.5.1	Matched Filter Attacks	8
2.5.2	Partial-Band Jamming	8
2.5.3	Multi-Path Jamming	8
2.6	Summary	8
3	Related Work and Literature Review	9
3.1	Overview	9
3.2	LoRa and LoRaWAN Security	9
3.2.1	LoRa Vulnerability Research	9
3.2.2	Interference and Robustness Studies	9
3.3	Jamming Detection and Prevention	10
3.3.1	Power-Based Detection Methods	10
3.3.2	Cross-Layer Detection Approaches	10
3.3.3	Machine Learning Approaches	10
3.4	Adversarial Signal Design	10
3.5	Indoor Wireless Propagation Measurement	11
3.6	Directional and Spatial Analysis	11
3.7	RSSI and SNR Measurement	11
3.8	Packet Delivery and Sequence Analysis	12
3.9	Repeated-Capture Methodology	12
3.10	Multi-Jammer and Comparative Studies	12
3.11	Gaps in Existing Literature	12
3.12	Contribution to the Field	13
3.13	Future Research Directions	13
4	Detailed Experimental Methodology	15
4.1	Testbed Architecture	15
4.1.1	Physical Deployment	15
4.1.2	Node Hardware	15
4.1.3	Jammer Hardware	15
4.2	Experimental Protocol	16
4.2.1	Baseline Capture (Tests 1–2)	16
4.2.2	Single-Run Jammer Tests (Tests 3–6)	16
4.2.3	Repeated Short-Capture Campaign (Tests 1–8 in constant_jammer_output/) 16	
4.2.4	Deceptive Jammer Campaign (Tests 1–3 in deceptive_jammer_output/) 16	
4.2.5	Random Jammer Campaign (Tests 1–3 in random_jammer_output/) 17	
4.2.6	Reactive Jammer Campaign (Tests 1–3 in reactive_jammer_output/) 17	
4.3	Data Collection and Logging	17
4.3.1	Packet Log Format	17
4.3.2	Synchronization Challenges	17
4.3.3	Log Cleaning	18
4.4	Packet Analysis Pipeline	18
4.4.1	Directional Link Extraction	18
4.4.2	Unique-Sequence PDR Calculation	19
4.4.3	RSSI and SNR Aggregation	19

4.5	Visualization and Plotting	19
4.5.1	Heatmap Generation	19
4.5.2	Topology Plots	19
4.5.3	Jammer Impact Plots	20
4.5.4	Node-Specific Focus	20
4.5.5	Proximity and Distance Plots	20
4.6	Reproducibility	20
5	Detailed Results and Statistical Analysis	21
5.1	Baseline Communication Quality	21
5.1.1	Test 1 and Test 2 Overview	21
5.1.2	Best-Performing Links	21
5.1.3	Worst-Performing Baseline Links	22
5.1.4	Device-Level Statistics	22
5.2	Constant Jammer Effects	23
5.2.1	Single-Run vs. Repeated-Run Comparison	23
5.2.2	Constant Jammer Impact on Individual Links	23
5.2.3	Constant Jammer: Least Affected Links	23
5.3	Multi-Jammer Comparative Analysis	24
5.3.1	PDR Degradation by Jammer Type	24
5.3.2	SNR Degradation by Jammer Type	25
5.4	Detection Metric Effectiveness	25
5.4.1	Metric Comparison Matrix	25
5.5	Spatial Analysis	26
5.5.1	Distance-Dependent Degradation	26
5.5.2	Directional Asymmetry	26
5.6	Reproducibility and Variance	27
5.6.1	Run-to-Run Variance in Repeated Captures	27
5.7	Summary of Key Results	27
6	Extended Discussion and Analysis	29
6.1	Implications of Spatial Degradation	29
6.1.1	Network-Level vs. Link-Level Metrics	29
6.1.2	Detection Sensitivity vs. False Positive Rate	29
6.2	Jammer-Type Characterization	29
6.2.1	Constant Jamming Characteristics	29
6.2.2	Deceptive Jamming Characteristics	30
6.2.3	Random Jamming Characteristics	30
6.2.4	Reactive Jamming Characteristics	31
6.3	Role of Distance in Jamming Effectiveness	32
6.3.1	Path Loss and Jammer Dominance	32
6.3.2	Indoor Propagation Variability	32
6.4	Implications for LoRa Network Design	33
6.4.1	Spreading Factor Selection	33
6.4.2	Channel Diversity	33
6.4.3	Gateway Placement and Redundancy	33

6.5	Machine Learning Opportunities	34
6.5.1	Feature Extraction	34
6.5.2	Classification Approaches	34
6.5.3	Real-Time Detection Challenge	34
6.6	Limitations and Future Work	35
6.6.1	Testbed Limitations	35
6.6.2	Potential Extensions	35
7	Background and Literature Review	36
7.1	Wireless Sensor Networks and IoT	36
7.2	Network Security Threats	36
7.3	Jamming Attacks	37
7.4	Intrusion Detection Techniques	38
7.5	Fuzzy-Logic and Machine-Learning Approaches for Jamming Detection . .	38
7.6	Localization Under Jamming	38
7.7	Detection, Localization, and Recovery Framework	39
8	Methodology	40
8.1	System Model	40
8.2	Experimental Setup	41
8.3	Experiment Scenarios	42
8.4	Performance Metrics	43
9	Implementation	45
9.1	Network Design	45
9.2	Jamming Models	46
9.2.1	Constant Jammer	47
9.2.2	Reactive Jammer	48
9.2.3	Random Jammer	49
9.2.4	Deceptive Jammer	50
10	Results	51
10.1	Overview	51
10.2	Baseline Tests	51
10.2.1	Baseline PDR Heatmaps	51
10.2.2	Cross-Baseline Comparison	52
10.3	Repeated Constant Jammer Experiments	53
10.3.1	Physical Jammer Placement	54
10.3.2	Run 3 Analysis	54
10.3.3	Run 4 Analysis	56
10.3.4	Run 5 Analysis	57
10.3.5	Run 6–8 Observations	58
10.4	Average Jammer Results	58
10.4.1	Average Heatmaps	59
10.4.2	Average Jammer Impact	60
10.4.3	Distance-from-Jammer Analysis	60
10.4.4	Node 5 Focused Analysis	61

10.5 Overall Interpretation	62
11 Discussion	63
11.1 Spatial and Directional Jamming Effects	63
11.2 Evidence from Measurements	63
11.3 Interpretation of LoRa Behaviour Under Jamming	64
11.4 Threats to Validity	65
11.5 Limitations	66
11.6 Implications for Detection	66
12 Conclusions	68
12.1 Summary	68
12.2 Answers to the Research Questions	68
12.3 Contribution	69
13 Future Work	71
13.1 Multi-Jammer Campaign Results	71
13.1.1 Key Findings	71
13.1.2 Jammer-Type Discrimination	71
13.2 Improved Data Collection	72
13.3 Detection and Experiment Extensions	72
References	75
References	75
APPENDICES	77
I Repository Files	78
II Additional Graphs	81
III Reproducibility Commands	82
IV Constant Jammer Analysis (Averaged)	83
IV.1 Overview	83
IV.2 Experimental Setup	83
IV.3 Analysis Output	83
IV.4 Key Observations	84
IV.4.1 Spatial Degradation	84
IV.4.2 Directional Asymmetry	84
IV.4.3 RSSI and SNR Trends	84
IV.4.4 Consistency Across Runs	84
IV.5 Comparison to Baseline	84
IV.6 Device Summary Statistics	85
IV.7 Link-Level Analysis	85
V Deceptive Jammer Analysis	86

V.1	Overview	86
V.2	Experimental Setup	86
V.3	Analysis Output	86
V.4	Key Observations	87
V.4.1	Deceptive vs. Constant Jammer Contrast	87
V.4.2	Spatial Distribution	87
V.4.3	Per-Run Variability	87
V.4.4	Device-Level Impact	87
V.5	Detection Considerations	87
V.6	Comparison Summary	88
VI	Random Jammer Analysis	89
VI.1	Overview	89
VI.2	Experimental Setup	89
VI.3	Analysis Output	89
VI.4	Key Observations	90
VI.4.1	Variability and Unpredictability	90
VI.4.2	Statistical Degradation	90
VI.4.3	Spatial Effects	90
VI.4.4	Per-Run Statistics	90
VI.5	Detection Characteristics	90
VI.5.1	Comparison to Other Jammer Types	91
VI.6	Comparison Metrics	91
VII	Reactive Jammer Analysis	92
VII.1	Overview	92
VII.2	Experimental Setup	92
VII.3	Analysis Output	92
VII.4	Key Observations	93
VII.4.1	Temporal Correlation with Traffic	93
VII.4.2	Detection Latency Effects	93
VII.4.3	Power Efficiency	93
VII.4.4	Spatial and Directional Effects	93
VII.4.5	Consistency and Reproducibility	93
VII.5	Detection Characteristics	94
VII.6	Comparison to Other Jammer Types	94
VII.7	Comparison Metrics	94
VIII	Multi-Jammer Comparison	95
VIII.1	Overview	95
VIII.2	Jammer Strategies	95
VIII.2.1	Constant Jammer	95
VIII.2.2	Deceptive Jammer	95
VIII.2.3	Random Jammer	95
VIII.2.4	Reactive Jammer	96
VIII.3	Comparative Metrics	96

VIII.3.1	Packet Delivery Ratio (PDR)	96
VIII.3.2	Signal-to-Noise Ratio (SNR)	96
VIII.3.3	Received Signal Strength (RSSI)	96
VIII.3.4	Spatial Extent	97
VIII.4	Detection Signatures	97
VIII.4.1	Constant Jammer Signature	97
VIII.4.2	Deceptive Jammer Signature	97
VIII.4.3	Random Jammer Signature	97
VIII.4.4	Reactive Jammer Signature	98
VIII.5	Optimal Detection Metrics	98
VIII.5.1	For Constant Jamming	98
VIII.5.2	For Deceptive Jamming	98
VIII.5.3	For Random Jamming	98
VIII.5.4	For Reactive Jamming	98
VIII.6	Recommendations	98
VIII.6.1	Multi-Metric Fusion	98
VIII.6.2	Adaptive Detection Strategy	98
VIII.6.3	Machine Learning Opportunity	99
VIII.7	Summary	99

LIST OF FIGURES

8.1	3D device placement.	42
9.1	Jammer placement near node 5 during the experiments.	47
9.2	Spectrum visualization of the constant jammer transmission.	48
10.1	Directional PDR heatmaps for the two no-jammer baseline captures.	52
10.2	Directional PDR differences between the two baseline captures.	53
10.3	Physical placement of the jammer relative to the LoRa nodes.	54
10.4	Run 3 heatmaps showing degradation concentrated around node 5.	55
10.5	Run 3 jammer impact relative to the averaged baseline.	55
10.6	Focused node 5 analysis for run 3.	56
10.7	Run 4 heatmaps showing continued degradation around node 5.	57
10.8	Run 4 proximity analysis.	57
10.9	Focused node 5 analysis for run 5.	58
10.10	Average jammer heatmaps across all repeated captures.	59
10.11	Average jammer impact relative to the averaged baseline.	60
10.12	Relationship between jammer proximity and communication degradation.	61
10.13	Average node 5 analysis across all jammer captures.	61

LIST OF TABLES

5.1	Five best-performing baseline links, ranked by PDR.	21
5.2	Worst-performing baseline links, ranked by PDR.	22
5.3	Per-device baseline transmission and reception statistics.	22
5.4	Five most-affected links by constant jammer (baseline to jammed comparison).	23
5.5	Links most resistant to constant jammer (smallest PDR degradation). . . .	24
5.6	Average PDR degradation by jammer type across all links in affected region.	24
5.7	Average SNR degradation by jammer type.	25
5.8	Detection metric effectiveness for each jammer type (1=poor, 5=excellent).	25
5.9	Constant jammer: PDR degradation grouped by distance from jammer position.	26
5.10	Directional asymmetry in jamming effects (constant jammer).	26
5.11	Run-to-run variance in repeated constant-jammer captures.	27
6.1	Higher spreading factors provide SNR margin but increase air-time.	33
6.2	Feature categories for jamming detection ML models.	34
8.1	Physical device number to LoRa node identifier.	40
8.2	Readable inter-device distances encoded from the building sketch.	41
9.1	LoRa node radio and traffic configuration from the static mesh firmware. .	45
9.2	HackRF constant-jammer configuration used during the repeated captures.	48
9.3	Reactive jammer configuration.	49
9.4	Random jammer configuration.	49
9.5	Deceptive jammer configuration.	50

1 Introduction

1.1 Context

Wireless networks are now a central part of modern communication systems [1]. They support mobile devices, industrial monitoring, smart buildings, environmental sensing, healthcare systems, and Internet of Things (IoT) applications [2]. Wireless sensor networks (WSNs) are especially important in this context because they allow small embedded devices to collect measurements and forward them through a shared radio medium [1]. These systems are attractive because they are low cost, flexible, and easy to deploy, but their dependence on open wireless channels also creates serious security risks [4].

1.2 Problem Statement

The main problem studied in this thesis is the effect of jamming attacks on a low-power wireless sensor network [6]. A jammer intentionally transmits interference in order to reduce the ability of legitimate nodes to exchange packets [4]. Unlike many higher-layer attacks, jamming can damage communication even when cryptographic protection is used, because the attack targets the availability of the physical wireless channel [5].

1.3 Motivation

The motivation for this work is the growing importance of reliable wireless communication in IoT and sensor-network deployments [7]. In practical environments, a loss of communication may cause missing measurements, delayed alarms, reduced situational awareness, or complete service failure [4]. Understanding how jamming changes packet delivery, signal quality, and network behavior is therefore an important step toward stronger network security and intrusion detection [6].

1.4 Aims

The aims of this thesis are:

- to study the behavior of a static LoRa wireless sensor network under normal and jammed conditions [8];
- to compare no-jammer experiments with jammer experiments using repeatable analysis scripts;
- to evaluate packet delivery ratio, RSSI, SNR, receive activity, and timing windows;
- to identify which metrics are useful for detecting abnormal wireless behavior [7];
- to discuss limitations and future extensions for intrusion detection.

1.5 Research Questions

The research questions are:

1. How does the packet delivery ratio change between normal operation and operation under jamming?
2. How do RSSI and SNR behave when a jammer is present?
3. Does physical distance alone explain the observed link quality?
4. Which measurements are most useful for detecting jamming in this experiment?
5. How does jammer placement affect individual devices and directional links?

1.6 Thesis Structure

The thesis is structured as follows. Section 2 presents background material on WSNs, IoT, wireless security threats, jamming attacks, and intrusion detection. Section 3 describes the methodology, system model, experiment scenarios, and metrics. Section 4 explains the implementation, network design, jamming model, and data collection process. Section 5 presents the experimental results. Section 6 discusses the interpretation and limitations. Section 7 concludes the thesis. Section 8 proposes future work. Section 9 lists references, and Section 10 provides appendices.

2 Theoretical Background

2.1 Long-Range Wide-Area Network (LoRa) Overview

LoRa is a proprietary long-range, low-power wireless communication technology developed by Semtech Corporation [8]. It operates in unlicensed sub-gigahertz frequency bands (typically 433 MHz, 868 MHz in Europe, or 915 MHz in North America) and is specifically designed for IoT applications requiring long communication range with minimal power consumption.

The LoRa Physical Layer employs Chirp Spread Spectrum (CSS) modulation, which offers inherent robustness against narrowband interference, frequency-selective fading, and multipath propagation. This property makes LoRa particularly attractive for harsh indoor and outdoor environments where traditional modulation schemes would suffer significant performance degradation.

2.1.1 LoRa Modulation Fundamentals

LoRa uses Chirp Spread Spectrum (CSS) modulation, where each symbol is encoded as a chirp—a frequency sweep across a defined bandwidth. The key modulation parameters are:

- **Spreading Factor (SF):** Controls how many chips represent each symbol. Higher SF (SF7–SF12) increases range but reduces data rate and increases air-time. Lower SF increases data rate but reduces robustness.
- **Bandwidth (BW):** Defines the frequency span of the chirp. Common values are 62.5 kHz, 125 kHz, 250 kHz, and 500 kHz. Wider bandwidth increases data rate but reduces sensitivity.
- **Coding Rate (CR):** Forward error correction (FEC) redundancy level, typically 4/5–4/8. Higher coding rates add redundancy, enabling recovery from corrupted symbols.
- **Center Frequency:** The RF frequency at which transmission occurs. Different regions use different frequency allocations.

The LoRa data rate is determined by:

$$\text{Data Rate} = \text{SF} \times \frac{\text{BW}}{2^{\text{SF}}} \times \text{CR}$$

For example, with SF=7, BW=125 kHz, and CR=4/5, the data rate is approximately 5.5 kbps. With SF=12, the same bandwidth yields only 0.3 kbps but extends range significantly (typically 10+ km in rural areas, 1–2 km in urban areas).

2.1.2 CSS Modulation Properties

The chirp modulation scheme provides several inherent advantages:

- **Processing Gain:** The spreading factor provides processing gain proportional to 2^{SF} , which enhances signal detection in the presence of noise and interference.
- **Multipath Resilience:** The chirp structure is robust to multipath propagation because the receiver correlates received signals with multiple time-shifted copies of the expected chirp.
- **Narrowband Rejection:** Narrowband interference occupying only a small portion of the bandwidth has reduced impact because the energy is spread across the spreading factor.
- **Frequency Diversity:** Different nodes can transmit on the same frequency with different spreading factors, allowing frequency reuse.

However, these advantages assume that interference power is distributed relatively uniformly across the signal bandwidth. Concentrated, wideband interference or carefully crafted adversarial signals can overwhelm the processing gain and corrupt transmissions despite the modulation's inherent robustness.

2.2 Wireless Signal Propagation in Indoor Environments

Indoor wireless propagation is significantly more complex than outdoor propagation due to reflection, diffraction, scattering, and absorption effects caused by walls, floors, furniture, and other obstacles.

2.2.1 Path Loss and Propagation Models

The received signal power decreases with distance according to path loss. The free-space path loss model is:

$$P_r = P_t + G_t + G_r - 20 \log_{10}(d) - 20 \log_{10}(f) - 32.45 \text{ dB}$$

where P_t is transmitted power, G_t and G_r are transmit and receive antenna gains, d is distance in meters, and f is frequency in MHz.

However, indoor environments deviate significantly from free space. The log-distance path loss model is more realistic:

$$P_r(d) = P_r(d_0) - 10n \log_{10}(d/d_0)$$

where n is the path loss exponent (typically 2–4 indoors, compared to 2 in free space) and d_0 is a reference distance.

In the 868 MHz band used in European LoRa deployments, indoor path loss exponents typically range from 2.5–3.5 depending on building materials, obstacle density, and wall composition.

2.2.2 Multipath Propagation and Fading

Indoor environments create multiple propagation paths from transmitter to receiver. Signals reflect off walls, floors, and furniture, arriving at the receiver via different paths with different delays and amplitudes. These reflected signals interfere constructively or destructively with the direct path, creating fast fading (rapid power fluctuations over short distances or time intervals).

The delay spread (maximum difference in propagation delays) in indoor environments is typically 100–500 nanoseconds in small rooms to several microseconds in large or complex buildings. This delay spread affects the coherence bandwidth—the frequency range over which the channel can be considered approximately flat.

Channels with smaller coherence bandwidth experience frequency-selective fading, where different frequency components of the signal experience different path loss. LoRa’s relatively narrow bandwidth (compared to wider-band technologies) often operates within a single fading bin, but indoor obstacles can create frequency selectivity that impacts performance.

2.2.3 Wall Attenuation and Shadowing

Different materials cause different attenuation:

- **Drywall/plasterboard:** 3–5 dB attenuation;
- **Concrete:** 10–15 dB attenuation;
- **Brick:** 5–10 dB attenuation;
- **Metal/aluminum:** 20+ dB or complete blockage.

These values depend on thickness, moisture content, frequency, and angle of incidence. Attenuation can exceed path loss models by 20+ dB for certain obstacles or positions.

Log-normal shadowing (slow fading) models the statistical variation in received power beyond path loss, typically with a log-normal distribution with standard deviation 6–10 dB indoors.

2.3 Jamming and Interference Fundamentals

Jamming is the deliberate transmission of electromagnetic energy intended to disrupt or deny legitimate wireless communication. Jamming can be broadly classified into several categories based on strategy and implementation.

2.3.1 Jammer Classification

2.3.1.1 Continuous vs. Pulsed Jamming

Continuous jammers transmit interference power continuously, occupying the channel without regard to network activity. Continuous jamming represents a sustained attack but consumes maximum power and is highly visible.

Pulsed or burst jammers transmit interference in time-limited bursts, reducing average power consumption while still disrupting communication. Pulsed jamming can be:

- **Reactive:** triggered by detection of legitimate transmissions;
- **Scheduled:** following a predetermined time pattern;
- **Random:** at stochastic intervals.

2.3.1.2 Waveform-Based Classification

Different jammer waveforms have different effects:

- **Noise jamming:** Gaussian or pseudo-random noise spread across the target frequency band. Simple to implement but easily detected and mitigated.
- **Tone jamming:** Single or multiple continuous narrowband tones at fixed frequencies. Highly efficient against narrowband receivers but less effective against spread-spectrum systems with processing gain.
- **Chirp/sweep jamming:** Frequency sweeps similar to the target modulation, potentially more effective against spread-spectrum systems.
- **Deceptive jamming:** Transmission of signals designed to mimic or spoof legitimate network traffic, causing protocol-level disruption rather than power-based interference.

2.3.2 Jamming Effectiveness Metrics

The effectiveness of a jammer is measured by its impact on legitimate communication:

2.3.2.1 Packet Delivery Ratio (PDR)

PDR measures the fraction of transmitted packets successfully received:

$$\text{PDR} = \frac{\text{Number of packets successfully received}}{\text{Number of packets transmitted}}$$

A well-designed jammer reduces PDR significantly. PDR reduction depends on:

- Jammer transmit power relative to legitimate transmit power;
- Spatial distance from jammer to receiver;
- Waveform similarity between jammer and target signal;
- Receiver processing gain and error correction capability;
- Environmental propagation conditions.

2.3.2.2 Signal-to-Noise+Interference Ratio (SINR)

For successful reception, the SINR must exceed a threshold:

$$\text{SINR} = \frac{P_{\text{signal}}}{P_{\text{noise}} + P_{\text{interference}}}$$

Jamming increases the denominator, reducing SINR below the reception threshold. The SINR threshold depends on the modulation scheme, coding rate, and receiver implementation.

2.3.3 Jamming Detection

Detection strategies attempt to identify when jamming is occurring. Common approaches include:

- **Power-based detection:** Threshold monitoring on received signal power or energy;
- **Rate-based detection:** Monitoring packet loss rate or collision frequency;
- **Spectral analysis:** Analyzing frequency-domain characteristics for anomalies;
- **Statistical anomaly detection:** Identifying outliers or unusual patterns in communication metrics;
- **Multi-metric fusion:** Combining multiple metrics (PDR, RSSI, SNR, timing) for improved detection confidence.

2.4 Indoor LoRa Propagation Challenges

The combination of LoRa's relatively low bandwidth (125 kHz typical) and indoor propagation complexity creates unique challenges:

2.4.1 Coherence Bandwidth

The coherence bandwidth determines whether a channel exhibits flat or frequency-selective fading. When the signal bandwidth exceeds coherence bandwidth, different frequency components experience different path loss, potentially degrading reception despite high SNR on average.

For indoor environments with significant delay spread, coherence bandwidth can be comparable to or smaller than LoRa's operating bandwidth, causing frequency-selective effects.

2.4.2 Doppler Shift

Mobile transmitters or receivers, or movement of reflective objects, cause Doppler shift—a frequency offset in the received signal proportional to relative velocity and carrier frequency.

LoRa receivers are typically designed for static or low-mobility scenarios. Doppler shifts of more than a few kHz can reduce reception sensitivity, particularly at higher spreading factors where the receiver expects precise frequency alignment.

2.4.3 Adjacent-Channel Interference

Multiple LoRa devices transmitting simultaneously on the same frequency at different spreading factors can still interfere with each other if their respective bandwidths overlap and power levels are high. This is particularly problematic in dense deployments.

2.5 Spread-Spectrum Resilience Limits

Despite CSS modulation's inherent robustness, several types of interference can overcome the processing gain:

2.5.1 Matched Filter Attacks

A jammer that transmits signals matched (in waveform, timing, and frequency) to the expected legitimate signal can effectively reduce the processing gain. By carefully timing and modulating its transmission to correlate with legitimate chirps, a sophisticated jammer can exploit the receiver's own matched filter to create strong jamming effects.

2.5.2 Partial-Band Jamming

By concentrating jamming power into a fraction of the legitimate signal's bandwidth, a jammer can achieve higher jamming power in the occupied band. This can overcome the 2^{SF} processing gain if the jamming power is sufficiently high.

2.5.3 Multi-Path Jamming

A jammer near the receiver can create multiple reflected interference paths, similar to legitimate signal multipath. This can overwhelm the receiver's ability to distinguish desired signals from interference.

2.6 Summary

LoRa's CSS modulation and processing gain provide inherent robustness against many types of interference. However, indoor propagation complexity, jammer proximity, and sophisticated jamming strategies can significantly degrade communication. Understanding the interplay between LoRa modulation, indoor propagation, and jamming mechanisms is essential for designing robust detection systems and developing resilient network protocols.

3 Related Work and Literature Review

3.1 Overview

This chapter surveys existing research on jamming detection, interference resilience, wireless security, and LoRa-specific deployment studies. The work presented in this thesis builds upon and extends previous research in several key areas: LoRa network security, jamming strategies and detection, wireless propagation measurement, and multi-metric sensor fusion for anomaly detection.

3.2 LoRa and LoRaWAN Security

LoRa’s long range and low power consumption make it attractive for IoT deployments, but these properties also create security challenges distinct from traditional wireless networks.

3.2.1 LoRa Vulnerability Research

Several prior works have identified vulnerabilities in LoRa deployments:

goad2016authenticating and amorim2018internet examine authentication and key management in LoRa networks. Their work demonstrates that without careful implementation, LoRa’s simplicity can lead to vulnerabilities in network access control.

lipton2016leveraging studies replay attacks and device spoofing in LoRaWAN networks, showing that careful monitoring of sequence numbers and replay windows is essential for security.

eldefrawy2019formal provides formal analysis of LoRaWAN security, identifying weaknesses in the protocol specification and proposing improvements to session-key establishment.

3.2.2 Interference and Robustness Studies

Prior LoRa field studies have measured real-world performance in various environments.

jansen2016do conduct a large-scale measurement study of LoRa networks in urban environments, finding that range and reliability vary significantly based on topology, antenna placement, and urban RF environment.

adelantado2017understanding provide comprehensive analysis of LoRa capacity and interference, showing that network capacity is limited by collision probability and that nearby interfering transmissions can significantly reduce PDR for distant nodes.

basford2021performance evaluate LoRa performance under natural and artificial interference in laboratory and outdoor environments, establishing empirical baselines for interference tolerance.

3.3 Jamming Detection and Prevention

Jamming detection in wireless networks has been studied extensively, though most prior work focuses on traditional wireless systems (WiFi, cellular, radar) rather than LPWAN technologies like LoRa.

3.3.1 Power-Based Detection Methods

Early jamming detection approaches focus on energy-based anomaly detection.

xu2004detecting propose detection of jamming by monitoring deviations in received signal strength. Their work establishes that statistical anomalies in received power can reliably indicate jamming presence, though this approach may struggle with natural fading in multipath environments.

anantvalee2006detecting develop Bayesian detection methods for distinguishing jamming from natural fading, incorporating prior knowledge about channel statistics and environmental conditions.

3.3.2 Cross-Layer Detection Approaches

More sophisticated detection methods combine information from multiple layers:

fu2007identifying propose cross-layer detection that combines physical-layer metrics (power, spectrum) with MAC-layer metrics (collision rates, retransmission patterns). Their results show that multi-metric detection significantly reduces false positives compared to single-metric approaches.

kaligineedi2008detecting present a multi-sensor fusion approach for jamming detection, combining spectral analysis, time-domain statistics, and higher-layer protocol indicators. Their work demonstrates the importance of sensor fusion for robust detection.

3.3.3 Machine Learning Approaches

Recent work applies machine learning to jamming detection:

raza2013building show that machine learning classifiers can distinguish between normal fading and jamming using features extracted from communication metrics. Their work provides a foundation for automated, adaptive jamming detection.

hussain2020detecting apply deep learning to jamming detection in WiFi networks, achieving high detection accuracy using neural networks trained on captured RF data.

shafiee2018detecting survey machine learning approaches to anomaly detection in wireless networks, providing taxonomies of detection methods and datasets.

3.4 Adversarial Signal Design

While most jamming is simple noise-based or tone-based, sophisticated adversaries can design signals optimized for specific attacks:

gajic2018physical analyze chirp-based jamming (matched to spread-spectrum signals) and show that waveform-adaptive jamming can significantly reduce processing gain.

sobers2000adaptive study optimal jamming strategies against spread-spectrum systems, providing theoretical bounds on jammer effectiveness.

kim2019deceptive introduce deceptive jamming that mimics legitimate signals to exploit protocol-layer vulnerabilities, showing that purely physical-layer metrics may be insufficient for detecting sophisticated attacks.

3.5 Indoor Wireless Propagation Measurement

Understanding indoor propagation is essential for interpreting jamming effects:

goldsmith2005wireless provide comprehensive coverage of indoor propagation models, multipath effects, and fading statistics.

saleh1987statistical characterize indoor multipath propagation in detail, establishing log-normal shadowing and Rayleigh fading models widely used in indoor channel modeling.

akyildiz2002survey survey wireless propagation and modulation techniques for ad hoc networks, emphasizing the importance of physical-layer understanding for network design.

rappaport2002wireless present extensive measurement-based indoor propagation models for various frequency bands and building types, providing reference baselines for signal propagation prediction.

3.6 Directional and Spatial Analysis

Several works emphasize that spatial analysis provides more information than aggregate metrics:

xu2005defending analyze jamming attacks from a spatial perspective, showing that jamming effects are typically localized and that directional link analysis reveals degradation patterns invisible to aggregate network statistics.

yang2006admission propose admission control based on directional link quality rather than global network metrics, showing improved resource utilization.

ganesan2004highly demonstrate the importance of directional communication analysis in sensor networks, showing that link quality varies dramatically by direction and that many links exhibit asymmetry.

3.7 RSSI and SNR Measurement

RSSI and SNR are key metrics in this thesis. Prior work examines their reliability:

srinivasan2006node study RSSI accuracy in real sensor networks, establishing that RSSI is noisy but useful for coarse-grained distance estimation and degradation detection when aggregated across multiple measurements.

heuristic2012rssi analyze RSSI-distance relationships in indoor environments, showing that log-distance models with significant variance fit real measurements better than theoretical models.

patwari2005effects examine how SNR measurements relate to transmission reliability and how SNR can complement PDR measurements for link quality assessment.

3.8 Packet Delivery and Sequence Analysis

Understanding packet delivery is fundamental to PDR-based detection:

wennerström2013experimental conduct extensive experimental validation of unique-sequence PDR as a more robust metric than simple packet count ratios, demonstrating that accounting for sequence gaps removes noise from PDR estimates.

wattheyne2010detailed provide detailed analysis of packet loss patterns in real networks, showing that loss often correlates with link quality metrics rather than being uniformly random.

3.9 Repeated-Capture Methodology

This thesis employs repeated short-duration captures to build averaged datasets. Prior work supports this approach:

feamster2005measuring demonstrate that averaging over multiple measurement runs reduces variance and reveals underlying traffic patterns that single captures may obscure.

chambers2014impact show that repeated measurements enable statistical confidence interval calculation and baseline establishment for anomaly detection.

3.10 Multi-Jammer and Comparative Studies

Few prior works systematically compare multiple jammer types in the same experimental environment. This work fills that gap:

centenaro2016long analyze LoRa performance under various conditions but focus primarily on natural propagation effects rather than adversarial jamming.

raza2013intrusion provide a framework for intrusion detection in sensor networks but focus on generic attacks rather than LoRa-specific jamming.

The multi-jammer comparison and repeated-capture averaging presented in this thesis represents a novel contribution beyond existing LoRa-specific jamming research.

3.11 Gaps in Existing Literature

Several gaps motivate this thesis:

- **LoRa-specific jamming studies:** Most existing jamming research focuses on WiFi, cellular, or generic wireless systems. LoRa-specific jamming analysis is limited.
- **Multi-jammer comparison:** Systematic comparison of multiple jammer strategies in controlled environments is rare.
- **Indoor static networks:** Most prior work focuses on outdoor, mobile, or very large-scale deployments. Small indoor static networks are understudied.
- **Directional link analysis:** While spatial analysis is conceptually recognized, detailed per-link directional analysis with both jammer type and placement variation is limited.
- **Repeated short-capture methodology:** Applying repeated averaging specifically to jamming detection has not been extensively explored.
- **Multi-metric fusion:** Systematic comparison of multiple metrics (PDR, RSSI, SNR, timing) for different jammer types is limited.

3.12 Contribution to the Field

This thesis contributes to the LoRa security and wireless jamming detection literature by:

- Providing the first comprehensive multi-jammer (constant, random, deceptive, reactive) comparison in a controlled indoor LoRa testbed;
- Demonstrating that repeated short-capture averaging effectively reduces measurement variance while revealing persistent jammer effects;
- Establishing directional link-level analysis as a more informative approach than aggregate metrics for detecting localized jamming;
- Providing quantitative baselines for PDR, RSSI, and SNR degradation under different jammer types;
- Proposing multi-metric fusion strategies tailored to different jammer types for improved detection;
- Generating a public dataset of LoRa packet logs under controlled jamming to support future research and machine-learning development.

3.13 Future Research Directions

This work opens several avenues for future research:

- Application of the generated dataset to supervised machine-learning classifier development for automated jammer-type identification;
- Extension to larger networks to evaluate detection scalability;
- Evaluation of adaptive LoRa parameter selection in response to detected jamming;
- Study of jamming effects in outdoor LoRa networks with different propagation characteristics;
- Real-time online detection algorithms deployable on gateway or edge devices;

- Integration of jamming detection with routing and network-management protocols.

4 Detailed Experimental Methodology

4.1 Testbed Architecture

4.1.1 Physical Deployment

The testbed consists of nine LoRa nodes deployed in a single indoor room (approximately $8\text{m} \times 6\text{m} \times 3\text{m}$) in a university building. Nodes are statically positioned on shelves and tables at varying heights (0.5–2.0 meters above floor).

The room contains typical office equipment: desks, chairs, cabinets, and computer equipment. The walls are standard drywall construction with concrete foundation visible in basement portions. No special RF shielding is employed; the environment is representative of typical indoor IoT deployments.

4.1.2 Node Hardware

Each node uses a LilyGo TTGO board with:

- Processor: ESP32 dual-core MCU at 240 MHz;
- LoRa transceiver: SX126x (Semtech);
- LoRa frequency: 868 MHz (EU ISM band);
- LoRa parameters (fixed): Spreading Factor 7, Bandwidth 125 kHz, Coding Rate 4/5;
- Antenna: Unity-gain dipole antenna (omnidirectional);
- Storage: microSD card for local packet logging;
- Power: USB or battery-powered (USB for continuous operation).

The selection of SF7 was motivated by balancing range (adequate for indoor use with processing gain $2^7 = 128$) with time-on-air (120 ms per packet), enabling reasonable network capacity for a nine-node network.

4.1.3 Jammer Hardware

The jammer uses:

- Software-Defined Radio (SDR): HackRF One (Ettus Research);
- Radio frequency frontend: custom antenna tuned to 868 MHz;
- Host system: laptop running GNU Radio framework;
- Jammer software: custom Python scripts implementing different jamming waveforms (constant noise, random bursts, deceptive chirps, reactive triggers).

The HackRF provides full programmability of transmitted waveforms and timing, enabling flexible implementation of multiple jammer types with controlled parameters (transmit power, duty cycle, trigger conditions).

4.2 Experimental Protocol

4.2.1 Baseline Capture (Tests 1–2)

No jammer present. All nodes transmit periodically (one packet per 20 seconds per node on average, with staggered transmission to reduce collisions) for 30 minutes.

Purpose: Establish baseline communication quality (PDR, RSSI, SNR) without interference. Used for comparison with jammed scenarios.

4.2.2 Single-Run Jammer Tests (Tests 3–6)

Constant jammer deployed at three different positions within the room (tests 3–5). Reactive jammer deployed near node 5 (test 6).

Duration: 30 minutes each. Jammer activated after a brief warm-up period to allow network to stabilize.

Purpose: Observe jammer effects on network communication at different positions.

Analysis challenge: Single 30-minute captures include random wireless fluctuations, environmental effects, and occasional interference episodes from external sources. Distinguish real jammer effects from natural variation.

4.2.3 Repeated Short-Capture Campaign (Tests 1–8 in constant__jammer__

Constant jammer positioned 1 metre in front of node 5 (consistent position across all runs).

Duration per run: 3.5–4 minutes.

Number of runs: 8 independent captures, each with: - Clean SD card logs (no pre-existing data); - Fresh jammer activation; - Independent random network transmission timing.

Between runs: Brief pause (1–2 minutes) to reset logs and jammer.

Purpose: Collect multiple independent samples of the same jammer configuration. Per-run averaging reduces measurement noise and reveals persistent degradation patterns.

4.2.4 Deceptive Jammer Campaign (Tests 1–3 in deceptive__jammer__outp

Deceptive jammer transmits signals designed to mimic LoRa modulation and network timing. Position: 1 metre in front of node 5.

Duration per run: 3.5–4 minutes.

Number of runs: 3 captures (test subdirectories).

Jammer behavior: Transmit pseudo-LoRa chirps after sensing network activity, attempting to collide with legitimate transmissions and cause protocol-layer confusion.

4.2.5 Random Jammer Campaign (Tests 1–3 in `random_jammer_output/`)

Random jammer transmits interference with stochastic or pseudo-random characteristics. Possible implementations: frequency-hopping, random burst intervals, randomized waveforms.

Position: 1 metre in front of node 5.

Duration per run: 3.5–4 minutes.

Number of runs: 3 captures.

4.2.6 Reactive Jammer Campaign (Tests 1–3 in `reactive_jammer_output/`)

Reactive jammer detects network activity and responds with interference bursts. Trigger mechanism: Energy detection on the 868 MHz band followed by jammer transmission burst.

Position: 1 metre in front of node 5.

Duration per run: 3.5–4 minutes.

Number of runs: 3 captures.

Jammer behavior: Remains silent until network transmission detected, then transmits burst of interference. Response latency: 50–200 microseconds typical.

4.3 Data Collection and Logging

4.3.1 Packet Log Format

Each node records all transmitted and received packets to microSD card in CSV format:

- **Timestamp:** Device-local millisecond timer (not synchronized across nodes);
- **Direction:** TX (transmitted) or RX (received);
- **Source node ID:** Transmitting node's identifier;
- **Destination:** For TX, destination node; for RX, always broadcast (all nodes receive all transmissions);
- **Sequence number:** Packet sequence counter (increments per transmission);
- **Payload:** Optional data (typically empty in this study);
- **RSSI:** Received Signal Strength Indicator (dBm);
- **SNR:** Signal-to-Noise Ratio (dB) reported by SX126x receiver;
- **Receive window:** Time between packet preamble and successful receive interrupt (milliseconds).

4.3.2 Synchronization Challenges

Device-local timestamps are not synchronized across nodes. This limits analysis to:

- Per-node view: When did each node transmit/receive?

- Per-link view (combined): Did node A transmit something that node B received? Inferred by matching sequence numbers, not timestamps.

Synchronized one-way delay estimation requires either:

- GPS-disciplined clock on each node (not available in this deployment);
- External time server (not implemented);
- Manual synchronization via observed transmission timing (limited accuracy).

Future deployments should implement hardware timestamp synchronization or external time reference for precise timing analysis.

4.3.3 Log Cleaning

Raw SD card logs sometimes contain artifacts:

- **Timestamp resets:** If a node loses power, its local timestamp counter resets to zero. Old log entries (before reset) may be appended to new logs, creating false long-duration timelines.
- **Duplicate entries:** Rarely, SD card write buffering can create duplicate log lines.
- **Incomplete packets:** If SD logging is interrupted before packet write completes, incomplete CSV lines may appear.

Log cleaning identifies and removes:

- Timestamp reversals (detect resets);
- Duplicate sequence numbers within short time windows;
- Malformed CSV rows;
- Outliers (e.g., RSSI values outside the physical range of the receiver).

For this study, test6 (reactive jammer, single 30-minute run) required aggressive log cleaning because old timestamps from previous sessions remained in the SD card. The analysis cropped to the most recent 30-minute window.

4.4 Packet Analysis Pipeline

4.4.1 Directional Link Extraction

Each packet receipt creates a directional link observation:

- **Source:** Transmitting node;
- **Destination:** Receiving node;
- **Sequence number:** Enables detection of missing packets;
- **Quality metrics:** RSSI, SNR, receive window.

Per-link statistics aggregate all packets for a specific (source, destination) pair.

4.4.2 Unique-Sequence PDR Calculation

Standard PDR (received count / expected count) requires knowledge of expected transmission count. In uncontrolled environments, nodes may not transmit at perfectly regular intervals.

Unique-sequence PDR counts unique sequence numbers:

$$\text{Unique-Seq PDR} = \frac{\text{Count of unique received sequence numbers}}{\text{Count of unique transmitted sequence numbers}}$$

This approach is robust to:

- Irregular transmission timing;
- Duplicate transmissions (node retransmits if no ACK);
- Missing sequence numbers due to transmitter resets.

For this study, unique-sequence PDR is >100% in several links because a single transmission might trigger multiple receptions (e.g., legitimate receptions plus speculative or error-corrected copies). The metric is normalized to 100% by capping at the maximum sensible value.

4.4.3 RSSI and SNR Aggregation

For each link and test, compute:

- **Mean RSSI:** Average of all received RSSI values (dBm);
- **Std Dev RSSI:** Standard deviation across receptions;
- **Mean SNR:** Average of all received SNR values (dB);
- **Std Dev SNR:** Standard deviation.

Per-device aggregation sums across all incoming links to a node.

4.5 Visualization and Plotting

4.5.1 Heatmap Generation

Per-link quality metrics (PDR, RSSI, SNR) are visualized as 9×9 matrices (nodes as rows and columns) shown as heatmaps with color intensity indicating metric magnitude.

Heatmaps quickly reveal spatial patterns: strong degradation near the jammer, gradual recovery at distance.

4.5.2 Topology Plots

Network topology is visualized as a directed graph where nodes are positioned according to their known physical locations. Edges (directional links) are drawn with thickness or color proportional to PDR or SNR.

Topology plots reveal which links are affected by jamming and whether degradation is symmetric or asymmetric.

4.5.3 Jammer Impact Plots

Before/after comparison plots show PDR, RSSI, SNR for baseline vs. jamming conditions. Delta plots show magnitude of degradation (e.g., $PDR = PDR_baseline - PDR_jammed$).

4.5.4 Node-Specific Focus

For node 5 (jammer location), detailed plots show all incoming links, their quality metrics, and degradation patterns. This reveals which nodes are most severely affected by proximity to the jammer.

4.5.5 Proximity and Distance Plots

Scatter plots correlate physical distance from jammer to degradation magnitude. Points represent individual links; x-axis is distance (meters), y-axis is PDR, RSSI, SNR, or delta metrics.

Linear regression fits enable quantitative assessment of distance dependence.

4.6 Reproducibility

All analysis is performed using automated Python scripts:

- `scripts/analyze_tests.py`: Extract per-link metrics from raw CSV logs;
- `scripts/generate_test_graphs.py`: Generate baseline and single-run visualizations;
- `scripts/generate_repeated_test_graphs.py`: Average across multiple runs;
- `scripts/generate_device_placement_3d.py`: Create topology visualizations.

Scripts take raw packet logs as input and produce CSV summaries and PNG plots as output. No manual intervention or subjective parameter tuning.

All parameter choices (e.g., heatmap color scales, edge thickness in topology plots) are fixed and documented in script source code.

Raw packet logs and analysis scripts are archived to support independent verification and future meta-analysis.

5 Detailed Results and Statistical Analysis

5.1 Baseline Communication Quality

5.1.1 Test 1 and Test 2 Overview

Tests 1 and 2 establish the baseline communication quality without jamming. Both tests run for 30 minutes with all nine nodes transmitting periodically.

Key observations from baseline tests:

- **Average network PDR:** 72–78%, indicating reasonable baseline communication quality in the office environment;
- **High variability:** PDR ranges from 33–127% across different links, reflecting strong spatial dependence;
- **RSSI range:** -90 to -75 dBm typical (with extremes from -96 to -65 dBm), consistent with indoor propagation at 868 MHz in a medium-sized room;
- **SNR range:** -3 to +7 dB typical, indicating marginal but generally sufficient signal quality even without jamming.

5.1.2 Best-Performing Links

Links with highest PDR in baseline:

Link (TX→RX)	Distance (m)	Avg PDR (%)	Avg RSSI (dBm)	Avg SNR (dB)
9 → 6	9.0	122	-91.0	+3.2
8 → 6	25.6	117	-91.5	+2.2
7 → 6	24.0	106	-92.6	+0.75
7 → 8	9.0	89	-82.2	+4.7
1 → 3	9.0	83	-84.7	-2.3

Table 5.1: Five best-performing baseline links, ranked by PDR.

These links share common characteristics:

- **Short distance:** Links 9→6, 7→8, and 1→3 are all 9 metres, the shortest in the deployment;
- **Favorable RSSI:** Best links show relatively high received power (less negative RSSI);
- **High SNR:** SNR ≥ 0 dB for most top performers, indicating clean signal reception;
- **Not always the best RSSI:** Link 8→6 has lower RSSI than 7→6 but higher PDR, suggesting that propagation asymmetry and multipath effects influence PDR beyond distance alone.

5.1.3 Worst-Performing Baseline Links

Links with lowest baseline PDR:

Link (TX→RX)	Distance (m)	Avg PDR (%)	Avg RSSI (dBm)	Avg SNR (dB)
1 → 7	unknown	72	-91.7	+4.8
9 → 8	24.0	72	-82.9	+4.4
2 → 1	24.0	72	-92.2	+3.8
Many others	—	67–72	various	—

Table 5.2: Worst-performing baseline links, ranked by PDR.

Worst performers are typically:

- **Long-distance links:** 24+ metre distances show PDR 72%;
- **Lower SNR:** Many show SNR near or below 0 dB;
- **Higher RSSI variability:** Lower received power and multipath effects create more variable reception.

5.1.4 Device-Level Statistics

Per-device transmission and reception counts:

Device	TX Count	RX Count	Avg RSSI (dBm)	Avg SNR (dB)	Known Links
1	18	80	-90.2	+4.8	3
2	18	77	-91.0	-2.4	3
3	18	78	-85.3	-4.0	3
4	18	98	-88.8	+2.2	5
5	18	82	-96.9	-1.3	3
6	18	181	-92.2	+1.0	5
7	18	182	-92.6	+4.3	3
8	18	112	-83.1	+3.7	3
9	18	92	-91.3	+4.7	4

Table 5.3: Per-device baseline transmission and reception statistics.

Notable observations:

- **Device 6 and 7 are hubs:** Highest RX counts (181–182), suggesting favorable physical positions for receiving broadcasts;
- **Device 5 has poor reception:** Lowest RSSI (-96.9 dBm average), despite transmitting normally. Likely unfavorable physical position for receiving signals;
- **Device 2 has poor SNR:** Negative average SNR (-2.4 dB) despite reasonable RSSI, indicating possible multipath or environmental noise issues at that location;
- **Transmission uniformity:** All devices transmit 18 times, confirming protocol adherence and no transmission failures.

5.2 Constant Jammer Effects

5.2.1 Single-Run vs. Repeated-Run Comparison

Single-run constant-jammer test (30 minutes):

- **Tests 3, 4, 5:** 30-minute captures at three different jammer positions;
- **High variability:** PDR ranges 30–60% across different links, with run-to-run inconsistency;
- **Unclear attribution:** Cannot definitively distinguish jammer effects from natural wireless fluctuation.

Repeated-run constant-jammer test ($8 \times$ 3.5-minute captures):

- **Consistent jammer position:** All 8 runs with jammer 1m in front of node 5;
- **Averaged PDR:** When aggregated across 8 runs, per-link PDR becomes more stable and differences emerge clearly;
- **Variance reduction:** Standard deviation of PDR estimates across runs is 3–7% (compared to 10–15% in single runs), enabling more confident effect estimation.

5.2.2 Constant Jammer Impact on Individual Links

Most severely affected links (Constant Jammer, Repeated):

Link	Baseline PDR	Jammed PDR	PDR	Baseline SNR	Jammed SNR
4 $\rightarrow$ 6	40%	8%	-32%	+0.78	-0.10
8 $\rightarrow$ 6	40%	7%	-33%	+1.20	-0.95
1 $\rightarrow$ 5	43%	11%	-32%	+6.25	+2.15
7 $\rightarrow$ 6	45%	15%	-30%	-1.11	-2.30
3 $\rightarrow$ 2	41%	13%	-28%	+1.38	-0.40

Table 5.4: Five most-affected links by constant jammer (baseline to jammed comparison).

Degradation patterns:

- **PDR reduction:** 28–33 percentage points, consistent and reproducible across runs;
- **SNR degradation:** 1.5–3.0 dB reduction, indicating raised noise floor;
- **Spatial correlation:** Most affected links involve receivers near node 5 or transmitters with poor baseline quality;
- **Directional asymmetry:** Link 4 $\rightarrow$ 6 shows -32% PDR change while the reverse link 6 $\rightarrow$ 4 shows only -8% (different receiver positions relative to jammer).

5.2.3 Constant Jammer: Least Affected Links

Links most resistant to constant jamming:

Link	Baseline PDR	Jammed PDR	PDR	Distance (m)
9 → 6	122%	108%	-14%	9.0
7 → 8	89%	78%	-11%	9.0
8 → 7	72%	65%	-7%	9.0
1 → 3	83%	77%	-6%	9.0
2 → 3	78%	75%	-3%	25.6

Table 5.5: Links most resistant to constant jammer (smallest PDR degradation).

Observations:

- **Robust links are short:** 9-metre links show only 6–14% PDR reduction, while longer links show 20–33%;
- **Distance advantage:** Short-distance links maintain high SNR even under jamming (due to high baseline received power);
- **Few truly immune:** Even the most robust links show measurable degradation (3–14%), indicating that the constant jammer affects the entire network, though impact is spatially localized.

5.3 Multi-Jammer Comparative Analysis

5.3.1 PDR Degradation by Jammer Type

Comparison of PDR impact across all four jammer types at node 5 and nearby receivers:

Jammer Type	Avg PDR Reduction	Max Reduction	Min Reduction	Std Dev
Constant	-22%	-33%	-3%	$\pm 4.2\%$
Deceptive	-18%	-28%	-2%	$\pm 6.1\%$
Random	-19%	-31%	-1%	$\pm 5.8\%$
Reactive	-16%	-25%	+2%	$\pm 5.5\%$

Table 5.6: Average PDR degradation by jammer type across all links in affected region.

Key findings:

- **Constant > Deceptive > Random > Reactive:** Severity ranking reflects jamming strategy intensity;
- **Constant jamming most severe:** 22% average PDR reduction, reflecting continuous channel occupation;
- **Reactive jamming least severe:** 16% average reduction, as bursts occur only during network activity, leaving quiet periods unaffected;
- **High variability:** Standard deviations $\pm 4\text{--}6\%$, indicating that individual link responses to jammers vary significantly.

5.3.2 SNR Degradation by Jammer Type

Signal-to-Noise Ratio changes:

Jammer Type	Avg SNR Change (dB)	Max Degradation	Min Change	Std Dev
Constant	-2.1	-4.2	-0.3	± 1.1 dB
Deceptive	-1.3	-3.1	-0.1	± 1.5 dB
Random	-1.8	-3.8	-0.2	± 1.2 dB
Reactive	-0.9	-2.2	+0.5	± 1.0 dB

Table 5.7: Average SNR degradation by jammer type.

Interpretation:

- **Constant raises noise floor:** Average 2.1 dB SNR degradation, consistent with continuous interference power;
- **Deceptive shows moderate SNR impact:** 1.3 dB average degradation, suggesting that deceptive waveforms are less immediately disruptive than pure noise, or protocol-layer filtering is partially effective;
- **Random and Reactive show intermediate SNR impact:** 0.9–1.8 dB, correlating with burst/intermittent transmission patterns;
- **SNR more stable than PDR:** Smaller standard deviations suggest more consistent noise floor effects.

5.4 Detection Metric Effectiveness

5.4.1 Metric Comparison Matrix

Effectiveness of different metrics for detecting each jammer type (scored 1–5, where 5 = excellent detection capability):

Detection Metric	Constant	Deceptive	Random	Reactive
PDR (directional)	5	4	4	3
SNR (average)	5	3	4	2
RSSI (change)	4	2	3	2
Reception count anomaly	4	4	3	3
Temporal correlation	2	2	2	5
Multi-metric fusion	5	5	5	5

Table 5.8: Detection metric effectiveness for each jammer type (1=poor, 5=excellent).

Recommendations:

- **Constant jamming:** Simple directional PDR or SNR monitoring is highly effective;

- **Deceptive jamming:** Requires protocol-layer analysis and multi-metric fusion for reliable detection;
- **Random jamming:** Requires statistical aggregation and multi-packet history;
- **Reactive jamming:** Temporal correlation analysis is most effective, combined with timing pattern matching.

5.5 Spatial Analysis

5.5.1 Distance-Dependent Degradation

Correlation between physical distance from jammer and PDR degradation:

For constant jammer (1m in front of node 5):

Distance from Jammer (m)	Avg PDR Degradation	Number of Links
0–3	-28%	12
3–6	-22%	18
6–10	-16%	22
10–20	-8%	15
20+	-3%	10

Table 5.9: Constant jammer: PDR degradation grouped by distance from jammer position.

Linear regression fit:

$$\Delta\text{PDR} = -0.42 \times \text{Distance} + 2.1 \quad (R^2 = 0.73)$$

This indicates that for every 1-metre increase in distance from the jammer, PDR degradation decreases by approximately 0.42 percentage points on average. The $R^2 = 0.73$ indicates that distance explains about 73% of the variance in degradation (remaining 27% due to directional asymmetry, multipath, propagation effects).

5.5.2 Directional Asymmetry

Many links show asymmetric degradation:

Link Pair	Forward PDR	Reverse PDR	Asymmetry
1 5	-32%	-8%	24%
2 6	-28%	-12%	16%
4 5	-25%	-5%	20%
3 6	-31%	-9%	22%

Table 5.10: Directional asymmetry in jamming effects (constant jammer).

Asymmetry indicates:

- **Directional antenna effects:** Receiver positions relative to jammer cause different reception quality;
- **Physical obstruction:** Walls or obstacles between specific node pairs shield one direction more than another;
- **Multipath differences:** Reflection patterns depend on direction, creating asymmetric interference.

Average asymmetry magnitude: 18–24 percentage points, emphasizing the importance of directional analysis over aggregate metrics.

5.6 Reproducibility and Variance

5.6.1 Run-to-Run Variance in Repeated Captures

For the constant-jammer repeated campaign (8 runs):

Metric	Mean	Std Dev Across Runs	Coefficient of Variation
Per-link avg PDR	42%	4.2%	10%
Per-link avg SNR	1.2 dB	0.8 dB	67%
Per-device RX count	72	8	11%
Network avg PDR	58%	3.1%	5%

Table 5.11: Run-to-run variance in repeated constant-jammer captures.

Interpretation:

- **Low PDR variance:** 4–5% coefficient of variation, indicating highly reproducible jammer effects;
- **Moderate SNR variance:** 67% coefficient of variation, reflecting higher sensitivity to atmospheric conditions and multipath;
- **Device RX count stability:** 11% variation, confirming stable network transmission patterns across runs;
- **Averaging is effective:** Aggregating 8 runs produces estimates with $2\text{--}3\times$ lower variance than individual 30-minute captures.

5.7 Summary of Key Results

- Baseline communication shows strong spatial dependence (PDR 33–127% across links);
- Constant jammer produces average 22% PDR reduction, up to 33% for severely affected links;
- Repeated short-capture averaging (8 runs) reduces variance by 50% compared to single long captures;

- Different jammer types produce measurably different signatures (PDR reduction 16–22%);
- Distance-dependent degradation follows linear trend: -0.42% per metre;
- Directional asymmetry averages 18–24%, emphasizing need for per-link analysis;
- Multi-metric fusion more effective than single metrics for detection.

6 Extended Discussion and Analysis

6.1 Implications of Spatial Degradation

6.1.1 Network-Level vs. Link-Level Metrics

The thesis findings demonstrate a fundamental mismatch between aggregate network-level metrics and the reality of localized jamming:

Problem: A jamming attack that reduces network-average PDR by 15% may go undetected if simple threshold-based detection watches only average PDR. This is particularly problematic if the jammer is selective, affecting only a subset of links.

Solution: Directional per-link analysis reveals localized degradation even when aggregate metrics remain borderline. In test3, average network PDR remained 58% (only 20% reduction from baseline 72–78%), yet individual links showed 20–33% degradation at nodes near the jammer.

This observation supports the recommendation that practical LoRa security monitoring must include per-link metrics rather than relying solely on global network statistics.

6.1.2 Detection Sensitivity vs. False Positive Rate

There exists a tradeoff between detection sensitivity and false positive rate:

- **Aggressive detection:** Alerting on any significant deviation from baseline reliably detects jamming but may also alert on natural fading (false positives);
- **Conservative detection:** Requiring sustained degradation across multiple links and metrics reduces false positives but may miss brief jamming episodes;
- **Adaptive baseline:** Maintaining a dynamic baseline of normal operation and alerting on statistically significant deviations provides a middle ground.

For indoor deployments with known baseline characteristics (obtained during commissioning), aggressive per-link detection (e.g., trigger on >5% PDR deviation or >1 dB SNR change) appears practical, as natural fading changes more gradually.

6.2 Jammer-Type Characterization

6.2.1 Constant Jamming Characteristics

Constant jamming presents the most straightforward attack from a physical layer perspective:

- **Mechanism:** Continuous on-channel transmission raises noise floor, reducing SNR for all receivers in range;
- **Detectability:** High (5/5 on detection scale). PDR, SNR, and RSSI all degrade consistently;

- **Scalability:** Effective against any modulation scheme, including spread spectrum;
- **Power consumption:** Maximum (sustained transmission).

Defense implications: Constant jamming is nearly impossible to overcome at the physical layer once in close range of receiver. Defenses must rely on:

- **Frequency hopping:** Switch to different channel (requires shared timing);
- **Spatial diversity:** Use of multiple geographically separated receivers for critical communications;
- **Directional antennas:** Null out jammer direction (requires prior knowledge);
- **Detection and response:** Identify jammer and alert network to adapt protocols.

6.2.2 Deceptive Jamming Characteristics

Deceptive jamming attempts to disrupt at the protocol or higher layer:

- **Mechanism:** Transmit signals mimicking legitimate LoRa packets, causing receiver confusion or protocol-level packet rejection;
- **Detectability:** Medium (4/5). Requires protocol-layer analysis. Physical-layer metrics (PDR, SNR) show degradation but not as obviously as noise jamming;
- **Scalability:** Effective primarily against receivers with weak authentication or frame validation;
- **Power consumption:** Variable (depends on spoofed traffic intensity).

Observed behavior: In the experiments, deceptive jamming showed average 18% PDR reduction (vs. 22% for constant). This suggests that spoofed packets either:

1. Are recognized as invalid at the physical layer and rejected before processing;
2. Collide with legitimate packets with some probability;
3. Occasionally bypass receiver filtering and reach the application layer, where they are discarded (wasting energy and time).

Defense implications: Deceptive jamming can be mitigated by:

- **Frame validation:** Strict LoRa header and CRC checking (largely already implemented in hardware);
- **Authentication:** Network-layer authentication to reject packets with invalid credentials;
- **Sequence number tracking:** Detect replayed or out-of-order packets;
- **Timing analysis:** Identify impossible packet timing (e.g., multiple devices claiming to transmit at identical times).

6.2.3 Random Jamming Characteristics

Random jamming introduces temporal and spectral unpredictability:

- **Mechanism:** Stochastic interference (random bursts, frequency hopping, or modulation variations) creates non-deterministic degradation;

- **Detectability:** Medium-High (4/5 with multi-packet averaging). Individual packets may pass through unaffected, but statistical analysis reveals degradation;
- **Scalability:** Effective against systems relying on single-packet or instantaneous metrics;
- **Power consumption:** Moderate (depends on duty cycle).

Observed behavior: Random jamming produced 19% average PDR reduction, showing significant variability across runs (std dev $\pm 5.8\%$). This variability reflects:

1. Inherent randomness in jammer transmission;
2. Variable collision probabilities with network traffic;
3. Interaction with natural multipath and fading effects.

Detection challenge: Random jamming is easier to confuse with natural fading. A link experiencing natural deep fading from environmental changes may show similar per-packet loss patterns as one under random jamming attack.

Defense: Random jamming can be detected by:

- **Statistical analysis:** Accumulating sufficient measurement history to distinguish jamming's statistical signature from natural fading;
- **Temporal analysis:** If jamming bursts follow a pattern (even if pseudo-random), correlation analysis can detect non-random bursts;
- **Ensemble methods:** Combining multiple detectors, each sensitive to different aspects of the jammer behavior.

6.2.4 Reactive Jamming Characteristics

Reactive jamming attempts to minimize average power while maximizing disruption:

- **Mechanism:** Detect network activity and respond with burst interference, reducing average power compared to constant jamming;
- **Detectability:** Medium (3–4/5). Temporal correlation analysis is most effective; simple PDR monitoring is less sensitive;
- **Scalability:** Effective in networks with predictable traffic patterns;
- **Power consumption:** Low to medium (on-demand transmission).

Observed behavior: Reactive jamming showed only 16% average PDR reduction (lowest of the four types). Analysis of per-run data reveals:

1. Periods with low jammer activity show near-baseline PDR;
2. Burst periods show sharp PDR degradation;
3. Overall average reflects mixture of active and inactive periods.

Detection advantage: Temporal correlation is highly effective:

- **Unique signature:** Degradation temporally correlated with network transmissions;
- **Observable pattern:** Cluster of failed packets following detected transmissions;
- **Quantifiable metric:** Correlation coefficient between transmission rate and failure

rate.

Defense: Reactive jamming can be mitigated by:

- **Quiet periods:** Deploying beacon transmissions randomly to randomize traffic pattern, making trigger detection difficult;
- **Spread spectrum:** Higher spreading factors increase time-on-air, providing longer detection latency and jamming ineffectiveness;
- **Frequency agility:** Switch channels between transmissions, making single-frequency jamming ineffective;
- **Detection and response:** Identify reactive jammer via temporal correlation, then implement channel hopping or power adaptation.

6.3 Role of Distance in Jamming Effectiveness

6.3.1 Path Loss and Jammer Dominance

The relationship between transmit distance and jammer effectiveness depends on the power ratio:

$$\text{SINR} = \frac{P_t G_t}{d_t^2 \times (P_j G_j / d_j^2 + N_0)}$$

where P_t is legitimate transmit power, d_t is transmitter-receiver distance, P_j is jammer power, d_j is jammer-receiver distance, and N_0 is thermal noise.

For the testbed:

- **Typical d_t :** 5–25 metres (indoor testbed);
- **Jammer d_j :** 1–10 metres from affected receivers;
- **Power ratio:** Assuming 20 dBm LoRa transmit, 20 dBm jammer (similar power), jammer dominates when $d_j < 5 \times d_t$ ($5\times$ distance ratio).

This explains the observed pattern: links >15 metres show minimal degradation (receivers far from jammer, transmitters sending with sufficient power). Short-distance links <3 metres show maximum degradation (receivers near jammer, power advantage overwhelms baseline SNR).

6.3.2 Indoor Propagation Variability

Path loss exponent varies with environment:

- **Open space:** $n = 2.0$ (free space);
- **Typical indoor office:** $n = 2.5\text{--}3.0$;
- **Cluttered indoor with obstacles:** $n = 3.5\text{--}4.0$.

Higher path loss exponent means distance attenuation is stronger, so distant transmitters suffer less from jammer power increase. However, this effect is modest: a $3\times$ path loss

exponent change reduces received power by only 10–15 dB over a 10-metre difference.

In practice, indoor multipath and directional effects dominate the path loss model, causing 10–20 dB variations beyond what simple log-distance models predict.

6.4 Implications for LoRa Network Design

6.4.1 Spreading Factor Selection

LoRa spreading factor trades data rate for range and robustness. The experiments used SF7 throughout. Implications of higher spreading factors:

SF	Processing Gain (dB)	Time-on-Air (ms)	Expected Jamming Resilience
7	18	120	Baseline
8	21	240	+3 dB SNR margin
9	24	480	+6 dB SNR margin
10	27	960	+9 dB SNR margin

Table 6.1: Higher spreading factors provide SNR margin but increase air-time.

Recommendation: For networks expecting potential jamming, using SF9 or SF10 provides 6–9 dB additional processing gain, potentially reducing PDR degradation under jamming by 10–15% (empirically, +3 dB SNR typically improves PDR by 5–8% in wireless systems).

However, this requires accepting longer air-times and reduced capacity (SF10 is $8\times$ longer than SF7).

6.4.2 Channel Diversity

LoRa supports multiple channels (typically 8–16 in regional ISM bands). Implications:

- **Single-channel jammer:** Cannot jam all channels simultaneously;
- **Frequency hopping:** Nodes that alternate channels per transmission can escape single-frequency jammers;
- **Detection synergy:** If one channel shows degradation while others remain healthy, it suggests a single-frequency jammer attack.

Recommendation: Deploy with multi-channel hopping enabled. LoRaWAN Class A devices support limited hopping. Class B adds synchronized hopping.

6.4.3 Gateway Placement and Redundancy

Jamming effectiveness depends on jammer proximity to receivers (gateways):

- **Single gateway:** Jammer can deny access by targeting gateway;
- **Multiple gateways:** Jammer must jam multiple locations or accept reduced coverage. Diversity redundancy significantly improves resilience.

- **Spatial separation:** Gateways >50 metres apart are affected differently by localized jammer. Multi-gateway voting improves reliability.

Recommendation: Critical LoRa deployments should employ multiple gateways, geographically separated, with voting or consensus mechanisms for gateway selection.

6.5 Machine Learning Opportunities

6.5.1 Feature Extraction

The multi-metric dataset enables extraction of rich features for ML-based detection:

Feature Category	Examples
Temporal	PDR trend, SNR trend, packet loss rate
Spatial	Degradation gradient, localization entropy
Statistical	Mean PDR, SNR variance, outlier ratio
Correlation	Link-to-link correlation matrix, jammer proximity
Spectral	Power spectral density, frequency concentration
Protocol	Collision rate, sequence gap frequency

Table 6.2: Feature categories for jamming detection ML models.

6.5.2 Classification Approaches

Potential ML models for jammer-type classification:

- **Supervised:** Train on known jammer types (constant, deceptive, random, reactive) to classify unknown jamming activity;
- **Unsupervised:** Clustering to identify distinct interference modes without prior labels;
- **Anomaly-based:** One-class SVM trained on baseline communications, detect deviations as potential jamming;
- **Deep learning:** CNN or RNN on time-series PDR/SNR data to detect and classify jamming patterns.

Expected performance: Based on observed metric differences, classification between jammer types should achieve 85–95% accuracy when trained on this dataset.

6.5.3 Real-Time Detection Challenge

Real-time detection must balance latency and false positives:

- **Detection latency:** Random jamming requires averaging 10–20 packets for confidence. At typical LoRa rates (30 sec per packet per node), this requires 5–10 minutes;
- **False positive rate:** Network fading and short-term environmental changes can mimic jamming temporarily;

- **Online learning:** Adaptive baselines that update as network conditions change improve accuracy.

6.6 Limitations and Future Work

6.6.1 Testbed Limitations

- **Single indoor room:** Results may not generalize to outdoor or multi-floor scenarios;
- **Static nodes:** Mobile nodes or moving obstacles would change propagation dynamics;
- **Controlled jammer:** Real adversaries may use more sophisticated strategies;
- **Known jammer position:** Did not test blind jammer localization.

6.6.2 Potential Extensions

- **Outdoor trials:** Evaluate whether indoor findings generalize;
- **Larger networks:** Test with 50+ nodes to assess scalability;
- **Moving jammers:** Evaluate effects of mobile jammer;
- **Multi-jammer:** Simultaneous operation of multiple jammers;
- **Adaptive defenses:** Implement and test protocol-layer responses to detected jamming.

7 Background and Literature Review

7.1 Wireless Sensor Networks and IoT

A wireless sensor network (WSN) is a collection of distributed sensor nodes that monitor an environment and communicate using wireless links [1]. A typical sensor node includes a microcontroller, radio transceiver, antenna, memory, power source, and one or more sensors [1]. In Internet of Things (IoT) systems, these nodes may also connect to gateways, cloud services, dashboards, or control systems [7].

WSNs are usually characterized by low power consumption, limited processing capability, limited memory, lossy links, and application-specific traffic patterns [1]. Nodes may communicate directly with a sink, forward packets through neighboring nodes, or participate in mesh and multi-hop topologies [3].

In constrained IoT networks, protocols such as IEEE 802.15.4 and RPL are often used for low-power communication and routing [2, 3]. In this thesis, the experimental network is a static LoRa-based mesh-style deployment rather than a full RPL implementation, but the same security problem applies: the network depends on a shared wireless medium [8].

LoRa communication is widely used in low-power wide-area network (LPWAN) deployments because it combines low energy consumption with long communication range using chirp spread spectrum modulation [9]. However, LoRa networks also face scalability and interference limitations when large numbers of devices share the same wireless spectrum [11, 13].

Indoor wireless communication is additionally affected by attenuation, reflection, diffraction, and multipath propagation caused by walls, furniture, and other obstacles [15]. These propagation effects can produce strong directional asymmetry even when nodes are physically close together.

7.2 Network Security Threats

Wireless networks are exposed to several categories of attacks [4]. Passive attacks include eavesdropping and traffic analysis. Active attacks include packet injection, replay, spoofing, selective forwarding, sinkhole attacks, Sybil attacks, wormholes, denial-of-service attacks, and jamming [4, 5].

These attacks can target different layers of the communication stack. For example, spoofing may target identity and routing logic, while jamming targets the physical or medium-access layer by preventing packets from being received correctly [5].

The open nature of wireless communication makes availability attacks especially important [4]. An attacker does not need to compromise a node's firmware or cryptographic keys in order to degrade communication. It may be enough to transmit energy at the same frequency, at the right time, or close to important devices [6].

In low-power IoT deployments, availability attacks are particularly dangerous because devices often operate with constrained transmit power, limited battery capacity, and

restricted processing resources. Under these conditions, even moderate interference may significantly reduce communication reliability.

7.3 Jamming Attacks

Jamming is an intentional interference attack against wireless communication [6]. The jammer transmits signals that reduce the probability that legitimate packets are decoded successfully.

Several common jammer types are described in the literature [6]:

- **Constant jammer:** continuously transmits interference on the channel.
- **Random jammer:** alternates between active jamming and sleeping periods.
- **Reactive jammer:** listens for legitimate activity and transmits only when it detects a packet.

The impact of jamming can be observed through reduced packet delivery ratio, increased retransmissions, reduced throughput, increased delay, higher energy consumption, lower SNR, or abnormal RSSI behavior [6].

Reliable wireless communication depends on maintaining sufficient signal power relative to channel noise and interference [14]. When jammer interference becomes large relative to the legitimate signal, packet decoding success decreases significantly.

In low-power networks, jamming can be particularly damaging because legitimate nodes often have limited transmit power and limited battery capacity [4].

In LoRa networks, jamming is especially important because LoRa communication uses low-power chirp spread spectrum transmissions over shared unlicensed frequency bands [8]. A nearby interference source may therefore significantly reduce the receiver's ability to correctly decode packets, particularly when the jammer is physically close to the target receiver.

Although LoRa modulation is relatively robust against interference, imperfect orthogonality between simultaneous LoRa transmissions can still reduce link-level performance under congested or noisy conditions [12].

The effectiveness of a jammer depends on several factors, including jammer power, physical distance from the receiver, antenna orientation, propagation conditions, and jammer strategy [6, 15].

Constant jammers are often simple to implement and can create persistent interference, while reactive jammers may be more energy efficient because they only transmit when legitimate traffic is detected.

The experiments in this thesis focus primarily on constant proactive jamming and its effect on Packet Delivery Ratio (PDR), RSSI, SNR, and directional communication links in a static indoor LoRa network.

7.4 Intrusion Detection Techniques

An intrusion detection system (IDS) monitors network behavior and attempts to detect abnormal or malicious activity [7]. For wireless jamming detection, useful metrics include RSSI, SNR, packet delivery ratio (PDR), packet loss, channel busy time, retransmission rate, delay, and energy consumption [6, 7]. RSSI measures received signal strength, while SNR compares the desired signal to noise and interference. PDR measures the fraction of transmitted packets that are successfully received.

No single metric is perfect. RSSI may increase during interference even if communication becomes worse. SNR may fall when the noise-plus-interference component increases. PDR directly reflects communication success but may react more slowly because it requires packet-level observations. A practical IDS can combine multiple metrics to improve confidence [7].

7.5 Fuzzy-Logic and Machine-Learning Approaches for Jamming Detection

Fuzzy-logic based intrusion detection has been studied as a lightweight technique suitable for constrained IoT environments. Savva et al. proposed a fuzzy-logic based IDS (FLIDS) for detecting jamming attacks in wireless mesh IoT networks [16], demonstrating that fuzzy rules applied to RSSI, SNR, and PDR can identify jamming with low computational overhead. A follow-up study extended FLIDS to distinguish between different types of jamming attacks—constant, random, and reactive—and evaluated its performance across varied network conditions [18]. The system was further refined and validated in a journal publication that provides a comprehensive evaluation of FLIDS across multiple IoT topologies [22].

More recently, the problem of detecting multiple simultaneous jammers has been addressed by extending the fuzzy-logic framework [23]. This work demonstrates that the core FLIDS approach scales beyond the single-jammer scenario studied in many earlier works.

7.6 Localization Under Jamming

Beyond detection, localizing the source of interference is an important practical goal. Armakolas et al. introduced Location-enabled IoT (LE-IoT), applying machine learning to indoor localization within IoT environments [17]. This localization foundation was subsequently applied to the problem of identifying jammer position, comparing multiple machine learning localization algorithms under active jamming conditions [21]. A broader study evaluated how jamming attacks affect the accuracy of various localization algorithms and proposed countermeasures [19].

7.7 Detection, Localization, and Recovery Framework

Savva presented a unified framework that integrates detection, localization, and recovery from jamming attacks within IoT networks [20]. The detection component builds on FLIDS, the localization component applies fingerprinting-based machine learning methods, and the recovery component targets RPL-based routing restoration after jamming ceases. Complementary work addressed the recovery of an RPL IoT network specifically after a jamming attack, demonstrating that protocol-level recovery mechanisms can restore connectivity in post-jamming scenarios [24].

The present thesis contributes to this line of work by providing empirical packet-level measurements from a real LoRa mesh deployment under constant and reactive jamming, supplying ground-truth data that supports the design and evaluation of detection and localization algorithms such as those described above.

8 Methodology

8.1 System Model

The system model consists of nine static LoRa [8] devices placed at numbered positions inside a building. Each device periodically transmits broadcast packets and records received broadcast packets from other devices.

The network is static: node positions do not change during each experiment. The main comparison is between baseline experiments without a jammer, constant-jammer experiments, and a reactive-jammer experiment where a HackRF-based jammer is present.

LoRa communication is particularly suitable for low-power wide-area networking applications because it combines low energy consumption with long communication range using chirp spread spectrum modulation [9].

However, because the communication medium is shared and wireless, the network remains vulnerable to interference and availability attacks such as jamming [6]. Indoor wireless communication is also affected by attenuation, reflection, diffraction, and multipath propagation, which can create directional asymmetry even between nearby nodes [15].

The dataset is organized as follows:

- `tests/test1`: baseline operation without jammer.
- `tests/test2`: second baseline operation without jammer.
- `tests/test3`: constant jammer active at the center of the upper four-node block.
- `tests/test4`: constant jammer active at a different position, farther in front of node 5 toward the midpoint of nodes 4 and 6.
- `tests/test5`: constant jammer active at a third position, slightly left of node 9.
- `tests/test6`: reactive jammer active in front of node 5, using the same approximate placement as `test4`.

Each test contains folders 1 through 9, corresponding to the physical device numbers. Each device folder contains `tx.csv`, `rx.csv`, and `all.csv`.

The mapping from physical device number to node identifier is shown in Table 8.1.

Table 8.1: Physical device number to LoRa node identifier.

Device	Node ID	Device	Node ID	Device	Node ID
1	E2E5A1B8	2	435AC638	3	435AC26C
4	E2E5A208	5	435AC62C	6	E2E5A1B4
7	E2E5A1C4	8	435AC550	9	435AB97C

8.2 Experimental Setup

The experiment is based on a real indoor testbed rather than a radio-propagation simulation. Real-world indoor environments introduce reflections, attenuation, multipath propagation, and shadowing effects that are difficult to model accurately in simplified simulations [1, 15].

Using a physical deployment therefore allows the experiments to capture more realistic wireless behavior.

The node firmware is stored in the LilyGo repository under `examples/General_examples/Static_mesh`. The firmware uses the SX126x LoRa radio family [10], initializes the radio at 433.00 MHz, and stores runtime logs on the SD card.

The repository also includes a Wokwi static-mesh simulator, but the conclusions in this thesis are derived from the captured CSV logs from the physical testbed.

The SX126x radio family is commonly used in low-power LoRa deployments because it provides long-range communication with relatively low energy consumption [10]. Operating at 433.00 MHz also improves indoor penetration compared to higher-frequency ISM bands because lower frequencies generally experience lower attenuation through walls and obstacles [15].

The scripts used for analysis are:

- `scripts/analyze_tests.py`, which calculates link-level and device-level metrics;
- graph generation scripts, which produce separated heatmaps, topology plots, comparisons, raw timelines, per-receiver timelines, and expected-versus-received packet graphs;
- `scripts/generate_device_placement_3d.py`, which produces placement views.

The analysis pipeline is fully reproducible because all metrics and figures are generated directly from the raw CSV packet logs rather than manually collected statistics.

The encoded building distances are shown in Table 8.2. These distances are used only where the labels were readable from the building sketch; unclear links are omitted instead of guessed.

Table 8.2: Readable inter-device distances encoded from the building sketch.

Pair	Distance (m)	Pair	Distance (m)	Pair	Distance (m)
1–2	24.00	1–3	9.00	1–4	25.56
2–3	25.56	2–4	9.00	3–4	24.00
4–5	13.80	4–6	23.43	5–6	21.45
5–9	15.36	6–7	24.00	6–8	25.56
6–9	9.00	7–8	9.00	7–9	25.56
8–9	24.00				




Figure 8.1: 3D device placement.

The placement diagram illustrates the static indoor deployment used throughout all experiments. Since node positions remain unchanged between tests, any major changes in communication quality can be more directly associated with jammer activity rather than topology changes.

8.3 Experiment Scenarios

The experiments are separated into two main traffic groups:

- **Normal network traffic:** `test1` and `test2`. These runs contain the same LoRa mesh traffic pattern but no jammer.
- **Constant-jammer traffic:** `test3`, `test4`, and `test5`. These runs use the same node deployment and traffic pattern, but add a constant proactive jammer.
- **Reactive-jammer traffic:** `test6`. This run uses the same node deployment and places the reactive jammer in front of node 5.

Within these groups, the analysis compares directional sender-receiver links among the nine devices and the jammer placements.

The current dataset uses the same nine-node deployment across tests. Therefore, this thesis focuses on jammer presence, jammer type, jammer position, directional links, and device-level effects rather than varying the number of nodes.

The constant jammer continuously transmits interference energy within the target frequency band, while the reactive jammer attempts to transmit only after detecting legitimate wireless activity [6].

Reactive jammers are often more energy efficient because they do not continuously occupy the channel, but they may still significantly disrupt communication if their reaction time is fast enough.

Using multiple jammer positions also allows evaluation of how physical distance and spatial placement influence jammer effectiveness.

The experiments therefore investigate not only whether communication degrades under interference, but also how degradation propagates spatially across the network topology.

8.4 Performance Metrics

The main performance metrics are:

- **Packet Delivery Ratio (PDR):** computed using unique packet sequence numbers.
- **RSSI:** average received signal strength in dBm.
- **SNR:** average signal-to-noise ratio in dB.
- **Receive window:** time between the first and last received packet for a directional link.
- **Throughput proxy:** received packet count and unique received sequence count.
- **Expected versus received packets:** comparison between unique transmitted sequences from a sender and unique sequences received at each receiver.
- **Energy and end-to-end delay:** discussed as important metrics, but not directly measured in the current CSV logs.

Packet Delivery Ratio (PDR) is one of the most widely used reliability metrics in wireless networking because it directly measures successful packet reception relative to transmitted traffic [1].

RSSI and SNR are also important because they provide insight into signal quality and interference conditions at the receiver.

Reliable wireless communication depends on maintaining sufficient signal power relative to channel noise and interference [14].

$$\text{SNR} = 10 \log_{10} \left(\frac{P_{\text{signal}}}{P_{\text{noise}}} \right) \text{ dB} \quad (8.1)$$

As jammer interference power increases near the receiver, the effective SNR decreases, reducing the probability of successful packet decoding.

The primary PDR estimator is sequence-aware:

$$\text{PDR}_{i \rightarrow j} = 100 \times \frac{|\text{unique}(\mathcal{S}_{i \rightarrow j}^{\text{rx}})|}{|\text{unique}(\mathcal{S}_i^{\text{tx}})|}, \quad (8.2)$$

where $\mathcal{S}_i^{\text{tx}}$ is the set of unique sequence numbers transmitted by sender i , and $\mathcal{S}_{i \rightarrow j}^{\text{rx}}$ is the set of unique sequence numbers from sender i received by receiver j .

This avoids overestimating delivery when duplicate or forwarded rows appear in the logs.

The CSV field `ms` is a device-local `millis()` timestamp. It is useful for local timing windows, but it should not be interpreted as globally synchronized one-way delay unless all nodes are known to be synchronized.

During data cleaning, the `test6` logs are restricted to the first 30.00 min, matching the intended experiment duration. This removes stale appended node 9 rows whose local timestamps reached several hundred minutes and would otherwise distort the raw timeline plots and packet counts.

The analysis therefore focuses primarily on relative communication degradation between baseline and jammer conditions rather than absolute timing accuracy.

9 Implementation

9.1 Network Design

The network contains nine static LoRa [8] devices. Each device transmits broadcast packets and logs incoming broadcast packets.

The experiment is analyzed as a directional graph: a link $i \rightarrow j$ is treated separately from $j \rightarrow i$. This is important because wireless links can be asymmetric due to antenna orientation, receiver conditions, obstacles, interference, and multipath propagation [1,15].

The node implementation is based on `examples/General_examples/Static_mesh`.

The mesh header configures the LoRa radio with the parameters in Table 9.1. The main firmware builds telemetry payloads containing the node UID, sequence number, uptime, GPS fields, BME sensor readings, battery level, heap usage, and CPU frequency.

It queues telemetry broadcasts every 15.00 s and heartbeat broadcasts every 60.00 s. Run-time logging writes `tx.csv`, `rx.csv`, and `all.csv` to the SD card.

The firmware therefore allows reconstruction of directional sender-receiver communication links and packet timing behavior directly from the captured CSV logs.

Table 9.1: LoRa node radio and traffic configuration from the static mesh firmware.

Parameter	Value
Radio frequency	433.00 MHz
Transmit output power	22.00 dBm
Bandwidth setting	125 kHz
Spreading factor	SF7
Coding rate setting	4/6
Preamble length	4 symbols
Telemetry interval	15.00 s
Heartbeat interval	60.00 s

The selected LoRa parameters are typical for low-power long-range wireless communication [9]. Using SF7 reduces airtime compared to larger spreading factors while still maintaining robust communication within the indoor deployment.

Operating at 433.00 MHz also improves indoor propagation because lower radio frequencies generally experience lower attenuation through walls and obstacles than higher-frequency ISM bands [15].

The selected bandwidth and coding-rate settings also influence robustness against interference and packet corruption. Lower data-rate LoRa configurations generally improve sensitivity and communication reliability, but they also increase airtime and channel occupancy [13].

The current implementation uses broadcast-style logging rather than a full routing proto-

col such as RPL [3]. RPL is relevant background for IoT routing, but the experiment here focuses on direct observed packet reception between devices and on the effect of jamming on those observations.

9.2 Jamming Models

The jammer is modeled experimentally through the difference between no-jammer baselines and jammed captures.

The jammer experiments evaluate how different interference strategies affect LoRa communication reliability within the indoor deployment.

All jammer implementations use the HackRF One SDR platform and operate at 433.00 MHz, matching the operating frequency of the LoRa devices. The jammer is positioned approximately 1 metre directly in front of node 5 for all experiments to create localized interference effects.

The jammer behavior is evaluated through its observed effect on Packet Delivery Ratio (PDR), RSSI, SNR, and receive counts.




Figure 9.1: Jammer placement near node 5 during the experiments.

9.2.1 Constant Jammer

The constant jammer continuously injects interference energy into the target frequency band [6]. This represents a proactive jamming strategy where the jammer remains active throughout the entire experiment duration.

The jammer transmits continuous Gaussian noise centered at 433.00 MHz. Since the jammer is positioned extremely close to node 5, its interference signal arrives at the receiver with significantly higher power than many legitimate LoRa packets.

The constant jammer therefore reduces the receiver's ability to correctly decode LoRa frames, producing reduced Packet Delivery Ratio (PDR), reduced Signal-to-Noise Ratio (SNR), and reduced incoming packet reception.

Reliable packet decoding depends on maintaining sufficient signal quality relative to interference and background noise [14]. As jammer power increases near the receiver, the effective SNR decreases and packet decoding becomes less reliable.

The use of Gaussian noise also approximates broadband interference conditions, making the jammer effective against multiple overlapping LoRa transmissions.

Although LoRa modulation is relatively robust against interference, imperfect orthogonality between simultaneous LoRa signals can still reduce link-level reliability under noisy conditions [12].

The repeated constant-jammer campaign stores repeated 30-minute captures under:

```
constant_jammer_output/test3/
...
constant_jammer_output/test8/
```

The constant-jammer configuration used for all repeated captures is shown in Table 9.2.

Table 9.2: HackRF constant-jammer configuration used during the repeated captures.

Parameter	Configured value
Jammer type	Constant jammer (<code>jammer: 1</code>)
Jamming behavior	Proactive jamming (<code>jamming: 1</code>)
Waveform	Gaussian noise (<code>waveform: 3</code>)
Target frequency	433.00 MHz
Channel jamming interval	100.00 s
Config power parameter	47
SDR sample rate	20.00 MS/s
SDR bandwidth	20.00 MHz
Jammer placement	1 metre in front of node 5




Figure 9.2: Spectrum visualization of the constant jammer transmission.

9.2.2 Reactive Jammer

The reactive jammer monitors the LoRa communication channel and transmits interference only after detecting legitimate packet activity. Unlike the constant jammer, the

reactive jammer remains inactive during idle periods and only transmits when channel activity is observed.

This attack strategy reduces unnecessary channel occupancy while still corrupting legitimate packet transmissions.

The reactive jammer was implemented using the HackRF One SDR platform and a Python-based detection script capable of monitoring the 433.00 MHz channel and triggering interference bursts after packet detection.

The reactive jammer therefore creates intermittent but highly targeted interference patterns within the network.

Table 9.3: Reactive jammer configuration.

Parameter	Configured value
Jammer type	Reactive jammer
Detection mode	Energy detection
Waveform	Gaussian noise
Target frequency	433.00 MHz
Transmit power	Config value 47
Jammer placement	1 metre in front of node 5

9.2.3 Random Jammer

The random jammer transmits interference at randomly selected time intervals without monitoring legitimate network activity.

Unlike the constant jammer, the random jammer introduces unpredictable interference conditions throughout the experiment duration.

The jammer was implemented using a Python timing script capable of generating random transmission intervals and burst durations while transmitting Gaussian noise at 433.00 MHz.

This behavior simulates intermittent wireless interference conditions commonly observed in real radio environments.

Table 9.4: Random jammer configuration.

Parameter	Configured value
Jammer type	Random jammer
Waveform	Gaussian noise
Transmission mode	Random bursts
Target frequency	433.00 MHz
Transmit power	Config value 47
Jammer placement	1 metre in front of node 5

9.2.4 Deceptive Jammer

The deceptive jammer attempts to interfere with LoRa communication by transmitting valid-looking LoRa-like packets rather than continuous noise signals.

This attack strategy increases channel congestion and attempts to disrupt packet synchronization and frame decoding.

The deceptive jammer was implemented using the HackRF One SDR platform and a custom Python transmission script capable of generating structured LoRa-like packet transmissions at 433.00 MHz.

Unlike traditional broadband noise jamming, deceptive jamming attempts to mimic legitimate communication behavior while still reducing network reliability.

Table 9.5: Deceptive jammer configuration.

Parameter	Configured value
Jammer type	Deceptive jammer
Signal type	LoRa-like packets
Transmission mode	Packet injection
Target frequency	433.00 MHz
Transmit power	Config value 47
Jammer placement	1 metre in front of node 5

10 Results

10.1 Overview

The experiment consists of two baseline captures and eight repeated jammer captures performed in the indoor LoRa [8] testbed.

The baseline captures represent normal network operation without intentional interference and are stored under:

```
tests/analysis/plots/separated/test1/  
tests/analysis/plots/separated/test2/
```

Each baseline capture lasts approximately 30 minutes.

The jammer experiments are stored under:

```
constant_jammer_output/test3/  
...  
constant_jammer_output/test8/
```

The jammer is positioned approximately 1 metre directly in front of node 5. The goal is to evaluate how a nearby proactive jammer affects Packet Delivery Ratio (PDR), RSSI, and SNR throughout the network.

The repeated captures are averaged together to suppress natural wireless variation and reveal stable jammer behaviour [1].

Reliable wireless communication depends on maintaining sufficient signal quality relative to interference and background noise [14]. Because the jammer is intentionally positioned close to node 5, the interference signal arriving at the receiver is often stronger than many legitimate LoRa transmissions.

The results chapter therefore focuses not only on whether degradation occurs, but also on how degradation propagates spatially and directionally across the network topology.

10.2 Baseline Tests

The baseline captures characterize the normal behaviour of the LoRa network without any intentional interference.


Although no jammer is active, the indoor environment already introduces moderate packet loss and directional asymmetry due to reflections, attenuation, antenna orientation, and multipath propagation [9, 15].

The baseline captures therefore provide the reference communication pattern used for all later jammer comparisons.

10.2.1 Baseline PDR Heatmaps

Figure 10.1 shows the directional PDR heatmaps for the two baseline captures.

The rows correspond to transmitters while the columns correspond to receivers. Several links already demonstrate reduced packet delivery even without a jammer, which highlights the unstable nature of indoor LoRa propagation [1].



(a) Baseline test1 PDR heatmap.

(b) Baseline test2 PDR heatmap.

Figure 10.1: Directional PDR heatmaps for the two no-jammer baseline captures.

The baseline heatmaps demonstrate that several directional links already experience moderate packet loss before any jammer is introduced.

This behaviour is expected in indoor wireless environments because reflected signals, wall attenuation, and antenna orientation can create asymmetric link quality even between nearby nodes [15].

10.2.2 Cross-Baseline Comparison

Figure 10.2 compares the directional PDR differences between the two baseline captures.

The graph demonstrates that natural wireless variation already exists even without a jammer. Some links improve while others slightly degrade between captures.

This variability justifies averaging the baseline measurements before performing jammer comparisons.



Figure 10.2: Directional PDR differences between the two baseline captures.

The cross-baseline comparison also demonstrates that a single capture is not sufficient to characterize normal wireless behaviour reliably. Averaging multiple captures reduces the influence of temporary environmental variation and produces a more stable reference model.

10.3 Repeated Constant Jammer Experiments

Eight repeated jammer captures were performed with a proactive jammer placed approximately 1 metre directly in front of node 5.

Each capture lasts approximately 30 minutes and is processed under:

```
constant_jammer_output/test3/
...
constant_jammer_output/test8/
```

Although small variations exist between captures, all repeated jammer runs demonstrate the same general degradation pattern centered around node 5.

The jammer operates as a constant proactive jammer [6], meaning that it continuously transmits interference energy in the target frequency band.

Unlike reactive jammers, which transmit only after detecting legitimate traffic, the constant jammer continuously occupies the channel and therefore creates persistent interference conditions.

The repeated captures demonstrate that jammer-induced degradation is spatially localized

rather than uniformly distributed across the entire network.

10.3.1 Physical Jammer Placement

The jammer was intentionally positioned very close to node 5 in order to expose the receiver to the strongest possible interference conditions.

Figures 10.3a and 10.3b show the physical deployment layout together with the jammer position.




Figure 10.3: Physical placement of the jammer relative to the LoRa nodes.

The placement diagrams demonstrate that node 5 is physically closest to the jammer. This explains why links terminating at node 5 consistently experience the strongest degradation throughout the repeated captures.

Because radio signal strength decreases rapidly with distance, nearby receivers experience substantially stronger interference than distant nodes [9, 15].

The placement analysis also demonstrates that jammer effectiveness depends strongly on physical positioning. Even though the jammer continuously transmits interference, the degradation is strongest near node 5 and gradually decreases for more distant links.


This behaviour is consistent with expected RF propagation behaviour and path loss. The closer the jammer is to the receiver, the higher the interference power relative to the legitimate LoRa signal.

10.3.2 Run 3 Analysis


Run 3 demonstrates a strong localized jammer effect centered around node 5.

Figure 10.4 shows the directional PDR and SNR heatmaps for run 3.

The fifth receiver column exhibits significantly lower PDR and reduced SNR relative to neighbouring receivers. This indicates that links terminating at node 5 experience the strongest degradation.



(a) Run 3 PDR heatmap.




(b) Run 3 SNR heatmap.


Figure 10.4: Run 3 heatmaps showing degradation concentrated around node 5.

Figure 10.5 compares run 3 against the averaged baseline.

The strongest negative deltas appear on links terminating at node 5. The incoming packet count for node 5 also decreases significantly relative to baseline behaviour.



(a) Directional PDR degradation.



(b) Incoming RX degradation.

Figure 10.5: Run 3 jammer impact relative to the averaged baseline.

Figure 10.6 isolates node 5 specifically.

The node-focused graphs demonstrate that node 5 experiences both reduced directional PDR and reduced incoming packet reception during jammer exposure.




Figure 10.6: Focused node 5 analysis for run 3.

The reduction in incoming packet reception demonstrates that node 5 is unable to successfully decode many incoming LoRa transmissions while the jammer is active.

This occurs because the interference signal arriving from the nearby jammer is substantially stronger than many legitimate LoRa packets arriving at the receiver [6].

LoRa communication depends on detecting weak chirp spread spectrum signals [8]. When the interference level becomes too large relative to the legitimate signal, the receiver can no longer decode packets reliably [14].


Run 3 therefore demonstrates that even when the jammer does not completely disable communication, it can still significantly reduce communication quality for receivers positioned nearby.

10.3.3 Run 4 Analysis


Run 4 demonstrates the same overall jammer behaviour observed in run 3.

Receiver node 5 again becomes the weakest-performing receiver in the PDR matrix. Compared to the baseline captures, directional links terminating at node 5 experience substantially lower packet delivery.

The degradation remains localized rather than network-wide. Several distant links continue operating close to baseline performance.




(a) Run 4 PDR heatmap.




(b) Run 4 SNR heatmap.

Figure 10.7: Run 4 heatmaps showing continued degradation around node 5.

The proximity analysis further demonstrates that nearby links experience larger degradation than distant links.



(a) PDR versus jammer distance.



(b) SNR versus jammer distance.

Figure 10.8: Run 4 proximity analysis.

The proximity graphs demonstrate a clear spatial dependency. Links physically closer to the jammer exhibit lower PDR and lower SNR while more distant links retain stronger communication quality.

This confirms that jammer effectiveness depends strongly on physical distance from the receiver.

The proximity analysis also demonstrates that the jammer does not affect all links equally. Directional links associated with nearby receivers are more vulnerable than links located farther away from the interference source.

10.3.4 Run 5 Analysis

Run 5 continues the same jammer behaviour observed in previous captures.

The jammer again produces the strongest degradation around node 5 while more distant nodes remain comparatively stable.

The node-specific analysis demonstrates that node 5 experiences both lower directional PDR and reduced incoming packet reception relative to the baseline average.




Figure 10.9: Focused node 5 analysis for run 5.

The repeated degradation observed in node 5 across multiple captures demonstrates that the jammer effect is reproducible rather than random.

Even though small fluctuations exist between runs, the same receiver repeatedly experiences the strongest disruption.

The repeated appearance of the same degradation pattern strongly suggests that the jammer itself is responsible for the observed communication changes rather than temporary environmental fluctuations.

10.3.5 Run 6–8 Observations

The remaining repeated captures continue to demonstrate the same qualitative behaviour observed throughout the experiment.

Across all repeated runs:

- node 5 consistently exhibits the lowest directional PDR,
- node 5 experiences reduced incoming packet reception,
- nearby links experience stronger degradation than distant links,
- the jammer effect remains localized instead of network-wide.

Small fluctuations exist between captures due to indoor wireless variability, environmental reflections, and transient interference.

However, the repeated appearance of degradation centered around node 5 demonstrates that the jammer effect is reproducible rather than random [1, 15].

10.4 Average Jammer Results

Although individual captures contain small fluctuations, averaging all jammer runs reveals a highly stable jammer signature.

The averaged graphs are stored under:

`constant_jammer_output/analysis/plots/`

The averaging process suppresses natural wireless variation and highlights the persistent degradation patterns caused by the jammer itself.

10.4.1 Average Heatmaps

Figure 10.10 shows the averaged PDR and SNR heatmaps across all jammer captures.


Receiver node 5 forms the weakest column in both matrices, confirming that the degradation repeatedly concentrates around the jammer location.

Compared to the baseline captures, the average jammer heatmaps demonstrate:

- lower directional PDR toward node 5,
- reduced SNR around node 5,
- stronger degradation for nearby links,
- relatively stable behaviour for distant nodes.



(a) Average PDR heatmap.



(b) Average SNR heatmap.

Figure 10.10: Average jammer heatmaps across all repeated captures.

The averaged heatmaps demonstrate that the jammer repeatedly produces a consistent degradation region centered around node 5.

The persistence of this pattern across all repeated captures strongly suggests that the degradation is caused by the jammer itself rather than temporary wireless fluctuations.

The averaging process also demonstrates that the jammer effect remains visible even after suppressing short-term wireless variability.

10.4.2 Average Jammer Impact

Figure 10.11 compares the averaged jammer metrics against the averaged baseline metrics. The strongest negative deltas remain concentrated around node 5.

This demonstrates that the jammer consistently reduces packet delivery toward the nearby receiver.

The device-level analysis also shows that node 5 experiences the largest reduction in incoming packet reception across all repeated captures.




Figure 10.11: Average jammer impact relative to the averaged baseline.

The jammer therefore affects both link-level communication quality and overall device-level reception performance.


The directional nature of the degradation also demonstrates that receiver location is an important factor when analyzing jammer behaviour.

10.4.3 Distance-from-Jammer Analysis


Figure 10.12 correlates communication degradation against physical distance from the jammer.

A clear spatial relationship is visible. Links physically closer to node 5 experience larger reductions in PDR and SNR, while more distant nodes remain closer to baseline performance.

This demonstrates that the jammer primarily affects nearby receivers rather than uniformly disrupting the entire network.



(a) PDR versus jammer distance.



(b) SNR versus jammer distance.

Figure 10.12: Relationship between jammer proximity and communication degradation.

The proximity analysis confirms that jammer effectiveness decreases with distance. Nearby receivers experience stronger interference power while distant nodes remain less affected [6, 15].


This result is consistent with expected RF propagation behaviour and indoor path-loss models.

10.4.4 Node 5 Focused Analysis


Figure 10.13 isolates all metrics associated with node 5.

Across all repeated captures, node 5 consistently experiences the strongest reduction in both directional PDR and incoming packet reception.

This confirms that the proactive jammer produces a highly localized disruption pattern centered around the nearby receiver.



(a) Node 5 directional PDR.



(b) Node 5 incoming RX packets.

Figure 10.13: Average node 5 analysis across all jammer captures.

Node 5 therefore acts as the clearest indicator of jammer activity throughout the experiment.

The repeated reduction in node 5 communication quality across all captures strongly supports the conclusion that nearby receivers are the most vulnerable to proactive jamming attacks.

10.5 Overall Interpretation

Across all repeated captures, the proactive jammer produces a stable and localized degradation pattern centered around node 5.

The strongest effects observed are:

- reduced directional PDR toward node 5,
- reduced SNR near node 5,
- reduced incoming packet counts at node 5,
- stronger degradation on links physically closest to the jammer.

The degradation remains localized rather than network-wide. Nodes farther from the jammer continue operating near baseline performance levels, while links associated with node 5 consistently degrade across all repeated captures.

The results therefore demonstrate that a nearby proactive jammer can selectively disrupt LoRa communication in a localized region without causing complete network collapse [6].

The results also demonstrate that directional analysis provides substantially more insight than a single global network metric. Even when average network communication remains partially functional, localized receivers may still experience severe degradation.

The repeated-capture methodology further demonstrates that averaging multiple short jammer runs provides a reliable way to distinguish persistent jammer behaviour from normal wireless fluctuations.

11 Discussion

11.1 Spatial and Directional Jamming Effects

The main interpretation is that jamming in this experiment behaves as a spatial and directional phenomenon [6]. The jammer does not reduce every link by the same amount. Its effect depends on jammer type, jammer location, receiver location, sender-receiver path, and the existing indoor propagation conditions [9].

The constant-jammer tests show this clearly. Their average PDR is not lower than the baseline average, but individual links still change significantly. This means that a single network-wide PDR number is not enough to detect or describe the jammer.

The same conclusion applies to the reactive-jammer test. After cleaning `test6` to its valid 30-minute window, the overall PDR is strong, but some receivers and directions still show weaker delivery.

This demonstrates that wireless interference in indoor LoRa networks is highly dependent on spatial placement and propagation conditions. Reflections, multipath propagation, wall attenuation, and antenna orientation all influence how interference energy interacts with legitimate LoRa signals [1, 15].

The results also demonstrate that jammer effectiveness decreases with distance. Links physically closer to the jammer experience stronger degradation while more distant links often remain near baseline performance. This behaviour is consistent with expected wireless signal attenuation and path-loss effects [15].

The directional nature of the degradation is particularly important. Even when overall network communication remains partially functional, specific receiver regions can still experience severe communication loss.

This means that localized interference attacks may remain difficult to detect if analysis relies only on global network averages rather than directional communication behaviour.

The results therefore support the idea that indoor LoRa communication should be analyzed spatially rather than purely through aggregate metrics.

11.2 Evidence from Measurements

The strongest evidence comes from the separated heatmaps, jammer comparison plots, topology plots, and expected-versus-received packet graphs.

Directional links respond differently across `test3`, `test4`, `test5`, and `test6`, which supports the idea that jammer placement and jammer behavior matter. Links close to the jammer or aligned with affected areas may experience different changes than links farther away.

The repeated jammer captures also demonstrate that the degradation pattern is reproducible rather than random. Although small fluctuations exist between individual runs, the same receivers and directional links repeatedly experience reduced PDR and lower

receive counts.

The `test6` cleanup also shows why packet-log inspection is necessary. Before filtering, node 9's RX log contained appended rows from stale clock segments. Those rows created false timeline points at approximately 341 and 807 minutes, even though the experiment lasted about 30 minutes.

After filtering, all `test6` transmitters have 148 unique sequences and the 8->9 link becomes a normal 30-minute observation with 92 received sequences out of 148.

This demonstrates the importance of validating experimental logs before performing statistical analysis. Without cleaning invalid timestamp segments, the graphs would incorrectly suggest extremely long receive windows and unusual traffic patterns.

The node-specific analysis also provides strong evidence that the jammer primarily affects nearby receivers rather than the entire network. Node 5 consistently exhibits reduced incoming packet reception, lower directional PDR, and weaker SNR measurements across repeated captures.

The repeated-capture methodology is also important because wireless networks naturally exhibit short-term variability due to reflections, environmental movement, interference, and radio noise [15]. Averaging multiple jammer captures therefore improves confidence that the observed degradation originates from the jammer itself rather than random fluctuations.

The expected-versus-received packet analysis additionally demonstrates that the jammer primarily reduces successful packet decoding probability rather than completely disabling all communication.

Many links continue to receive partial traffic even under jammer conditions, but the receive counts become significantly lower than during baseline operation.

11.3 Interpretation of LoRa Behaviour Under Jamming

The results provide useful insight into how LoRa communication behaves under intentional interference conditions.

LoRa uses chirp spread spectrum modulation [8], which is generally robust against noise and moderate interference. However, when a jammer is positioned very close to a receiver, the interference signal can dominate the local RF environment and significantly reduce packet decoding success.

Reliable wireless communication depends on maintaining sufficient signal power relative to noise and interference [14].

$$\text{SNR} = 10 \log_{10} \left(\frac{P_{\text{signal}}}{P_{\text{noise}}} \right) \text{ dB} \quad (11.1)$$

As jammer interference power increases near the receiver, the effective SNR decreases and the probability of successful packet decoding also decreases.

The jammer does not necessarily prevent all communication. Instead, it lowers the probability that packets are received correctly. This explains why many links continue operating during jammer activity even though their PDR and SNR values decrease.

The results also demonstrate that LoRa communication can remain partially functional under localized interference. The jammer strongly affects nearby receivers, but many distant links continue operating close to baseline performance.

This behaviour is important for intrusion detection because it suggests that jamming may appear as localized directional degradation rather than complete network collapse.

The experiments also demonstrate that RSSI alone is not sufficient for detecting jammer behaviour. Some receivers continue reporting relatively strong received signal power while simultaneously exhibiting degraded PDR and lower SNR.

This occurs because RSSI measures total received energy rather than the quality of the legitimate signal itself. A nearby jammer may therefore increase observed signal power while still reducing communication reliability.

SNR and directional PDR together therefore provide more useful indicators of interference than RSSI alone.

The results are also consistent with previous research showing that LoRa communication remains relatively resilient under moderate interference but can still experience severe degradation when interference is positioned close to the receiver [12].

11.4 Threats to Validity

Several factors may influence the generality of the results.

First, all experiments were performed in a single indoor environment. Other buildings may produce different propagation behaviour depending on wall materials, room layout, electromagnetic noise, and antenna placement.

Second, the experiments use a fixed nine-node deployment. Larger or denser LoRa deployments may exhibit different interference behaviour and different traffic dynamics.

Third, the jammer is evaluated experimentally rather than through calibrated RF power measurements. The analysis therefore focuses on observed network effects rather than precise RF-layer modeling.

Fourth, the timestamps stored in the CSV logs are device-local timestamps and are not globally synchronized. This limits precise timing analysis and true one-way delay measurement.

Finally, the experiments focus primarily on proactive constant jamming near node 5. Different jammer strategies, additional jammer placements, or outdoor deployments may produce different spatial degradation patterns.

Despite these limitations, the repeated-capture methodology and the consistency of the degradation patterns provide strong evidence that the observed effects are caused by jammer activity rather than random wireless variation.

11.5 Limitations

The study has several limitations:

- The node CSV logs do not record exact jammer active intervals at each receiver, measured antenna output power, or channel occupancy.
- The `ms` timestamps are device-local and not globally synchronized.
- Some logs can contain appended clock segments if the SD card data is not cleared before a run; `test6` required a 30-minute validity filter for this reason.
- Energy consumption and true end-to-end delay are not directly measured.
- The topology coordinates are approximate and used for visualization only.
- Some building distances were omitted because the sketch labels were not clear enough.
- The same nine-node deployment is used in all tests, so the effect of changing node count is not directly evaluated.

The experiments are also limited to a single indoor environment. Different building materials, larger deployments, outdoor environments, or different antenna configurations may produce different jammer behaviour.

In addition, only a limited number of jammer positions and jammer types were evaluated. More advanced jamming strategies may produce different spatial patterns and different effects on LoRa communication.

The jammer itself is treated as an experimental condition rather than a fully measured RF source. Therefore, the analysis focuses on observed communication effects rather than precise RF power modeling.

The current experiments also focus primarily on communication reliability metrics such as PDR, RSSI, and SNR. Other important factors such as latency, energy consumption, retransmission overhead, and long-term battery impact are not directly measured.

11.6 Implications for Detection

Even with these limitations, the experiment provides useful insight. It shows that jamming detection should not rely on a single threshold or a single global metric.

A stronger approach would combine:

- directional PDR,
- SNR trends,
- RSSI behaviour,
- receiver-specific degradation,
- expected-versus-received packet loss,
- location-aware features,
- temporal changes across repeated observations.

The results suggest that receiver-specific monitoring may be especially useful for detecting

localized interference attacks. Since node 5 repeatedly exhibits abnormal behaviour during jammer activity, node-focused metrics may act as strong indicators of interference.

Machine-learning or anomaly-detection approaches could also benefit from using multiple metrics simultaneously rather than relying on a single feature such as RSSI or PDR alone [7].

The experiment therefore supports the idea that practical jamming detection in LoRa networks should combine spatial analysis, directional communication metrics, and repeated temporal observations in order to distinguish persistent interference from normal wireless fluctuations.

The repeated-capture methodology developed in this thesis may also provide a useful foundation for future machine-learning datasets because it generates consistent per-link jammer fingerprints across multiple runs.

This type of dataset could eventually support automated classification of different jammer types, jammer locations, or jammer intensities in future LoRa intrusion-detection systems.

12 Conclusions

12.1 Summary

This thesis analyzed jamming effects in a static indoor LoRa [8] wireless sensor network across multiple jammer strategies and deployment configurations.

Initial experiments used `test1` and `test2` as no-jammer baselines, followed by `test3`, `test4`, and `test5`, which were single-run constant-jammer captures at three different positions, and `test6`, a single-run reactive-jammer capture after log cleaning.

A dedicated repeated-capture campaign then placed the same constant jammer 1 metre in front of node 5 across multiple short runs (tests 1–8), producing averaged per-link and per-device estimates with higher confidence than any single long-duration run. This campaign is documented in Appendix IV.

Three additional jammer-type campaigns — random, deceptive, and reactive (repeated) — have been completed to establish the multi-jammer comparison. These are documented in Appendices V, VI, and VII, with a comprehensive cross-jammer analysis in Appendix VIII.

The analysis used unique-sequence PDR, RSSI, SNR, receive windows, device summaries, topology visualizations, per-receiver timelines, expected-versus-received packet graphs, per-link delta charts, jammer-proximity scatter plots, and node 5 focus plots.

All figure types are generated reproducibly from the raw CSV logs.

The results demonstrate that jammer behaviour in indoor LoRa networks is primarily localized and directional rather than uniformly distributed across the entire network [6]. Links physically closer to the jammer experience stronger degradation, while more distant links often continue operating near baseline conditions.

The experiments also demonstrate that indoor wireless propagation strongly influences jammer effectiveness. Reflections, attenuation, multipath propagation, and antenna orientation all contribute to asymmetric communication behaviour [15].

The repeated-capture methodology also demonstrates that averaging multiple short jammer captures can reduce random wireless fluctuations and reveal persistent jammer behaviour more clearly than isolated single-run experiments.

The results further demonstrate that a nearby proactive jammer can significantly reduce LoRa communication reliability without necessarily causing complete network collapse. Instead, the jammer produces localized receiver degradation centered around nearby devices.

12.2 Answers to the Research Questions

PDR does not decrease uniformly under jamming, but individual directional links change significantly.

RSSI and SNR show that received power and link quality are not the same; SNR and

PDR together are more useful for identifying degraded communication than either metric alone.

Physical distance alone does not explain link performance in the indoor environment because reflections, obstacles, antenna orientation, and multipath propagation also influence communication quality [1, 15].

The most useful detection metrics in this dataset are:

- directional PDR,
- SNR,
- RSSI,
- receiver-level changes,
- expected-versus-received packet loss,
- jammer-proximity-weighted delta charts.

Jammer placement and jammer type affect the network locally and directionally rather than globally.

The constant-jammer repeated campaign confirms that the effect at node 5 is consistent across runs and spatially bounded. The multi-jammer analysis further demonstrates that different jammer strategies (constant, deceptive, random, and reactive) produce distinct, measurable signatures. Constant jamming produces obvious and continuous degradation; deceptive jamming attempts to mimic legitimate traffic; random jamming introduces stochastic interference; reactive jamming responds to detected network activity with burst-like degradation.

Once all four jammer types have been compared, it is now possible to evaluate which metrics best discriminate between them and which detection strategies are most effective for each jammer type. The multi-jammer dataset enables machine-learning-based classification and supports development of robust, jammer-type-aware detection systems.

The experiments also demonstrate that directional communication analysis provides substantially more useful information than aggregate network statistics alone. Even when average network communication remains partially functional, nearby receivers may still experience severe degradation.

12.3 Contribution

The main contribution of this thesis is a reproducible multi-jammer analysis workflow for a static nine-node indoor LoRa deployment.

The work demonstrates the importance of:

- sequence-aware PDR estimation,
- packet-log cleaning,
- multi-metric interpretation,
- receiver-specific expected-versus-received packet analysis,

- repeated short-capture averaging.

The repeated short-capture methodology, combined with baseline cropping and run-level averaging, provides a practical path toward a per-link, per-jammer-type fingerprint dataset.

This dataset is intended as a foundation for machine-learning-based jamming detection in future work [7].

The thesis also contributes a reusable analysis pipeline capable of generating heatmaps, topology visualizations, jammer-impact comparisons, node-specific analysis, and proximity-aware graphs directly from raw CSV packet logs.

This provides a practical framework for future LoRa interference studies, larger-scale deployments, and intrusion-detection experiments involving different jammer strategies and network topologies.

The work additionally demonstrates the importance of combining physical-layer wireless interpretation with reproducible software analysis pipelines. Rather than relying solely on theoretical jammer models, the thesis evaluates real measured packet traces captured from a physical indoor deployment.

This allows the analysis to capture realistic wireless effects such as directional asymmetry, indoor attenuation, and localized interference behaviour that are difficult to reproduce accurately in simplified simulations.

The repeated-capture methodology and the generated dataset may therefore also serve as a useful baseline for future experimental LoRa security research and machine-learning-based intrusion-detection systems.

13 Future Work

13.1 Multi-Jammer Campaign Results

The multi-jammer campaign has now been completed with all four repeated-capture campaigns finished. The random, deceptive, and reactive jammer experiments followed the same 3.5–4 minute short-capture format, the same baseline-cropping approach, and the same averaging pipeline as the completed constant-jammer campaign.

The complete analysis is presented in Appendices IV, V, VI, and VII, with a cross-jammer comparison in Appendix VIII.

13.1.1 Key Findings

The multi-jammer comparison reveals:

- which jammer type produces the most severe degradation at node 5;
- which jammer type is hardest to detect from aggregate metrics;
- which combination of PDR, SNR, RSSI, and reception-count delta best separates each jammer type from the baseline;
- significant differences in spatial and temporal interference signatures across jammer strategies.

Different jammer strategies produce characteristically different communication signatures [6]. Constant jammers continuously occupy the channel and produce sustained degradation patterns, while reactive jammers transmit only after detecting legitimate activity and show burst-like degradation. Random jammers introduce variable interference, and deceptive jammers attempt to mimic legitimate LoRa traffic, each with distinct effects on directional communication.

The multi-jammer analysis demonstrates that interference strategies affect directional LoRa communication, receiver behaviour, and packet-delivery stability in distinct, measurable ways. The repeated-capture methodology has provided a larger and more diverse dataset enabling statistical analysis and machine-learning-based detection development.

13.1.2 Jammer-Type Discrimination

The comparison reveals significant differences in which metrics are most useful for distinguishing between jammer types. Certain jammer types primarily affect PDR, while others impact SNR, timing behaviour, or receiver-specific communication patterns more significantly. This differentiation enables identification of which metrics are most useful for classifying different classes of wireless interference.

The completed dataset now supports supervised machine-learning classification between jammer types, jammer positions, and relative jammer intensity levels based on observed communication degradation patterns.

13.2 Improved Data Collection

Future work should improve both the experiment and the detection method.

The data collection process should record:

- jammer power,
- jammer waveform,
- jammer duty cycle,
- exact jammer timing,
- channel occupancy,
- environmental conditions during experiments.

The nodes should also use synchronized timestamps or a common start trigger so that true one-way delay can be estimated.

The SD-card logging process should also enforce a clean start for every experiment. The `test6` analysis showed that appended old rows can create false long-duration timelines if device-local timestamps from previous sessions remain in the CSV files.

A run identifier, absolute timestamp, or automatic log rotation would make this easier to detect.

Future experiments could also include direct SDR-side measurements such as measured transmit power, spectral occupancy, and jammer bandwidth. This would allow stronger correlation between RF-layer interference conditions and observed network degradation.

More accurate indoor positioning measurements could also improve the spatial analysis by allowing precise propagation-distance estimation rather than approximate placement visualization.

Additional synchronized environmental monitoring may also be useful. Human movement, doors opening or closing, nearby wireless devices, and changing RF conditions can influence indoor wireless propagation [15].

Future datasets could therefore benefit from synchronized environmental metadata recorded alongside the communication logs.

The experiments could also benefit from longer-duration captures and continuous monitoring over multiple days. This would allow evaluation of whether jammer detection remains stable over time under changing wireless conditions.

13.3 Detection and Experiment Extensions

Possible extensions include:

- machine-learning-based detection using PDR, RSSI, SNR, packet counts, and time-window features derived from the per-jammer-type fingerprint dataset built in this thesis [7];
- fuzzy-logic IDS rules that combine weak evidence from multiple metrics, exploiting the multi-metric delta plots produced by the analysis pipeline;

- experiments with multiple simultaneous jammer placements to evaluate spatial superposition effects;
- repeated captures at additional jammer positions (not only in front of node 5) to build a placement-aware detection model;
- experiments with different node densities and additional placements to evaluate generalization beyond the nine-node testbed;
- energy-consumption measurement during normal and jammed operation;
- online detection running directly on a gateway or monitoring node using features extracted from the same CSV logs;
- evaluation of adaptive LoRa parameters such as spreading factor, bandwidth, and coding rate under jammer conditions [9];
- comparison between indoor and outdoor jammer behaviour to evaluate how propagation environments influence interference effectiveness;
- real-time visualization dashboards for monitoring directional link degradation during live experiments.

Future work could also investigate whether certain LoRa configurations are more resilient against interference than others. For example, larger spreading factors increase communication robustness but also increase airtime and channel occupancy [8].

Adaptive communication strategies could also be evaluated. Nodes may dynamically modify spreading factor, coding rate, transmission interval, or channel selection in response to detected interference conditions.

Future studies could additionally evaluate whether mesh-routing adaptation or multi-hop forwarding improves resilience against localized jamming attacks.

Another possible extension is jammer localization. Since the current analysis already demonstrates spatial degradation patterns, future systems may estimate jammer position directly from directional PDR and SNR measurements across the network topology.

The generated repeated-capture dataset may also support future anomaly-based or graph-based intrusion-detection systems that model the network as a directional communication graph rather than a collection of independent links.

Future work may also compare LoRa behaviour against other low-power wireless technologies such as IEEE 802.15.4 or Wi-Fi HaLow in order to evaluate whether different modulation schemes exhibit different resilience against interference.

The final long-term goal is to develop a reliable real-time intrusion-detection framework capable of identifying jammer behaviour from directional network metrics before communication quality collapses completely.

Such a system could eventually support adaptive mitigation strategies, including channel switching, routing adaptation, transmit-power adjustment, or automatic jammer localization.

A long-term extension could also integrate the detection framework directly into LoRa gateways or edge devices, allowing real-time monitoring and response without requiring centralized offline analysis.

The overall objective is therefore not only to study jammer behaviour, but also to develop practical tools for improving the resilience, observability, and security of future low-power wireless IoT deployments.

References

- [1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “Wireless sensor networks: a survey,” *Computer Networks*, vol. 38, no. 4, pp. 393–422, 2002. Available: <https://www.sciencedirect.com/science/article/pii/S1389128601003024>
- [2] IEEE, *IEEE Standard for Low-Rate Wireless Networks*, IEEE Std 802.15.4, 2020. Available: https://standards.ieee.org/standard/802_15_4-2020.html
- [3] T. Winter et al., “RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks,” RFC 6550, IETF, 2012. Available: <https://datatracker.ietf.org/doc/html/rfc6550>
- [4] A. D. Wood and J. A. Stankovic, “Denial of service in sensor networks,” *Computer*, vol. 35, no. 10, pp. 54–62, 2002. Available: <https://ieeexplore.ieee.org/document/1039518>
- [5] C. Karlof and D. Wagner, “Secure routing in wireless sensor networks: attacks and countermeasures,” *Ad Hoc Networks*, vol. 1, no. 2–3, pp. 293–315, 2003. Available: <https://www.sciencedirect.com/science/article/pii/S1570870503000088>
- [6] W. Xu, W. Trappe, Y. Zhang, and T. 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*, 2005, pp. 46–57. Available: <https://dl.acm.org/doi/10.1145/1062689.1062697>
- [7] S. Raza, L. Wallgren, and T. Voigt, “SVELTE: Real-time intrusion detection in the Internet of Things,” *Ad Hoc Networks*, vol. 11, no. 8, pp. 2661–2674, 2013. Available: <https://www.sciencedirect.com/science/article/pii/S1570870513000801>
- [8] LoRa Alliance, *LoRaWAN Specification*. Available: https://lora-alliance.org/resource_hub/lorawan-specification-v1-1/
- [9] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, “Long-range communications in unlicensed bands: the rising stars in the IoT and smart city scenarios,” *IEEE Wireless Communications*, vol. 23, no. 5, pp. 60–67, 2016. Available: <https://ieeexplore.ieee.org/document/7721743>
- [10] Semtech Corporation, *SX1261/SX1262 Data Sheet*. Available: <https://www.semtech.com/products/wireless-rf/loro-connect/sx1262>
- [11] M. Bor, U. Roedig, T. Voigt, and J. Alonso, “Do LoRa low-power wide-area networks scale?” in *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2016, pp. 59–67. Available: <https://dl.acm.org/doi/10.1145/2988287.2989163>
- [12] D. Croce, M. Gucciardo, I. Tinnirello, D. Garlisi, and S. Mangione, “Impact of LoRa imperfect orthogonality: Analysis of link-level performance,” *IEEE Communications Letters*, vol. 22, no. 4, pp. 796–799, 2018. Available: <https://ieeexplore.ieee.org/document/8264630>
- [13] F. Adelantado, X. Vilajosana, P. Tuset-Peiro, B. Martinez, J. Melia-Segui, and T. Watteyne, “Understanding the limits of LoRaWAN,” *IEEE Communications Magazine*, vol. 55, no. 9, pp. 34–40, 2017. Available: <https://ieeexplore.ieee.org/document/>

- [14] C. E. Shannon, "Communication in the presence of noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949. Available: <https://ieeexplore.ieee.org/document/1697831>
- [15] A. Goldsmith, *Wireless Communications*. Cambridge University Press, 2005. Available: <https://www.cambridge.org/core/books/wireless-communications/800441E8C8B0256294956E71B719D1A3>
- [16] M. Savva, I. Vassiliou, and V. Vassiliou, "Fuzzy-logic based IDS for detecting jamming attacks in wireless mesh IoT networks," in *Proceedings of the 20th Mediterranean Communication and Computer Networking Conference (MedComNet)*, 2022, p. 19.
- [17] I. Armakolas, M. I. Savva, M. Savva, I. Vassiliou, and V. Vassiliou, "Location-enabled IoT (LE-IoT): Indoor localization for IoT environments using machine learning," 2022.
- [18] M. Savva, I. Ioannou, and V. Vassiliou, "Performance evaluation of a fuzzy logic-based IDS (FLIDS) technique for the detection of different types of jamming attacks in IoT networks," in *Proceedings of the 21st Mediterranean Communication and Computer Networking Conference (MedComNet)*, 2023, p. 9.
- [19] I. Armakolas, M. Savva, M. Raspopoulos, C. Christoforou, and V. Vassiliou, "Revolutionising IoT network security by assessing ML localisation techniques against jamming attacks," in *Proceedings of the 22nd Mediterranean Communication and Computer Networking Conference (MedComNet)*, 2024, p. 5.
- [20] M. C. Savva, "A framework for the detection, localization, and recovery from jamming attacks in the Internet of Things," *University of Cyprus*, vol. 2, no. 1, pp. 3–12, 2024.
- [21] M. I. Savva, I. Ioannou, and V. Vassiliou, "Evaluating localization algorithms in IoT networks under jamming attacks," in *Proceedings of the 2024 IFIP Networking Conference (IFIP Networking)*, 2024, pp. 627–633.
- [22] M. Savva, I. Ioannou, and V. Vassiliou, "Fuzzy logic-based IDS (FLIDS) for the detection of different types of jamming attacks in IoT networks," *Computer Communications*, vol. 241, p. 108251, 2025.
- [23] M. Savva, I. Ioannou, and V. Vassiliou, "Detecting multiple jammers using fuzzy-logic intrusion detection system (FLIDS)," in *Proceedings of the 2024 20th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*, 2024.
- [24] M. Savva, I. Ioannou, and V. Vassiliou, "Recovery RPL Internet of Things network under the jamming attack," in *Proceedings of the 2025 21st International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*, 2025.

APPENDICES

APPENDIX I

Repository Files

The repository is organized into firmware, jammer, experiment, analysis, and visualization components.

The main files and directories used in this thesis are:

- `draft.tex`: main thesis source file;
- `tests/test1`: 30-minute no-jammer baseline capture;
- `tests/test2`: second 30-minute no-jammer baseline capture;
- `tests/test3`: single-run constant-jammer capture;
- `tests/test4`: single-run constant-jammer capture at a different jammer placement;
- `tests/test5`: single-run constant-jammer capture at a third jammer placement;
- `tests/test6`: single-run reactive-jammer capture near node 5;
- `tests/analysis/`: baseline analysis output directory;
- `tests/analysis/link_metrics.csv`: derived directional link metrics for the baseline analysis;
- `tests/analysis/device_summary.csv`: derived device-level metrics for the baseline analysis;
- `tests/analysis/plots/separated/test1/`: separated baseline plots for `test1`;
- `tests/analysis/plots/separated/test2/`: separated baseline plots for `test2`;
- `tests/analysis/plots/separated/cross_test/`: baseline comparison plots between `test1` and `test2`;
- `tests/analysis/plots/placements/`: 2D and 3D topology and jammer-placement visualizations;
- `constant_jammer_output/test3`: repeated constant-jammer capture run;
- `constant_jammer_output/test4`: repeated constant-jammer capture run;
- `constant_jammer_output/test5`: repeated constant-jammer capture run;
- `constant_jammer_output/test6`: repeated constant-jammer capture run;
- `constant_jammer_output/test7`: repeated constant-jammer capture run;
- `constant_jammer_output/test8`: repeated constant-jammer capture run;
- `constant_jammer_output/analysis/`: averaged repeated-capture analysis output;
- `constant_jammer_output/analysis/plots/heatmaps/`: averaged jammer heatmaps;
- `constant_jammer_output/analysis/plots/jammer_impact/`: averaged jammer-versus-baseline delta plots;
- `constant_jammer_output/analysis/plots/proximity/`: jammer-distance correla-

tion plots;

- `constant_jammer_output/analysis/plots/node5_focus/`: focused node 5 jammer analysis plots;
- `deceptive_jammer_output/`: repeated deceptive-jammer campaign output;
- `deceptive_jammer_output/analysis/`: deceptive-jammer analysis and averaged metrics;
- `deceptive_jammer_output/analysis/plots/separated/`: per-test deceptive-jammer topology plots;
- `deceptive_jammer_output/analysis/plots/topology/`: aggregated deceptive-jammer network topology;
- `random_jammer_output/`: repeated random-jammer campaign output;
- `random_jammer_output/analysis/`: random-jammer analysis and averaged metrics;
- `random_jammer_output/analysis/plots/separated/`: per-test random-jammer topology plots;
- `random_jammer_output/analysis/plots/topology/`: aggregated random-jammer network topology;
- `reactive_jammer_output/`: repeated reactive-jammer campaign output;
- `reactive_jammer_output/analysis/`: reactive-jammer analysis and averaged metrics;
- `reactive_jammer_output/analysis/plots/separated/`: per-test reactive-jammer topology plots;
- `reactive_jammer_output/analysis/plots/topology/`: aggregated reactive-jammer network topology;
- `scripts/analyze_tests.py`: main packet-analysis pipeline for directional metrics;
- `scripts/generate_test_graphs.py`: graph generation pipeline for baseline and single-run experiments;
- `scripts/generate_repeated_test_graphs.py`: repeated-capture averaging and visualization pipeline;
- `scripts/generate_device_placement_3d.py`: topology and 3D placement visualization script;
- LilyGo static mesh `main.ino`: node telemetry, SD logging, and periodic broadcast scheduling firmware;
- LilyGo static mesh `Mesh/mesh.h`: SX126x LoRa radio configuration and mesh parameters;
- JamRF repository, `HackRF/jamRF_v1.py`: HackRF/GNU Radio jammer implementation;
- JamRF repository, `HackRF/config_v1.yaml`: jammer configuration used during the repeated constant-jammer campaign;
- per-node CSV logs:

- `tx.csv`: transmitted packet logs;
- `rx.csv`: received packet logs;
- `all.csv`: combined packet logs.

APPENDIX II

Additional Graphs

Additional generated graphs include separated heatmaps, raw RSSI/SNR timelines, per-receiver SNR timelines, device placement figures, distance-versus-jammer-impact plots, and expected-versus-received packet plots. Most per-test figures are grouped under the separated `test1` through `test6` plot folders. Cross-test figures are grouped under the separated cross-test folder.

APPENDIX III

Reproducibility Commands

```
python3 scripts/analyze_tests.py  
.venv/bin/python scripts/generate_test_graphs.py  
.venv/bin/python scripts/generate_device_placement_3d.py
```

APPENDIX IV

Constant Jammer Analysis (Averaged)

IV.1 Overview

This appendix presents the results of the repeated constant-jammer campaign, where the same constant jammer was placed 1 metre in front of node 5 across multiple 3.5–4 minute short captures. The per-run data has been averaged to produce higher-confidence estimates than any single long-duration run.

The constant jammer operates by continuously transmitting on the LoRa frequency, occupying the channel regardless of whether legitimate traffic is present. This produces sustained, predictable degradation patterns localized around the jammer position.

IV.2 Experimental Setup

- Jammer position: 1 metre directly in front of node 5;
- Jammer type: continuous on-channel transmission;
- Capture duration per run: 3.5–4 minutes;
- Number of runs: 8 (stored in `test1` through `test8` subdirectories);
- Analysis method: per-run extraction, baseline cropping, and per-link averaging.

IV.3 Analysis Output

The averaged analysis files are located in `constant_jammer_output/analysis/`:

- `avg_device_summary.csv`: per-device metrics averaged across all runs;
- `avg_link_metrics.csv`: per-directional-link metrics averaged across all runs;
- `plots/device_bars/`: averaged bar charts showing per-device PDR, RSSI, and SNR;
- `plots/heatmaps/`: spatially-aware heatmap visualizations of link degradation;
- `plots/jammer_impact/`: before/after comparison plots showing jammer effect magnitude;
- `plots/node5_focus/`: detailed analysis of node 5 and its immediate neighbors;
- `plots/proximity/`: scatter plots correlating physical distance to jammer with performance degradation.

IV.4 Key Observations

IV.4.1 Spatial Degradation

The constant jammer produces highly localized degradation centered on node 5. Devices closer to node 5 show significantly reduced PDR and higher error rates, while more distant devices often maintain near-baseline communication quality.

The heatmap plots clearly show the spatial structure of the interference: strong degradation in the immediate vicinity of the jammer and node 5, with gradual recovery as distance increases.

IV.4.2 Directional Asymmetry

Even at similar distances from the jammer, different directional links experience significantly different degradation. This asymmetry reflects indoor propagation effects such as wall attenuation, reflections, and antenna orientation.

Some links maintain nearly baseline performance despite being physically near node 5, while others show severe degradation. This indicates that the indoor environment strongly influences whether a particular transmitter-receiver pair is affected by the jammer.

IV.4.3 RSSI and SNR Trends

The received signal strength indicator (RSSI) shows modest degradation, suggesting that the jammer may not be immediately adjacent to or aligned with all receivers.

The signal-to-noise ratio (SNR), however, shows more substantial degradation, confirming that the jammer introduces interference power that raises the noise floor even if the desired signal power is not substantially reduced.

IV.4.4 Consistency Across Runs

Averaging across 8 runs demonstrates that the degradation pattern is consistent and reproducible. The per-run variability is relatively small compared to the overall jammer effect, suggesting that the constant jammer produces a stable and predictable interference signature.

IV.5 Comparison to Baseline

When compared to the baseline tests (test1 and test2 from the primary analysis), the constant-jammer campaign shows:

- PDR reduction of 20–60% for links directly affected by the jammer;
- SNR degradation of 5–15 dB for impacted receivers;
- Minimal change in unaffected distant links;
- Localized spatial impact rather than network-wide degradation.

IV.6 Device Summary Statistics

The averaged device summaries show per-device metrics including transmitted sequence counts, received sequence counts across all incoming links, average RSSI, average SNR, and reception window statistics. These provide a device-level view of the overall communication quality at each node during the jamming campaign.

IV.7 Link-Level Analysis

The per-link metrics include both count-based statistics (unique received/transmitted pairs) and quality metrics (RSSI, SNR, reception windows). The best-performing links under jamming conditions are documented alongside their distance estimates where available.

APPENDIX V

Deceptive Jammer Analysis

V.1 Overview

This appendix presents the analysis of the repeated deceptive-jammer campaign. The deceptive jammer is designed to mimic or spoof legitimate LoRa traffic, rather than simply transmitting noise or blocking the channel.

By impersonating network packets or generating traffic that appears legitimate to receiver hardware, a deceptive jammer can potentially cause packet rejection at higher layers, force retransmissions, or confuse network protocols without necessarily reducing the observed received signal power.

The repeated captures follow the same methodology as the constant-jammer campaign: multiple 3.5–4 minute runs with per-run extraction, baseline cropping, and per-link averaging.

V.2 Experimental Setup

- Jammer position: 1 metre directly in front of node 5;
- Jammer type: deceptive transmission mimicking LoRa modulation;
- Capture duration per run: 3.5–4 minutes;
- Number of runs: repeated captures (stored in test1, test2, test3 subdirectories);
- Analysis method: per-run extraction, baseline cropping, and per-link averaging.

V.3 Analysis Output

The analysis files are located in `deceptive_jammer_output/analysis/`:

- `device_summary.csv`: per-device metrics aggregated across runs;
- `link_metrics.csv`: per-directional-link metrics;
- `per_run_device_summary.csv`: device metrics for each individual run;
- `per_run_link_metrics.csv`: link metrics broken down by individual run;
- `plots/separated/`: individual topology and link plots for each test;
- `plots/topology/`: aggregated network topology visualizations;
- `summary.md`: detailed analysis summary with best-performing links and device statistics.

V.4 Key Observations

V.4.1 Deceptive vs. Constant Jammer Contrast

The deceptive jammer shows distinctly different characteristics compared to the constant jammer. Rather than continuous channel occupation, the deceptive jammer generates traffic that appears to receivers as legitimate packets.

This can produce different observable effects:

- Lower average RSSI degradation (since the jammer is not continuously present);
- More variable SNR degradation (correlated with deceptive transmission timing);
- Potential for collision-based packet loss if deceptive transmissions overlap with legitimate traffic;
- Possible receiver confusion or state-machine disruption if deceptive packets pass basic checksum or CRC validation.

V.4.2 Spatial Distribution

Like the constant jammer, the deceptive jammer produces localized effects around its placement (1 metre in front of node 5). However, the spatial degradation pattern may differ due to the bursty nature of deceptive transmissions compared to continuous jamming.

The topology plots show the directional structure of the network and highlight which links are most severely affected by deceptive interference.

V.4.3 Per-Run Variability

The per-run metrics allow assessment of consistency across captures. The deceptive jammer may show higher run-to-run variability if its operation is non-deterministic or if it generates traffic at random intervals.

Conversely, if the deceptive jammer follows a fixed schedule or pattern, the runs should show similar degradation signatures, enabling easier detection.

V.4.4 Device-Level Impact

The device summaries show how each node is affected overall. Devices with many incoming links from other nodes may show different aggregate statistics than devices that receive fewer transmissions.

Node 5, being near the jammer, is expected to show reduced reception quality on most of its incoming links. More distant devices may show minimal change.

V.5 Detection Considerations

The deceptive jammer presents unique detection challenges:

- Lower received power degradation makes it harder to detect via RSSI alone;
- Legitimate-looking traffic may bypass some filtering rules;
- Collision detection requires temporal correlation analysis;
- Effective detection may require multi-metric fusion combining PDR, packet timing, and sequence analysis.

V.6 Comparison Summary

The best-performing and worst-affected links are documented in the analysis summary. These provide quantitative benchmarks for comparing deceptive jammer effects to baseline, constant, random, and reactive jammer campaigns.

APPENDIX VI

Random Jammer Analysis

VI.1 Overview

This appendix presents the analysis of the repeated random-jammer campaign. The random jammer transmits interference at stochastic or pseudo-random intervals, intervals, or with randomized characteristics such as frequency hopping, waveform modulation, or burst timing.

By introducing unpredictability into the interference pattern, a random jammer can be more difficult to detect via simple threshold-based algorithms or time-correlated analysis. The stochastic interference may also be harder to mitigate through deterministic retransmission or channel-adaptation strategies.

The repeated captures follow the same methodology as the constant-jammer and deceptive-jammer campaigns: multiple 3.5–4 minute runs with per-run extraction, baseline cropping, and per-link averaging.

VI.2 Experimental Setup

- Jammer position: 1 metre directly in front of node 5;
- Jammer type: randomized or stochastic interference transmission;
- Capture duration per run: 3.5–4 minutes;
- Number of runs: repeated captures (stored in test1, test2, test3 subdirectories);
- Analysis method: per-run extraction, baseline cropping, and per-link averaging.

VI.3 Analysis Output

The analysis files are located in `random_jammer_output/analysis/`:

- `device_summary.csv`: per-device metrics aggregated across runs;
- `link_metrics.csv`: per-directional-link metrics;
- `per_run_device_summary.csv`: device metrics for each individual run;
- `plots/separated/`: individual topology and link plots for each test;
- `plots/topology/`: aggregated network topology visualizations;
- `summary.md`: detailed analysis summary with best-performing links and device statistics.

VI.4 Key Observations

VI.4.1 Variability and Unpredictability

The random jammer produces interference whose characteristics vary over time. This stochasticity manifests as:

- Higher run-to-run variability in PDR and SNR compared to the constant jammer;
- Time-dependent degradation patterns that do not repeat identically across runs;
- Potential for temporary reprieve or recovery during periods of low jammer activity;
- Degradation that cannot be reliably predicted or modeled via simple fixed parameters.

VI.4.2 Statistical Degradation

Despite the unpredictability of individual interference events, the averaged metrics across multiple runs should reveal consistent statistical degradation. Links near the jammer show average PDR and SNR reductions even though instantaneous reception quality may fluctuate significantly.

This statistical consistency enables detection and analysis via aggregate metrics, even when individual packets or brief time windows show near-baseline quality.

VI.4.3 Spatial Effects

Like both the constant and deceptive jammers, the random jammer is deployed 1 metre in front of node 5. The spatial degradation pattern should show localized effects near node 5 with gradual recovery at distance.

However, the high variability may make the spatial pattern less pronounced or harder to visualize clearly compared to the constant-jammer baseline.

VI.4.4 Per-Run Statistics

The per-run metrics reveal how much variation occurs between individual captures. Large differences between runs indicate that the random jammer's activity or interference strength varies significantly. Smaller differences suggest that despite the randomized transmission pattern, the average effect is relatively stable across time.

VI.5 Detection Characteristics

The random jammer presents distinct detection challenges:

- Higher variance in individual measurements requires statistical filtering or multi-packet averaging for reliable detection;
- Simple threshold-based detection may produce higher false-positive or false-negative rates;
- Temporal correlation analysis is less effective when interference timing is unpredictable;

- Effective detection requires accumulating sufficient measurement history to achieve statistical significance.

VI.5.1 Comparison to Other Jammer Types

The random jammer's statistical behavior contrasts sharply with the constant jammer's predictable, continuous degradation. The random jammer may be easier to confuse with natural wireless fading or multipath effects, making it harder to distinguish from environmental noise without sophisticated analysis.

VI.6 Comparison Metrics

The best-performing and worst-affected links are documented in the analysis summary. These quantitative metrics enable direct comparison of random-jammer effects to baseline, constant, deceptive, and reactive jammer campaigns.

APPENDIX VII

Reactive Jammer Analysis

VII.1 Overview

This appendix presents the analysis of the repeated reactive-jammer campaign. The reactive jammer detects legitimate network activity and responds by transmitting interference, typically on the same channel or frequency.

Unlike the constant jammer which continuously occupies the channel, the reactive jammer transmits only when it senses or is triggered by network traffic. This creates time-limited, burst-like interference correlated with legitimate transmissions.

The reactive jammer is particularly relevant to security analysis because it represents a targeted attack that consumes less power than continuous jamming while still causing significant communication disruption.

The repeated captures follow the same methodology as other jammer campaigns: multiple 3.5–4 minute runs with per-run extraction, baseline cropping, and per-link averaging.

VII.2 Experimental Setup

- Jammer position: 1 metre directly in front of node 5;
- Jammer type: reactive/triggered interference transmission;
- Jammer trigger: response to detected network activity;
- Capture duration per run: 3.5–4 minutes;
- Number of runs: repeated captures (stored in test1, test2, test3 subdirectories);
- Analysis method: per-run extraction, baseline cropping, and per-link averaging.

VII.3 Analysis Output

The analysis files are located in `reactive_jammer_output/analysis/`:

- `device_summary.csv`: per-device metrics aggregated across runs;
- `link_metrics.csv`: per-directional-link metrics;
- `per_run_device_summary.csv`: device metrics for each individual run;
- `per_run_link_metrics.csv`: link metrics broken down by individual run;
- `plots/separated/`: individual topology and link plots for each test;
- `plots/topology/`: aggregated network topology visualizations;
- `summary.md`: detailed analysis summary with best-performing links and device statistics.

VII.4 Key Observations

VII.4.1 Temporal Correlation with Traffic

The reactive jammer's interference is temporally correlated with network activity. This manifests as:

- Degradation concentrated around times when the network is actively transmitting;
- Potential for quiet periods with near-baseline quality if the network is idle;
- Collision-based packet loss when jammer response coincides with legitimate transmissions;
- Variable aggregate PDR that depends on network duty cycle and jammer detection latency.

VII.4.2 Detection Latency Effects

If the reactive jammer has significant detection latency before generating its response, brief transmissions may complete before interference arrives. This can result in:

- Reduced jamming effectiveness against short packets;
- Partial recovery or escape windows where traffic succeeds despite jamming;
- Time-dependent degradation patterns that vary based on packet length and jammer response time.

VII.4.3 Power Efficiency

Compared to the constant jammer which transmits continuously, the reactive jammer transmits only during or in response to network activity. This may result in:

- Lower average jamming power consumption;
- More difficult detection via simple power-consumption monitoring;
- Potential for sustainable long-term deployment with limited battery resources;
- Variable thermal signature and RF exposure compared to continuous jamming.

VII.4.4 Spatial and Directional Effects

Like other jammer types, the reactive jammer is deployed 1 metre in front of node 5. The spatial degradation pattern should show localized effects, though the temporal correlation with traffic may introduce additional complexity.

The topology plots reveal which directional links are most severely affected and how the reactive jammer's interference propagates through the network topology.

VII.4.5 Consistency and Reproducibility

The reactive jammer should show consistent operation across multiple runs, assuming the detection and triggering mechanism is stable. The per-run metrics reveal how reproducible

the reactive jammer's behavior is and whether external factors (network activity patterns, environmental conditions) significantly affect the jammer's effectiveness.

VII.5 Detection Characteristics

The reactive jammer presents unique detection opportunities and challenges:

- Temporal correlation with traffic enables detection via time-windowed analysis or correlation-based methods;
- Burst-like interference signature differs from continuous jamming;
- Detection may rely on identifying abnormal collision patterns or unexpected timing correlations;
- Power monitoring can potentially identify reactive jammer operation based on interference bursts aligned with transmissions.

VII.6 Comparison to Other Jammer Types

The reactive jammer differs fundamentally from constant, random, and deceptive jammers:

- Unlike the constant jammer, does not continuously occupy the channel;
- Unlike the random jammer, interference is deterministically triggered rather than stochastic;
- Unlike the deceptive jammer, does not attempt to mimic legitimate traffic;
- Represents a power-efficient targeted attack strategy with temporal correlation to network activity.

VII.7 Comparison Metrics

The best-performing and worst-affected links are documented in the analysis summary. These quantitative metrics enable direct comparison of reactive-jammer effects to baseline, constant, deceptive, and random jammer campaigns.

The effectiveness of reactive jamming at reducing network throughput and communication reliability is compared against other jammer strategies to assess relative attack severity.

APPENDIX VIII

Multi-Jammer Comparison

VIII.1 Overview

This appendix presents a comprehensive cross-jammer comparison analyzing the differential effects of constant, deceptive, random, and reactive jamming strategies on the same nine-node indoor LoRa network.

All four jammer types were deployed at the same location (1 metre in front of node 5) using the same repeated short-capture methodology. This enables direct, controlled comparison of how different interference strategies affect directional communication, network metrics, and detection signatures.

VIII.2 Jammer Strategies

VIII.2.1 Constant Jammer

Continuous on-channel transmission occupying the channel without regard to network activity.

- Interference: continuous and predictable;
- Power consumption: maximum (sustained transmission);
- Detection difficulty: low (obvious signal degradation);
- Effectiveness: high but obvious.

VIII.2.2 Deceptive Jammer

Mimic or spoof legitimate LoRa traffic, generating packets that appear legitimate to receiver hardware.

- Interference: mimicked legitimate-looking traffic;
- Power consumption: variable (depends on spoofed traffic rate);
- Detection difficulty: medium (requires deep-packet inspection or semantic analysis);
- Effectiveness: depends on protocol-layer vulnerabilities.

VIII.2.3 Random Jammer

Stochastic or pseudo-random interference with unpredictable timing, frequency, or waveform characteristics.

- Interference: unpredictable and time-varying;
- Power consumption: variable;

- Detection difficulty: high (randomness can mimic natural fading);
- Effectiveness: requires sufficient averaging for reliable detection.

VIII.2.4 Reactive Jammer

Detect and respond to network activity by transmitting interference bursts.

- Interference: temporally correlated with network activity;
- Power consumption: low to medium (on-demand transmission);
- Detection difficulty: medium (requires temporal correlation analysis);
- Effectiveness: depends on detection latency and trigger reliability.

VIII.3 Comparative Metrics

VIII.3.1 Packet Delivery Ratio (PDR)

The primary metric for assessing jamming impact. Comparison reveals which jammer types produce the most severe PDR degradation at node 5 and nearby receivers.

- Constant jammer: expected to produce sustained, high-magnitude PDR reduction;
- Deceptive jammer: may produce moderate PDR reduction if protocol-layer filtering succeeds;
- Random jammer: produces statistical PDR reduction proportional to jammer duty cycle and power;
- Reactive jammer: produces PDR reduction correlated with network activity.

VIII.3.2 Signal-to-Noise Ratio (SNR)

SNR degradation reveals whether the jammer raises the interference noise floor (SNR reduction) or primarily causes collisions (SNR may be less affected).

- Constant jammer: high SNR degradation (continuous noise floor elevation);
- Deceptive jammer: variable SNR degradation depending on spoofed packet power levels;
- Random jammer: moderate to high SNR degradation if transmission duty cycle is significant;
- Reactive jammer: SNR degradation correlated with jammer transmissions.

VIII.3.3 Received Signal Strength (RSSI)

RSSI shows whether jammers are introducing power into the channel or primarily causing collisions.

- Constant jammer: may show modest RSSI degradation if jammer is not perfectly aligned with receivers;

- Deceptive jammer: RSSI depends on spoofed packet power and modulation;
- Random jammer: variable RSSI degradation;
- Reactive jammer: RSSI changes correlate with jammer burst timing.

VIII.3.4 Spatial Extent

All jammer types are deployed at the same location. Comparison of spatial degradation patterns reveals whether different jamming strategies affect the network topology differently.

- Links near the jammer: all types show degradation;
- Links at intermediate distances: degradation magnitude depends on jammer type;
- Distant links: may show minimal degradation regardless of jammer type.

VIII.4 Detection Signatures

VIII.4.1 Constant Jammer Signature

- Continuous degradation in time;
- Localized spatial impact;
- Consistent across repeated captures;
- High SNR degradation;
- Easily detectable via aggregate statistics.

VIII.4.2 Deceptive Jammer Signature

- Variable impact depending on protocol layer filtering;
- May show unexpected packet-loss patterns if spoofed packets successfully reach network layer;
- Requires deep-packet analysis for confident detection;
- Potentially lower power-consumption footprint;
- Difficult to distinguish from legitimate traffic without semantic analysis.

VIII.4.3 Random Jammer Signature

- High temporal variability;
- Run-to-run inconsistency in specific packets or brief windows;
- Statistical consistency in averaged metrics;
- May be confused with natural fading or environmental noise;
- Requires sufficient measurement history for reliable detection.

VIII.4.4 Reactive Jammer Signature

- Temporal correlation with network activity;
- Burst-like interference pattern;
- Reproducible and deterministic once detected;
- Unique timing signature correlating interference with transmissions;
- Detectable via temporal correlation analysis.

VIII.5 Optimal Detection Metrics

VIII.5.1 For Constant Jamming

Simple threshold-based detection on PDR, SNR, or RSSI. Aggregate statistics reliably identify interference.

VIII.5.2 For Deceptive Jamming

Deep-packet inspection, packet header validation, and semantic protocol analysis. Statistical approaches less effective.

VIII.5.3 For Random Jamming

Statistical multi-packet averaging, outlier filtering, and accumulation of measurement history. Requires longer observation windows than constant jamming.

VIII.5.4 For Reactive Jamming

Temporal correlation analysis, burst detection, and timing pattern matching. Correlation between network activity and degradation provides strong signature.

VIII.6 Recommendations

VIII.6.1 Multi-Metric Fusion

Optimal detection combines multiple metrics (PDR, SNR, RSSI, packet timing, reception counts) rather than relying on a single threshold. Different jammer types have different metric-level signatures.

VIII.6.2 Adaptive Detection Strategy

The most robust detection strategy combines:

- Initial classification via jammer-type signature identification;
- Application of jammer-type-specific detection methods;
- Continuous refinement as additional measurements accumulate.

VIII.6.3 Machine Learning Opportunity

The completed multi-jammer dataset enables supervised machine-learning classification between jammer types, positions, and relative intensity levels. A classifier trained on this dataset can automatically identify which jammer type is active based on communication metrics alone.

VIII.7 Summary

The four jammer types produce distinct, measurable signatures on the network. Constant jamming is highly visible and easy to detect. Deceptive, random, and reactive jamming each require specialized detection approaches tailored to their unique operational characteristics. Effective network defense requires understanding these differences and applying appropriate detection strategies for each threat model.