

Acoustic Backscatter Communication in Metal Structures

Vom Promotionsausschuss der
Technischen Universität Hamburg

zur Erlangung des akademischen Grades

Doktor-Ingenieur (Dr.-Ing.)

genehmigte Dissertation (Monografie)

von

Peter Oppermann

aus

Osnabrück, Deutschland

2024

Author ORCID:

<https://orcid.org/0000-0001-6936-2307>

Thesis DOI:

<https://doi.org/10.15480/882.9646>

Date of Oral Examination | April 3rd, 2024

Chair of Examination Board | Prof. Dr.-Ing. Görschwin Fey
Institute of Embedded Systems
Hamburg University of Technology

First Examiner | Prof. Dr.-Ing. Bernd-Christian Renner
Institute for Autonomous Cyber-Physical Systems
Hamburg University of Technology

Second Examiner | Prof. Dr.-Ing. Gerhard Bauch
Institute of Communications
Hamburg University of Technology

Acknowledgment

I extend my deepest gratitude to my supervisor, Prof. Christian Renner, for his unwavering guidance and support over the past five years. His mentorship has been instrumental in my development as a researcher. I look forward to continuing my journey at the Institute for Autonomous Cyber-Physical Systems.

My heartfelt appreciation goes to the diligent examination board for their meticulous review and evaluation of my thesis. The final examination provided a platform for stimulating discussions and inspired many new ideas and insights.

I am grateful to my colleagues at TUHH, whose contributions have been invaluable. Special thanks to Prof. Marcus Rutner and Lennart Dorendorf from the I³-Project VAM for their tremendous assistance and insightful feedback. Collaborations with Prof. Robert Meißner and Benjamin Boll have opened up exciting new avenues for research.

I am thankful for the camaraderie and intellectual exchange shared with my colleagues and fellow PhD students at the former smartPORT group and the Institute of Autonomous Cyber-Physical Systems. Our shared experiences have been marked by stimulating discussions and memorable celebrations.

I also wish to express my gratitude to Prof. Volker Turau for welcoming me to the Institute of Telematics after my previous lab closed. I am especially thankful to Christoph Weyer for his invaluable mentorship during this period.

Finally, my deepest thanks go to my wife, Sonja, for her unwavering support and encouragement. Her motivation and strength have been a constant source of inspiration during challenging times.

Abstract

Structural Health Monitoring promises to limit maintenance costs and increase the safety of new and aging infrastructure, such as pipelines, bridges, or ship and airplane hulls. Many sensors must be distributed over structures to efficiently fuel monitoring and predictive maintenance applications with the essential data. Unfortunately, many locations in structures are hard to reach and shielded by metal, impeding ready-to-use radio frequency-based wireless communication techniques and complicating regular battery replacement. Alternatively, acoustically communicating with and wirelessly transferring power to sensor nodes can penetrate metal barriers, enabling in-structure networks to run maintenance-free for decades. However, several shortcomings hinder widespread adoption: Passive backscatter communication—known for its ultra-low-power requirements—has only been demonstrated over a few centimeters between coaxially aligned transducers, restricting its usefulness. Moreover, transmission over longer distances suffers from severe reverberation, limiting data rates with conventional methods to only a few hundred bits per second. Lastly, existing approaches were often tailored to specific structures, requiring tedious setup by experts, and were only evaluated with costly, bulky, and power-hungry laboratory equipment.

This dissertation builds the foundation for batteryless and acoustically-powered sensors embedded in structures, taking a holistic approach, including active downlink, passive uplink, and wireless power transfer. The feasibility of backscatter communication from a sensor node to a readout device over more than 3 m distance is demonstrated for the first time—showing potential extensions to tens of meters. Higher-order modulation schemes and equalization in the non-linear backscatter channel are analyzed and demonstrated to increase data rates several-fold despite reverberation. These methods require knowledge about the specific channel. Plug-and-play-like deployments are realized by deriving channel models and proposing in-situ estimation procedures for parameter estimation. Similarly, a wideband modulation scheme is proposed and compared with commonly used narrowband schemes to speed up communication from the active readout device to the passive sensor, increasing robustness and alleviating the dependency on suitable carrier frequencies. Furthermore, the strong coupling between structural and electrical impedances is identified as a significant limitation for wireless power transfer, which requires impedance matching. We investigate using maximum power point tracking (MPPT) to maximize efficiency without manual adaption per deployment.

The proposed schemes are evaluated with simulation and practical experiments, for which custom prototypes were constructed from commercially off-the-shelf (COTS) components. These reveal practical non-idealities that severely impact communication reliability, e. g., inaccurate and unstable clocks and constrained computational resources. Estimation, tracking, and correction strategies for inaccurate oscillators are proposed and analyzed. The prototypes are published under an open-source license, enabling other researchers to explore acoustic in-structure sensor networks in diverse applications.

Table of Contents

List of Figures	vii
List of Tables	ix
List of Symbols	xi
List of Acronyms	xiii
1 Introduction	1
1.1 Contributions	3
1.2 Organization	4
2 Fundamentals	7
2.1 Notation	7
2.2 Wireless Communication	7
2.2.1 Encoding and Decoding	8
2.2.2 Mapping and Constellations	9
2.2.3 Equalization	14
2.2.4 Synchronization	15
2.3 Piezoacoustic Backscatter	20
2.3.1 Piezoelectric Transducers	20
2.3.2 Reflection Coefficient	21
2.4 Wireless Power Transfer	23
2.4.1 Maximum Efficiency and The Maximum Power Transfer Theorem	24
2.4.2 Rectification and Voltage Multiplication	26
2.4.3 Power Management	27
2.5 Elastic Waves in Solids	27
3 State of the Art	29
3.1 Applications and Requirements	29
3.2 Backscatter Communication	31
3.2.1 Radio-Frequency Identification	32
3.2.2 RFID Physical Layer	33
3.2.3 Beyond the RFID Standards	38
3.2.4 Piezoacoustic Backscatter	39
3.3 Through-Metal Backscatter Systems	40
3.3.1 Active Communication through GW Channels	40
3.3.2 Power Transfer through GW Channels	42
3.3.3 Sandwich Plated Channel	46
4 Goals and Preliminary Analysis	47
4.1 Dissertation Goals	47
4.1.1 Contributions	48
4.2 Channel Analysis	49
4.2.1 Piezoelectric Transducer	50
4.2.2 Metal Channel	51
4.3 Summary and Discussion	55

5	Increasing the Modulation Order	57
5.1	Overview	57
5.1.1	Problem Analysis	57
5.1.2	General Approach	59
5.1.3	Structure	59
5.2	Communication System	60
5.2.1	Steady-State Channel Model	61
5.2.2	Setup Protocol	63
5.2.3	Message Transmission	65
5.3	Implementation	68
5.3.1	Reader	68
5.3.2	Tag	72
5.4	Evaluation	75
5.4.1	Setup	75
5.4.2	Dynamic Frequency Selection	75
5.4.3	Dynamic Load Picking	77
5.4.4	Transient and Measurement Noise	79
5.4.5	Transmission Speed and Reliability	80
5.4.6	Refined Channel Estimation	82
5.5	Summary and Discussion	83
6	Battling Inter-Symbol Interference	85
6.1	Overview	85
6.1.1	Problem Analysis	85
6.1.2	General Approach	87
6.1.3	Structure	87
6.2	Channel Analysis	88
6.2.1	Channel Modeling	88
6.2.2	Equivalent Linear Channel	89
6.3	Communication System	91
6.3.1	Protocol and Message Structure	91
6.3.2	Receiver Structure	92
6.3.3	Equalizer	93
6.4	Prototype	96
6.4.1	Analog Demodulator	96
6.5	Evaluation	98
6.5.1	Training Sequence and Parameters	98
6.5.2	Symbol Timing Error	100
6.5.3	Noise and Channel Coding	102
6.5.4	Achievable Data Rates	103
6.5.5	Timing Imperfections	103
6.5.6	Time-Variant Channels	105
6.6	Summary and Discussion	108
7	Wirelessly-Powered Sensor Nodes	111
7.1	Overview	111
7.1.1	Problem Analysis	111
7.1.2	General Approach	112
7.1.3	Structure	113
7.2	Wireless Power Transfer	113
7.2.1	Harvester Design	114

7.3	Backscatter Communication	116
7.3.1	Modulation Scheme	116
7.3.2	Analog Receiver	119
7.3.3	Software-Defined Receiver	120
7.4	Evaluation	122
7.4.1	Wireless Energy Transfer	123
7.4.2	Low-Power Oscillator	126
7.4.3	Tag Power Consumption	127
7.4.4	Backscatter Communication	129
7.4.5	Larger Structures	132
7.5	Summary and Discussion	134
8	Wideband Downlink Communication	137
8.1	Overview	137
8.1.1	Problem Analysis	138
8.1.2	General Approach	139
8.1.3	Structure	140
8.2	Communication System	140
8.2.1	Modulation	141
8.2.2	Receiver	144
8.3	Simulation Study and Evaluation	145
8.3.1	Wideband Pulse Characteristics	146
8.3.2	Achievable Data Rates	147
8.3.3	Samples Per Symbol	149
8.3.4	Sampling Frequency Offset	149
8.3.5	Noise Tolerance	151
8.4	Prototype Implementation	152
8.4.1	Sample Frequency Offset Mitigation	152
8.4.2	Reader	153
8.4.3	Tag Receiver	154
8.5	Real-World Evaluation	156
8.5.1	Achievable Data Rates	156
8.5.2	Wideband Power Transfer	159
8.5.3	Power Consumption and Availability	160
8.6	Summary and Discussion	162
9	Conclusion and Outlook	165
9.1	Outlook	167
	Bibliography	169
	Index	177

List of Figures

1.1	Pipeline monitoring as an exemplary application.	2
2.1	Communication system overview.	8
2.2	Common constellations.	10
2.3	Pulse shaping.	10
2.4	Illustration of the demodulator.	14
2.5	Digital PLL to track phase and frequency errors.	17
2.6	Frame synchronization.	18
2.7	Square timing error detector.	19
2.8	Timing-locked loop and derivative timing error detector.	19
2.9	Direct and inverse piezoelectric effect.	21
2.10	Transducer resonance modes.	21
2.11	Wave reflection and acoustic impedance.	22
2.12	Wireless power transfer system.	23
2.13	Impedance matching between source and load.	25
2.14	Rectifier topologies.	26
2.15	Wave types in solid media.	28
3.1	Applications for GW ultrasonic power transfer and communication.	30
3.2	Terminology of a backscatter system.	32
3.3	Downlink modulation in EPCglobal standard.	34
3.4	Uplink modulation in EPCglobal standard.	35
3.5	WISP platform's analog frontend.	37
3.6	Higher-order modulation with multiple load impedances.	38
3.7	Sandwich-plated vs. guided-wave channel.	41
4.1	Transducer-to-transducer channel.	50
4.2	Piezoelectric transducer impedance.	52
4.3	Test specimen used throughout this thesis.	52
4.4	Guided-wave channel gain.	53
4.5	Guided-wave channel RMS delay spread.	54
4.6	Guided-wave channel transducer impedance.	54
4.7	Sandwich-plated channel gain and impedance.	55
5.1	Distorted constellation in GW channels.	58
5.2	Communication system overview and notation.	60
5.3	Equivalent circuit model for GW channel.	62
5.4	Frequency calibration procedure.	64
5.5	Estimated channel function.	65
5.6	ISI and symbol duration in constellation diagram.	67
5.7	Reader printed circuit board.	68
5.8	Simplified schematic of the reader circuit.	69
5.9	Undersampling in frequency spectrum.	71
5.10	The modulator circuit of a tag.	73

List of Figures

5.11	The tag circuit board.	74
5.12	Picture of setup with tag and reader.	76
5.13	SNR versus carrier frequency.	76
5.14	Modulation range with tag device.	78
5.15	Poor generality of selected loads.	79
5.16	SINR is dominated by ISI.	80
5.17	Achievable data rate with higher-order modulation.	81
5.18	Shortest feasible symbol durations in multiple channels.	82
5.19	Channel variation over time.	82
6.1	Signal flow in backscatter channel.	86
6.2	Groups of propagation paths.	89
6.3	Equivalent linear channel.	90
6.4	Training sequence to learn steady-state values.	92
6.5	Receiver structure with equalizer.	92
6.6	MMSE equalizer structure.	93
6.7	DF equalizer structure.	94
6.8	Adaptive DFE equalizer structure.	96
6.9	Analog demodulator.	97
6.10	EVM: exemplary sample distribution.	99
6.11	Required number of equalizer taps.	99
6.12	Effect of training sequence length on EVM.	100
6.13	Effect of symbol-timing synchronization on EVM.	101
6.14	PER with and without decoding.	102
6.15	Maximum throughput.	103
6.16	Effect of clock offset on EVM.	105
6.17	Constellation point adaptation.	106
6.18	Experiment setup with a time-variant channel.	107
6.19	Adaptive equalization in time-variant channel.	108
7.1	System architecture with harvester and modulator.	113
7.2	Harvester circuit.	114
7.3	Harvester principle with MPPT and duty-cycling.	115
7.4	PSK modulation at tag.	118
7.5	Modulator circuit.	118
7.6	Analog frontend with self-interference suppression.	119
7.7	Source of sampling frequency offset.	120
7.8	Software-defined radio pipeline.	121
7.9	Potential of WPT in GW channels.	124
7.10	Characterization of harvester efficiency.	125
7.11	Stability of low-power oscillators.	127
7.12	Power consumption of tag prototype.	128
7.13	Investigation of selected channels.	130
7.14	SFO mitigation with DPLL.	130
7.15	EVM with SFO tracking and mitigation.	131
7.16	Effect of harvester decoupling.	131
7.17	Simulation study on impulse response and PLL gain.	132
7.18	WPT on larger structures.	133
8.1	Demodulation with envelope detector at large delay spreads.	139
8.2	Interdependence of harvesting and signal reception.	139

8.3	Downlink message structure.	141
8.4	Investigated continuous wave modulation schemes.	142
8.5	Wideband modulation schemes.	143
8.6	Tag receiver structure.	144
8.7	Simulation signal flow.	145
8.8	Spectral density functions.	147
8.9	Maximum data rates with different modulation schemes.	148
8.10	Narrowband data rates statistics.	148
8.11	Reliability versus the number of samples per symbol.	149
8.12	Clock frequency offset effects and extent.	150
8.13	Receiver noise tolerance.	151
8.14	SFO estimation and mitigation.	152
8.15	Reader experiment setup.	154
8.16	Modified harvester circuit.	156
8.17	Achieved data rates with the prototype.	157
8.18	Accuracy of SFO detection mechanism	158
8.19	Power transfer efficiency with wideband pulses.	160
8.20	Receiver power consumption.	161
8.21	Maximum data rate of a batteryless tag.	162

List of Tables

3.1	Platforms for computational RFID.	34
3.2	State of the art in active GW communication.	43
3.3	State of the art in GW power transfer.	45
5.1	Available digital potentiometer ICs.	73

List of Symbols

Signals	
$m[k]$	Transmitted message (bits)
$m_c[k]$	Encoded message (bits)
$c[k]$	Signal after symbol mapping
$b(t)$	Transmitted baseband signal
$x(t)$	Transmitted passband signal
$y(t)$	Received passband signal
$\hat{b}(t)$	Received baseband signal
$\hat{b}_I(t), \hat{b}_Q(t)$	In-Phase and quadrature part of received baseband signal
$\tilde{b}(t)$	Equalized baseband signal
$\hat{m}[k]$	Symbol decision at receiver
$n(t)$	Noise signal
$b_p[k]$	Transmitted preamble signal
Γ	Reflection coefficient
LLR	The log-likelihood ratio
Transceiver Parameters	
S	Modulation order, i. e. number of distinct symbols
N_b	Number of symbols per message
N_a	Number of samples in acknowledgement
N_m	Number of samples per symbol
T_m	Symbol duration
T_s, f_s	Sampling interval and rate
f_c, f_i	Carrier and intermediate frequency
ω_c, ω_i	Carrier and intermediate angular frequencies
f_{co}	Cutoff frequency of lowpass filter

Channel Description	
τ_{rms}	RMS delay spread
$h_{\text{pb}}(t), h_{\text{c}}(t)$	Pass and baseband channel impulse response
$h_{\text{eq}}(t)$	Impulse response of linear equalizer
$h_{\text{si}}(t)$	Impulse response of self-interference channel
$h_{\text{ul}}, h_{\text{dl}}, h_{\text{rs}}$	Impulse response of uplink, downlink, rescattering channel
$p(t)$	Pulse-shaping impulse response
Estimation and Tracking	
$\tilde{f}_{\text{CFO}}, f_{\text{o}}$	Estimated and actual carrier frequency offset
ϕ_{o}	Initial phase offset
α	The symbol time error
$\tilde{T}_{\text{m}}, T'_{\text{m}}$	Estimated and actual symbol duration
$\tilde{f}_{\text{i}}, f'_{\text{i}}$	Estimated and actual intermediate frequency
$\Delta\phi_{\text{i}}$	Average phase shift between consecutive samples
\tilde{a}	Amplitude estimate
Equalization	
$c(t)$	Generated carrier
y_{ss}	Steady-state convergence signal
\mathcal{Y}_{ss}	Set of steady-state constellation points
N_{T}	Number of symbols in training sequence
$\mathbf{c}, \mathbf{c}^{\text{MMSE}}, \mathbf{c}^{\text{DF}}$	Equalizer coefficients (MMSE and DF equalizer)
$M, M_{\text{f}}, M_{\text{b}}$	Number of equalizer taps (forward and feedback filter for DF)
$b[k], b^{\text{DF}}[k]$	Transmitted symbols in baseband
Electrical	
z_{l}	Load impedance at the tag
$z_{\text{r}}, \tilde{z}_{\text{r}}$	Observed impedance at the reader transducer (continuous and time-discrete)
z_{o}	Transducer characteristic impedance
$i_{\text{R}}, v_{\text{R}}$	Current and voltage through reader transducer
$C_{\text{in}}, V_{\text{in}}$	Input capacitance and voltage over input capacitor
$C_{\text{st}}, V_{\text{st}}$	Storage capacitance and voltage over storage capacitor
$E_{\text{st}}, E_{\text{MCU}}$	Stored and consumed energy

List of Acronyms

ACK	Acknowledgement
ADC	Analog to Digital Converter
AM	Amplitude Modulation
ASK	Amplitude Shift Keying
AWGN	Additive White Gaussian Noise
BP	Belief Propagation
CFO	Carrier Frequency Offset
CPA	Constellation Point Adaption
CW	Continuous Wave
DAC	Digital to Analog Converter
DF	Decision Feedback
DPLL	Digital Phase Locked Loop
EVM	Error Vector Magnitude
EWMA	Exponentially-Weighted Moving-Average
FM	Frequency Modulation
FSK	Frequency Shift Keying
GW	Guided-Wave
IC	Integrated Circuit
ISI	Inter-Symbol Interference
LDPC	Low-Density Parity-Check
LMS	Least-Mean-Square
LTI	Linear Time-Invariant
MCU	Microcontroller Unit
MMSE	Minimum Mean Square Error
MPPT	Maximum Power Point Tracking
OFDM	Orthogonal Frequency Division Multiplexing
OOK	On-/Off-Keying
PDM	Pulse-Duration Modulation
PIE	Pulse-Interval Encoding
PM	Phase Modulation
PMU	Power Management Unit
PO	Phase Offset
PPM	Pulse Position Modulation

List of Acronyms

PSK	Phase Shift Keying
QAM	Quadrature Amplitude Modulation
RFID	Radio-Frequency Identification
RMS	Root-Mean-Square
SFO	Sample Frequency Offset
SHM	Structural Health Monitoring
SINR	Signal to Interference and Noise ratio
SNR	Signal to Noise ratio
STE	Symbol Timing Error
STFT	Short-Time Fourier Transform
SWP	Sandwich-Plated
WB-OOK	Wideband On-/Off-Keying
WB-PPM	Wideband Pulse Position Modulation
WISP	Wireless Identification and Sensing Platform
WPT	Wireless Power Transfer

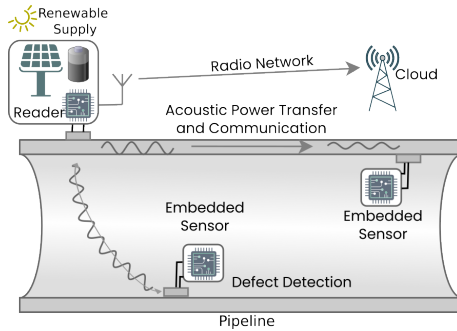
Introduction

Large parts of the infrastructure of many industrialized countries were constructed in the 1950s and 1960s—and are increasingly deteriorating today. For example, in 2021, more than 25 percent of German highway bridges received a 4 (sufficient, equivalent to a D in the ECTS/US classification system) in the German grading system [23b]. As a result, monitoring infrastructure health has become more important than ever. The current process of structural monitoring involves experts that perform routine checks in fixed intervals anywhere from three to six years [Dig22]. These checks are time-consuming, tedious, and expensive. Furthermore, the human eye often only detects faults when they have progressed to a significant extent. Hence, automated monitoring of civil infrastructure comes with significant advantages:

- It helps to identify damage early when repairs can be performed at low cost and complexity.
- It prevents potentially catastrophic failure by timely countermeasures.
- It frees experts from tedious, time-consuming, and expensive manual routine checks.

Likewise, monitoring other safety-critical structures, such as aircrafts and industrial facilities, can reduce overall maintenance costs. Continuous data from densely deployed sensors in structures is vital to fuel predictive maintenance techniques in such applications.

There are many promising technologies for detecting malfunctions and damage in various structures [Ric+22; Jia+20], but the deployment of the necessary sensors is still a challenge: Sensors must be distributed in large quantities on large structures, often in hard-to-reach or completely inaccessible places. Cable connections, on the one hand, require a large amount of material and considerably increase the deployment cost of sensor networks. Moreover, cables may even add critical weight or require maintenance by themselves, and if holes must be drilled to feed cables into enclosed spaces, the structural integrity may be compromised. On the other hand, battery-driven sensors



■ **Figure 1.1:** Pipeline monitoring is one out of many envisioned applications for acoustic through-metal power transfer and backscatter communication. An outside device (reader) wirelessly transfers power and data to one or more embedded and backscattering sensors.

require regular battery replacement, causing considerable maintenance costs and environmental impact. Sensor nodes that are wireless and battery-free, i.e., solely rely on harvested energy, can last for decades without manual maintenance and, thus, alleviate this burden.

Fueled by the Internet of Things (IoT)—the trend to connect devices in all sectors, from industry to private homes—several low-cost and low-power communication technologies have reached a mature state and are readily available. However, most of these technologies rely on radio-frequency waves. Radio-frequency waves are unfavorable, if not infeasible, in applications where sensors are shielded by or embedded in metal structures. Acoustic waves are a possible remedy to transmit both data and power wirelessly in a broad range of different applications, such as structural health monitoring, pipeline monitoring, nuclear power plant monitoring, aerospace vehicles, and cargo container monitoring. Figure 1.1 presents the example of a pipeline monitoring system similar to Radio-Frequency Identification (RFID) that uses wireless power transfer to supply a sensor node within a pipeline and acoustic backscatter communication to transmit sensor reading out of the pipe.

Acoustically powered and connected sensors have a huge potential for structural and industrial monitoring. However, they must be cheap and easy to deploy to be economically attractive. Acoustic communication through metals is an emerging area, and many exciting research paths are yet unexplored. Current solutions only touch upon the surface of this potential and are often very specific to a single application and tailored to the specific specimen tested in a study. We envision networks of passively communicating and wirelessly powered sensor nodes within all critical infrastructure. RFID has already fulfilled this vision in the realm of electromagnetic waves and delivers promising approaches that we can transfer to acoustic waves. However, the nature of

acoustics has a unique set of challenges and opportunities in store that differ significantly from the established [RFID](#) scenarios: While metal structures are excellent conductors of acoustic waves, they usually also exhibit extreme multipath-propagation with strong reverberation for several milliseconds, leading to frequency dependency. Furthermore, usable bandwidths are three to four orders of magnitude smaller than in comparable electromagnetic communication systems.

1.1 Contributions

This dissertation investigates communication and power transfer with [RFID](#)-like passive modulation using acoustic waves through structural elements. Because of the excellent propagation of acoustic waves in metals, the dissertation concentrates on metal structures. However, the explored techniques are also feasible in other solid media, such as carbon-fiber composites or reinforced concrete. In detail, we make the following contributions:

- A better understanding of the nature of metal channels. Therefore, we analyze and test our developments in several distinct physical channels on different specimens, while most prior investigations tailored their solutions to a single channel only, with no generality. Furthermore, we investigate the influence of transducers and the coupling between the electrical and mechanical behavior of the system, which significantly impacts system design.
- A method to achieve dynamic higher-order modulation with passive backscatter communication, using the scarce available bandwidth more efficiently. While straightforward higher-order modulation has already been proposed for [RFID](#), the coupling between structural and transducer impedance in through-metal communication requires custom hardware modification for every deployment. We develop automatic channel estimation procedures and hardware that can dynamically adapt carrier frequencies and symbol constellations to every new deployment based on a mathematical model of the through-metal channel in steady-state.
- A model of the nonlinear transient backscatter channel, on which we show that linear equalizers can be applied despite the nonlinearity of the backscatter channel. We compare two popular equalization techniques and demonstrate that decision feedback equalizers increase data rates the most.
- A wireless power transfer system that uses maximum power point tracking to dynamically adapt to the characteristic impedance of the channel to increase transfer efficiency. We also provide a detailed analysis of harvesting potential in through-metal channels.

- An investigation of wideband modulation for the communication towards the resource-constrained passive sensor. Despite the need to use simplistic demodulation techniques because of the constrained resources, modulation with wideband pulses achieves, on average, a several-fold increase in data rate and removes the need to select a suitable carrier frequency dynamically.

In this dissertation, we focus on practical solutions. In contrast to many prior studies, we develop real prototypes based on microcontrollers and custom circuit boards constructed from commercially available components. Resources on the passive nodes are strictly limited to achieve ultra-low-power consumption. Furthermore, as typical for wireless communication, oscillators at the sender and receiver node are never perfectly synchronized, and frequency offsets occur. We investigate how these offsets influence the performance of acoustic backscatter and evaluate possible remedies.

1.2 Organization

The remainder of this dissertation is organized as follows. Chapter 2 summarizes the fundamentals of wireless communication, the piezoacoustic backscatter phenomenon, wireless power transfer, and elastic waves in solids. Then, Chapter 3 reviews the state of the art of passive backscatter communication and power transfer in RFID and then focuses on through-metal communication schemes. Section 4.2 defines the goals of this dissertation and explores the properties of the transmission channels and transducers used for evaluation throughout this dissertation.

Next, Chapter 5 focuses on expanding the modulation order in backscatter communication from simple binary modulation to up to 16-QAM. Furthermore, we present a reader and tag prototype constructed from commercially available off-the-shelf components. The hardware design also serves as the base for prototypes in the later chapters, where some adjustments to the prototype follow.

Chapter 6 explains why conventional equalizers can not directly be used to mitigate Inter-Symbol Interference (ISI) in acoustic through-metal channels because of the channel nonlinearity. We will characterize the transient channel and present a technique to linearize it and apply equalization, which increases the achievable bitrates several times.

In Chapter 7, we combine wireless energy transfer with acoustic backscatter communication, showing that energy transfer is also practically feasible over longer distances. We also strive to increase the ranges over which such sensors can be powered and read out. We, therefore, introduce a method to cancel out the self-interference in the reader node by isolating the backscattered signal before sampling.

Chapter 8 focuses on the communication from the active to the passive device, where the main challenge is that the resource-constrained receiver can not apply complex filters or equalizers. We utilize wideband pulses to mitigate ISI without requiring complex receiver algorithms. Furthermore, we analyze the interdependence between the energy

harvester and receiver at the same device. Finally, Chapter 9 concludes the dissertation by summarizing the results and discussing future research paths.

Fundamentals

This chapter reviews the indispensable foundation of wireless communication and energy harvesting that this thesis builds upon. We start with a brief summary of wireless communication, then move on to piezoelectric transducers and the piezoelectric backscatter effect. We then discuss wireless power transfer and energy harvesting and conclude with a short wrap-up of elastic wave propagation in solids.

2.1 Notation

The remainder of this thesis uses the following mathematical notation:

x^*	Complex conjugate of x
$\text{Re}\{x\}, \text{Im}\{x\}$	Real and imaginary part of x
$\mathcal{F}\{x(t)\}$	Fourier transform of $x(t)$
$\mathcal{F}^{-1}\{X(f)\}$	Inverse Fourier transform of $X(f)$
$x(t) * y(t)$	Convolution of signal $x(t)$ with $y(t)$
$x[k]$	Time-discrete signal x evaluated at sample index k
$x(t)$	Time-continuous signal x evaluated at time t
$\angle(x)$	Phase of x
$\ x\ $	Euclidean norm of x
\mathbf{x}, x_k	Vector \mathbf{x} , and the k -th element of the vector
$\delta(t)$	Dirac delta function

2.2 Wireless Communication

Since the first transmission of a wireless message in the late 19th century, the field has made tremendous progress—to the point of ubiquitous connectivity that we enjoy

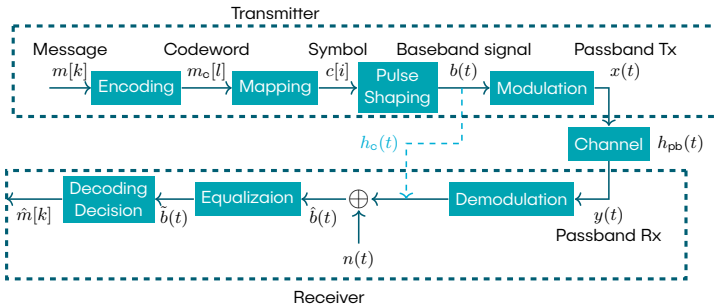


Figure 2.1: Block diagram showing the stages of a communication system.

today. Most of this progress concentrated on the transmission of information via electromagnetic waves. However, the established theoretical background is not limited to electromagnetism but also applies to other forms of waves, e.g., acoustic waves. This section discusses the fundamental steps in a wireless communication system required to follow the work presented in this dissertation, as depicted in Fig. 2.1. As a detailed review of all aspects of wireless communication is beyond the scope of this dissertation, we cover only the parts most relevant to this dissertation. If no explicit sources are provided, we refer the reader to [Gol05].

The input to the transmitter is a digital message, i.e., a bitstream $m[k]$. In the encoding step, the transmitter applies a channel code to the original bitstream, which adds redundancy to the message, aiding the receiver in detecting and correcting transmission errors. The encoded bitstream $m_c[l]$ is then mapped to symbols, where each symbol $c[i]$ encodes one or more bits. After mapping, pulse shaping converts the symbols into time-continuous analog signals, yielding the *baseband signal* $b(t)$. The baseband signal is then modulated, i.e., up-converted to the *passband*—a higher frequency band. The passband signal $x(t)$ is emitted into the wireless channel, where it is distorted with the channel's impulse response $h_{pb}(t)$. The receiver senses the distorted signal $y(t)$ from the channel and performs the above steps in reverse. That is, it first demodulates it, down-converting the signal back to the baseband, yielding $\hat{b}(t)$. An equalizer then compensates for the distortion introduced by the channel, yielding $\hat{b}(t)$. The equalized signal is then decoded, i.e., the symbols are mapped back to bits, yielding $\hat{m}[k]$. Within this process, noise $n(t)$ occurs in various stages of the system. These steps are explained in more detail in the following sections.

2.2.1 Encoding and Decoding

Encoding adds redundancy to the message to enable the receiver to detect and correct transmission errors. An important property of a code is the code rate R , which is the

ratio of the number of information bits to the total number of bits in a codeword.

One common type of channel code are Low-Density Parity-Check (LDPC) codes, which are used in many modern communication systems. LDPC codes are linear block codes, which means that the encoding process can be represented as a matrix multiplication of the original message with a generator matrix G , i.e.,

$$\mathbf{m}_c = \mathbf{m} \cdot \mathbf{G}. \quad (2.1)$$

While encoding a message is straightforward, decoding is more computationally demanding. Directly deriving the codeword with maximum probability from a received message is infeasible on realistic hardware. Instead, a feasible decoding algorithm is the Belief Propagation (BP) algorithm, which approximates the optimal solution iteratively. A detailed description of the decoder and derivation of the underlying equations is out of the scope of this dissertation, but we refer the reader to [Joh09].

2.2.2 Mapping and Constellations

Mapping translates a sequence of bits to a sequence of symbols in the baseband. The mapping is usually defined by a *constellation*, a set of valid baseband signals. We represent the constellation as a set of complex numbers, where each complex number is associated with a symbol. The number of symbols in the constellation S is usually a power of two, i.e., $S = 2^k$. Then, one symbol can represent k bits. The choice of the constellation has a significant impact on the performance of the communication system. The most common constellations are Phase Shift Keying (PSK), Amplitude Shift Keying (ASK), and Quadrature Amplitude Modulation (QAM). Often, the number of constellation points is prefixed, e.g., 4-PSK, 4-ASK, and 4-QAM all consist of four distinct constellation points.

In PSK, all constellation points have equal amplitude but differ in the phase angle of the complex number, while in ASK, all constellation points have the same phase angle but differ in amplitude. QAM combines both PSK and ASK, i.e., the constellation points differ in both amplitude and phase angle. Figure 2.2 shows examples for these constellations. In any case, it is desirable to maximize the distance between adjacent constellation points, as this increases robustness against noise. The smallest distance between two constellation points is called the *minimum distance* of the constellation. In some cases, the same constellation is known under different terms, e.g., a 2-ASK constellation where one constellation point is zero is also known as On-/Off-Keying (OOK). Similarly, 4-QAM can yield the same constellation as 4-PSK.

Mapping results in a time-discrete sequence of constellation points $c[i]$. To transform this into a time-continuous baseband signal, the transmitter uses a *pulse shaping filter*, which convolves the sequence of constellation points with a time-continuous pulse. The

2 Fundamentals




Figure 2.2: Typical forms of constellations are ASK, PSK and QAM. The graphs also show the mapping from bits to symbols.




Figure 2.3: Convolution with the pulse shape transforms the time-discrete baseband signal into a time-continuous signal with a given symbol duration. This figure shows the results for two different pulse shapes: Rectangular and raised cosine.

pulse shaping filter is also responsible for the *symbol interval* T_m of the communication system. With pulse shape $p(t)$, the baseband signal is described by

$$b(t) = \left(\sum_i c[i] \cdot \delta(t - iT_m) \right) * p(t), \quad (2.2)$$

where $\delta(t)$ is the dirac delta function, i.e., the unit impulse. Figure 2.3 shows the convolution of an OOK signal with two common pulse shapes, a rectangular pulse and a raised cosine pulse. The rectangular pulse can be implemented purely with digital logic, but a drawback is the infinite bandwidth because of the sharp edges. The raised cosine pulse, in contrast, has a limited bandwidth but is usually implemented with a digital filter with many taps and requires a Digital to Analog Converter (DAC) with a high update rate to generate the corresponding analog signal.

Many communication technologies apply pulse shapes with a narrow bandwidth, such as the raised cosine pulse. These *narrowband* pulses typically occupy bandwidths inversely proportional to the symbol rate, i.e., $B \approx 1/T_m$. In contrast, other schemes deliberately use *wideband* pulses, e.g., chirps or spreading codes that have advantageous auto-correlation properties, to increase the resilience against noise or jamming

in specific frequency bands. Furthermore, wideband pulses are typically more resilient against multipath fading.

Modulation

Modulation mixes the baseband signal with a carrier, converting it into the passband—a favorable frequency range in the transmission medium. The carrier is a complex sinusoid with a frequency f_c , i.e.,

$$x(t) = \text{Re}\{b(t) \cdot e^{j2\pi f_c t}\}, \quad (2.3)$$

which is equivalent to

$$x(t) = \text{Re}\{b(t)\} \cos(2\pi f_c t) - \text{Im}\{b(t)\} \sin(2\pi f_c t). \quad (2.4)$$

Hence, the carrier's amplitude is proportional to the magnitude of the baseband signal, and the carrier's phase offset is proportional to the phase of the baseband signal. Modulation is often distinguished into Amplitude Modulation (AM) and Phase Modulation (PM). That directly relates to the constellation in the digital communication system. If ASK is employed, the baseband signal is real-valued, leading to pure AM, while PSK is pure PM, and QAM combines AM and PM.

We want to note that the discussion of modulation schemes is not exhaustive. Modern communication schemes employ numerous other modulations, such as Frequency Modulation (FM), where the symbols in baseband are sinusoids with different frequencies, or Orthogonal Frequency Division Multiplexing (OFDM), where the baseband signal is split into multiple subcarriers, each transmitting different message sections. However, these modulation schemes are not relevant to this dissertation.

The wireless channel

Once the passband signal has been generated, the sender emits it into the wireless medium using an antenna or other transducer. While traveling through the wireless medium to the receiver over a single path, the physical waves experience distortions. First, attenuation decreases the wave's amplitude. Second, the wave travels with a finite speed, requiring a delay proportional to the distance between the transmitter and receiver. In the received signal, the attenuation is represented by a scaling. In general, attenuation and time shift are also frequency-dependent.

Additionally, the signal may travel to the receiver over multiple paths when reflected at obstacles in the environment, leading to interference at the receiver. As each path i covers a different distance and delay, it is scaled with an attenuation factor a_i ($|a_i| \leq 1$) and experiences a time delay $\tau_i > 0$. We model the channel with its channel impulse

2 Fundamentals

response $h_{\text{pb}}(t)$, which is the sum of the attenuated and time-shifted components of each path, i.e.,

$$h_{\text{pb}}(t) = \sum_i a_i \delta(t - \tau_i), \quad (2.5)$$

and the received signal is then represented as the convolution of the transmitted signal with the channel impulse response, i.e.,

$$y(t) = x(t) * h_{\text{pb}}(t). \quad (2.6)$$

The multipath propagation leads to distortions of the received signal, e.g., Inter-Symbol Interference (ISI). Due to the time delay of the paths, a late path of a previous symbol can interfere with an early path of a later symbol, which may lead to decoding errors at the receiver. The extent of ISI can be quantified with the channel delay spread—the time difference between the first and last significant multipath component. A more clearly defined version of the channel delay spread is the Root-Mean-Square (RMS) delay spread, which we use throughout this dissertation. It weights the delay of propagation paths by their power, given by

$$\tau_{\text{rms}} = \sqrt{\frac{\int_0^\infty (\tau - \bar{\tau})^2 \|h_{\text{pb}}(\tau)\| d\tau}{\int_0^\infty \|h_{\text{pb}}(\tau)\| d\tau}}, \quad (2.7)$$

where $\bar{\tau}$ is the mean delay of the channel given as

$$\bar{\tau} = \frac{\int_0^\infty \tau \|h_{\text{pb}}(\tau)\| d\tau}{\int_0^\infty \|h_{\text{pb}}(\tau)\| d\tau}. \quad (2.8)$$

A related property is the *coherence bandwidth* of a channel, which is defined as a bandwidth around the center frequency in which the channel impulse response is roughly flat. The coherence bandwidth is approximately the inverse of the delay spread.

Another important property is the signal-to-noise ratio (SNR), which is the ratio of the signal power to the noise power. The SNR is usually given in decibels (dB) and is defined as

$$\text{SNR} = 10 \log_{10} \left(\frac{\int_0^T \|\hat{b}(t)\|^2 dt}{\int_0^T \|n(t)\|^2 dt} \right) \quad (2.9)$$

over a period T , with $\hat{b}(t)$ as received baseband signal and $n(t)$ as noise signal.

Demodulation

After the signal propagates through the wireless channel, the receiver demodulates it, i.e., inverts the modulation to recover the baseband signal. In pure AM, the envelope,

the signal's instantaneous amplitude, can be recovered with rectification and low-pass filtering. However, a more complex approach is required when the carrier's phase is required, too. Processing the signal directly in the passband requires high sampling rates. Therefore, the receiver first down-converts the signal to a lower frequency band. The most common approach to achieve this is to mix the received signal with an unmodulated carrier with the same frequency f_c .

Using the fundamental trigonometric identities

$$\cos(\phi_0) \cdot \cos(\phi_1) = \frac{1}{2} [\cos(\phi_0 + \phi_1) + \cos(\phi_0 - \phi_1)] , \text{ and} \quad (2.10)$$

$$\sin(\phi_0) \cdot \cos(\phi_1) = \frac{1}{2} [\sin(\phi_0 + \phi_1) + \sin(\phi_0 - \phi_1)] , \quad (2.11)$$

we can rewrite the transmitted signal from Eq. (2.4), mixed with a cosine of the same frequency as

$$\begin{aligned} \hat{b}_I(t) &= \underbrace{[\text{Re}\{b(t)\} \cos(2\pi f_c t) - \text{Im}\{b(t)\} \sin(2\pi f_c t)]}_{\text{Transmitted}} \cdot \cos(2\pi f_c t + \phi_o) \\ &= \frac{\text{Re}\{b(t)\}}{2} [\cos(4\pi f_c t + \phi_o) + \cos(-\phi_o)] \\ &\quad - \frac{\text{Im}\{b(t)\}}{2} [\sin(4\pi f_c t + \phi_o) + \sin(-\phi_o)] . \end{aligned} \quad (2.12)$$

The mixed signal contains a baseband component and a component at twice the carrier frequency. The high-frequency components can be removed with a lowpass filter, with only the baseband component remaining, i.e.,

$$\hat{b}_I(t) \doteq \frac{1}{2} [\text{Re}\{b(t)\} \cos(-\phi_o) - \text{Im}\{b(t)\} \sin(-\phi_o)] , \quad (2.13)$$

where ϕ_o is a phase offset introduced because the transmitter and receiver are usually not synchronized, and the receiver starts sampling at an arbitrary point in time. Similarly, mixing the received signal with a sine at the same frequency yields


$$\hat{b}_Q(t) \doteq \frac{1}{2} [-\text{Re}\{b(t)\} \sin(-\phi_o) - \text{Im}\{b(t)\} \cos(-\phi_o)] . \quad (2.14)$$

Interpreting \hat{b}_I and \hat{b}_Q as a complex number yields the demodulated signal

$$\hat{b}(t) = \hat{b}_I(t) - j \cdot \hat{b}_Q(t) = \frac{1}{2} e^{-j\phi_o} \cdot b(t) . \quad (2.15)$$

Hence, the demodulated signal is a rotated and scaled version of the transmitted baseband signal. This derivation did not consider the distortion in the passband due to the channel. Instead of describing this with the passband channel impulse response, the channel is often characterized by the baseband channel impulse response $h_c(t)$,

2 Fundamentals



■ **Figure 2.4:** Illustration of (left) the demodulator at the receiver and (right) the frequency domain representations of the signals throughout the stages. After demodulation, the distorted baseband signal is recovered. The higher-frequency components are removed with a lowpass filter.

which describes the transformation between the baseband signal at the transmitter and the demodulated signal at the receiver. The baseband impulse response is essentially the downconverted and filtered passband impulse response. Figure 2.4 illustrates the demodulator and shows the frequency domain of the signal during the transmission stages.

2.2.3 Equalization

We discussed that the multipath propagation distorts the received signal significantly, and especially ISI may lead to decoding errors. Equalizers are designed to mitigate the distortion by learning and reverting the effect of multipath propagation. The *linear equalizer* is a filter with impulse response $h_{eq}(t)$. Without noise, the ideal filter response is such that convolution with the channel impulse response yields a dirac delta impulse, which may incorporate a time delay t_0 , i.e.,

$$h_{eq}(t) * h_c(t) = \delta(t - t_0). \quad (2.16)$$

Consequently,

$$\begin{aligned} \tilde{b}(t) &= \hat{b}(t) * h_{eq}(t) \\ &= (b(t) * h_c(t) + n(t)) * h_{eq}(t) \\ &= b(t - t_0) + (n(t) * h_{eq}(t)). \end{aligned} \quad (2.17)$$

Initially, the receiver can not know the channel impulse response and must estimate it. Therefore, the transmitter sends a *training sequence*, a predefined sequence of symbols known to the receiver. The receiver can compare the known transmitted and received signals and estimate the optimal filter coefficients that transform one into another.

Unfortunately, since noise compromises the received training sequence, the equalizer can only approximate the channel impulse response. Furthermore, during the reception

of the main message, noise will also distort the equalizer's output. In channels with a low Signal to Noise ratio (SNR), or highly frequency-selective channels, the equalizer can amplify the noise in the received signal significantly, leading to substantial noise in the equalizer's output. In such channels, choosing the equalizer impulse response according to Eq. (2.16) leads to bad performance. Instead, a better choice is to minimize the mean square error between the expected training sequence and the equalizer output, known as the Minimum Mean Square Error (MMSE) criterion [Gol05, p.357f].

Different types of equalizers have been developed to reduce the influence of noise on the equalizer output. This dissertation employs the MMSE equalizer and the Decision Feedback (DF) equalizers. The MMSE equalizer is linear and optimizes its weights to minimize the mean square error between the expected training sequence and the equalizer output, balancing precise channel inversion and noise enhancement.

The DF equalizer is a nonlinear equalizer, which further reduces the noise amplification compared to linear equalizers. It consists of a short feedforward filter and a feedback filter. After every equalized symbol, it makes a symbol decision, i. e., extracts the constellation point closest to its output. This constellation point feeds into the feedback filter. Hence, the feedback filter contains noise-free information about the previously transmitted symbols. However, a severe disadvantage of the DF equalizer is that once it makes a wrong decision, this decision error propagates through the feedback filter to the following symbols, which can lead to burst errors.

We discuss the detailed implementation and equations for equalizer training used within this dissertation in the context of the designed backscatter communication system in Section 6.3.3.

2.2.4 Synchronization

In wireless communication, the transmitter and receiver are not synchronized, which has several implications. First, the receiver must detect the start of a message, known as *frame synchronization*. Second, the oscillators used in the devices are not ideal. For modulation, the transmitter needs an oscillator to generate the carrier frequency, and the receiver needs an oscillator at *exactly* the same frequency for ideal demodulation. In reality, oscillators are inaccurate, and their frequencies deviate between devices. The task of *frequency synchronization* is to detect and compensate for the Carrier Frequency Offset (CFO). Lastly, mismatching oscillators also cause a Sample Frequency Offset (SFO), i. e., the sampling rate at the receiver does not perfectly match the symbol rate generated at the transmitter. This section discusses techniques to detect and compensate for these synchronization errors.

Phase and Frequency Synchronization

With realistic oscillators, the local oscillator at the receiver has a frequency offset f_o , such that the demodulator mixes it with a frequency of $f_c + f_o$. This results in a periodic

2 Fundamentals

rotation of the received baseband signal after demodulation, i.e.,

$$\hat{b}(t) = \underbrace{e^{j2\pi f_o t + \phi_o}}_{\text{Rotation}} \cdot (b(t) * h_c(t)) + n(t). \quad (2.18)$$

While many approaches exist to detect and mitigate the CFO [LP17], we focus on two methods in this dissertation. We distinguish between static detection at the beginning of a message—which is sufficient when the CFO is approximately constant—and a dynamic correction continuously updating its CFO estimate throughout the message.

Static detection transmits a CFO synchronization sequence at the beginning of the message. When transmitting an unmodulated carrier, $b(t)$ in Eq. (2.18) is constant, and the receiver detects only the rotation caused by the CFO (and noise). To determine the frequency offset, it calculates the phase difference between two or more samples. The more samples are considered in the estimation, and the longer the estimation sequence is, the more accurate the estimate is. More formally, the receiver estimates the CFO from the demodulated signal using

$$f_o = \frac{f_s}{2\pi} \angle \left(\frac{1}{N} \sum_{i=0}^{N-1} \hat{b} \left(\frac{i}{f_s} \right) \cdot \hat{b}^* \left(\frac{i+1}{f_s} \right) \right), \quad (2.19)$$

with a sampling rate f_s . Once f_o is known, the receiver can compensate it by multiplying the received signal with a complex sinusoid $\exp(-j2\pi f_o t)$ in the inverse direction, recovering the frequency-synchronized signal.

A Digital Phase Locked Loop (DPLL) can be used for dynamic CFO detection and mitigation. The core idea of the DPLL is that after making a symbol decision for each symbol, the phase offset between the detected sample and the ideal constellation point serves as an error signal. According to the error signal, the DPLL adapts its current phase offset estimate as shown in Fig. 2.5. The error signal is additionally passed through a loop filter with gain k . The loop filter can take various forms depending on the characteristics of the communication system. The filter gain must be adapted to the SNR to find a balance between quick response and stability of the feedback loop.

Frame Synchronization

A frame is a packet of information, i.e., a sequence of symbols that follow a protocol known to the transmitter and receiver. To correctly decode such a packet, the receiver must detect the start of a frame. Frame synchronization is often realized by prepending the message with a preamble $b_p(t)$, a predefined sequence of symbols known to the receiver. The receiver can then perform a cross-correlation between the received signal and the expected preamble to determine the part of the received signal that matches

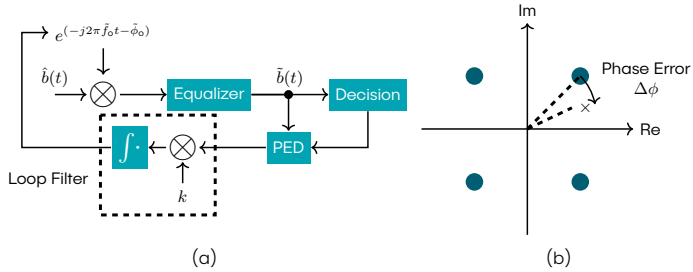


Figure 2.5: Principle sketch of (a) a DPLL with proportional loop filter to track and compensate phase errors, and (b) a possible Phase Error Detector (PED) implementation working on a sample and a decision.

the preamble the closest. The correlation between the received baseband signal and the preamble is given by

$$\text{corr}_{\hat{b}, b_P}[i] = \sum_{k=0}^{\infty} \hat{b}^*[k] \cdot b_P[k - i]. \quad (2.20)$$


The most probable start index i_0 of the correlation sequence in the received signal is then

$$i_0 = \arg \max_i \|\text{corr}_{\hat{b}, b_P}[i]\|. \quad (2.21)$$

As an additional benefit, the cross-correlation also yields the phase shift between the received signal and the preamble, i.e., the phase shift ϕ_o introduced due to the channel delay

$$\phi_o \approx \angle (\text{corr}_{\hat{b}, b_P}[i_0]). \quad (2.22)$$

Care must be taken to choose a suitable preamble sequence, i.e., one that produces a high and uniquely identifiable peak when correlated with itself. A common choice is a Barker-Sequence, where every correlation with a version of itself shifted by one or more symbols is zero, and only the exact overlap produces a high correlation peak. In contrast, randomly chosen symbol sequences often produce multiple correlation peaks. For example, Fig. 2.6 shows the autocorrelation of a 13-symbol Barker-Sequence compared to a repetitive clocklike signal with many peaks. In a noisy channel, the peaks of the clocklike signal may be ambiguous, while the peak of the Barker-Sequence is more pronounced.



■ **Figure 2.6:** Illustration of frame synchronization. The receiver correlates the received signal with the preamble to determine the start of the message. The figure shows two exemplary preamble sequences, where the Barker-Sequence is much better suited for correlation.

Symbol Timing

As the receiver starts sampling at an arbitrary point in time, the sampling instants are usually not ideal, i.e., a sample is not taken at the center of a symbol, which we call *Timing Error*. Furthermore, the sampling clock at the receiver usually deviates from the clock at the transmitter, causing a drift of the sampling instants over time, which we refer to as *Sampling Frequency Offset* (SFO). Throughout this dissertation, we use a static approach to detect the timing error once at the beginning of a message, and a dynamic approach to track and correct the time-variant SFO throughout the message.

The static approach uses a part of the signal to detect the symbol timing using a square timing error detector first described in [OM88]. This detector is based on the observation that a signal of alternating ones and zeros in the baseband is similar to a sine with a frequency of half the symbol rate. Then, the square of this signal has a period equal to T_m (see Fig. 2.7). With such a signal, the digital Fourier transform evaluated for the frequency bin $1/T_m$ yields the phase offset $\Delta\phi_m$, which can be used to correct the timing error, e.g., using interpolation. To resolve the frequency component in the digital Fourier transform, the sampling rate of the baseband signal must be more than twice the symbol rate. The square timing error detector also works on signals that are not purely alternating ones and zeros, as long as the signal contains sufficient zero-crossings, usually ensured by encoding with a channel or source code. However, the square timing error detector may fail when the signal contains long streaks of ones or zeros.

Like the frequency offset, diverging clock frequencies at the transmitter and receiver also cause the sample timing to deviate throughout the message. The dynamic approach tracks the timing error over time using a *timing-locked loop*. The differences to the previously discussed DPLL are

1. the error detector, which needs to detect if the sampling instants are too early or too late, and
2. the compensation mechanism, which interpolates the ideal sample between the original samples $\hat{b}(l/f_s)$, as displayed in Fig. 2.8 (left).




Figure 2.7: Illustration of the square timing error detector with an alternating signal (blue) of ones and zeros and the squared signal (orange), having a period of T_m .




Figure 2.8: A timing-locked loop (left) takes feedback from a timing error detector (TED) and interpolates the samples in the baseband signal. The Derivative timing error detector (right) uses the derivative and sample value to determine if the sample has been taken too early or too late.

For 1), a *derivative timing error detector* is employed. This timing error detector assumes that the derivative of the received signal is zero at the center of a transmitted pulse. If the pulse is positive, a positive derivative indicates that the sampling instant is too early, while a negative derivative indicates it is too late. In the case of a negative pulse, the opposite is the case (see Fig. 2.8 right). Therefore, the timing error signal is proportional to

$$e_T(t) \propto \hat{b}(t) \cdot \frac{d}{dt} \hat{b}(t). \tag{2.23}$$

Similarly to the frequency-locked loop, the error signal is multiplied with a gain and integrated, which reduces noise and stabilizes the feedback loop. The accumulated timing error then controls the interpolator, which will correct subsequent samples accordingly, synthetically recreating the samples at the center of a symbol. However, in a software-defined receiver, the derivative must be estimated, usually by taking the central difference of the digital samples, i. e., with sampling interval T_s , the derivative is approximated as

$$\frac{d}{dt} \hat{b}(l/f_s) \approx \frac{1}{2T_s} \cdot (\hat{b}([l+1]/f_s) - \hat{b}([l-1]/f_s)). \tag{2.24}$$

This way of approximating the derivative requires a sampling rate several times the symbol rate, e.g., 4 or 8 times as fast. Like the static method, the timing error can only be detected when the symbol value changes. Long streaks of ones or zeros can lead to a loss of synchronization.

2.3 Piezoacoustic Backscatter

This section reviews the piezoelectric effect, the transducers fabricated from piezoelectric materials, and the acoustic reflection coefficient, which are critical to understanding the physical foundation of backscatter communication.

2.3.1 Piezoelectric Transducers

The *direct piezoelectric effect* describes a property of certain materials in which a mechanical strain forced upon the material induces a change of an electric charge on the material surfaces. Figure 2.9 (a) shows a simplified model of a quartz crystal—which has a regular inner structure of positively and negatively charged particles—with electrodes applied to both sides. In equilibrium, the center of the positive and negative charges coincide, and the crystal is electrically neutral. If a mechanical force F_m deforms the crystal, as shown in Fig. 2.9 (b), the charge centers shift, inducing an electric field between the electrodes.

Similarly, an externally generated voltage V_{e1} between the electrodes generates an electric field, which acts on the charged particles with a force F_{e1} . Upon the force, the crystal is deformed. This is called the *inverse piezoelectric effect*. A detailed description of the fundamental physics of piezoelectricity is out of scope for this thesis, but a comprehensive summary is given in [Rup19].

For our purposes, it suffices to know that the coupling between mechanical strain and charge on the transducer is linear, i.e., the induced charge ΔQ is linearly coupled to the particle displacement Δx in the crystal, and the electrically generated force F_{e1} is linearly coupled with the voltage drop over the crystal. Hence,

$$\Delta Q \propto \Delta x \quad \text{and} \quad F_{e1} \propto -V_{e1} . \tag{2.25}$$

A piezoelectric transducer consists of a piezoelectric material coated with two electrodes. As piezoelectric materials are always dielectrics, the transducer resembles a capacitor, but energy can be stored in the electric and mechanical fields.

The energy exchange between the electric and mechanic fields makes any transducer a resonator. The exact resonance frequency of a transducer depends on its material parameters and geometry. In this dissertation, we used only disk-shaped transducers, which have two fundamental resonance modes: a longitudinal mode, where the mechanic oscillation occurs in the thickness direction of the transducer, and a transverse




Figure 2.9: The direct and inverse piezoelectric effect shown for a quartz crystal. In equilibrium (a), the center of positive and negative charges (C_+ , C_-) overlap. On deformation (b), the charge centers move, creating an electric field. Applying an external electric field (c) exerts a mechanical force on the charges, causing deformation. The illustration is adapted from [Rup19].





Figure 2.10: A transducer can resonate in transverse (radial) mode or in longitudinal (thickness) mode, yielding two distinct resonant frequencies.

mode, where the oscillation occurs in the radial direction (see Fig. 2.10). Driving the transducer close to its resonance is essential for high-efficiency conversion of mechanical energy into electrical energy and vice versa. This is especially important for wireless power transfer systems, where transfer efficiency is critical. A detailed inspection of the resonances and electrical impedance of the transducers used in this dissertation follows in Section 4.2.1.

2.3.2 Reflection Coefficient

When a wave travels through a medium and encounters an impedance discontinuity, e.g., a junction of two different materials, part of the wave is reflected. Analogously to electromagnetic waves, this is also true for acoustic waves, where the *acoustic impedance* measures a system's opposition to particle movement when excited with



■ **Figure 2.11:** The transducer (left) can change the acoustic impedance of the cross-section to which it is attached. The load impedance (right) influences the relation between the forces from the direct and inverse piezoelectric effect.

a periodic force. The acoustic impedance Z_a is a complex and frequency-dependent quantity, which is defined as

$$Z_a(f) = \frac{\text{Pressure}}{\text{Velocity}} = \frac{1}{A} \frac{\mathcal{F}\{F_m(t)\}}{\mathcal{F}\{dx(t)/dt\}}, \quad (2.26)$$


with mechanical force F_m acting on area A , and particle velocity dx/dt [Bru06].

When excited with a sinusoidal force, the particle velocity will also be sinusoidal. However, the relation between magnitudes and phase differences of force and velocity is influenced by the material properties. When assuming a thin sheet from a homogeneous material as a waveguide (see Fig. 2.11, left), the acoustic impedance is the same everywhere, and a wave can travel through the material without reflection. When a transducer is attached to the metal surface, it changes the stiffness and mass of the cross-section to which it is attached. Furthermore, due to the active piezoelectric material, the transducer can generate a force on the surface.

Due to the coupling of mechanical and electric fields in a piezoelectric transducer, the load impedance—the electric impedance connected between a transducer’s terminals—influences the acoustic impedance of the transducer. Figure 2.11 (right) is a simplified model of how the forces from direct and inverse piezoelectric effects are related. In this model, a charge source represents the induced charge proportional to the external mechanical force, and the voltage over the load impedance is proportional to the force induced from the electric field. Hence, short-circuiting the load results in no electric field ($V_l = 0 \text{ V}$) and no force F_{el} from the inverse piezoelectric effect is acting on the transducer. On the other hand, an infinite load impedance results in the maximum V_{el} and, hence, a force F_{el} acting on the transducer in the opposite direction of the driving mechanical force. Similarly, capacitive or inductive loads result in a phase shift between the mechanical and electric forces up to $\pm 90^\circ$.

Wave theory provides a simple characterization of this effect, the complex reflection coefficient Γ . At a transmission from one medium with characteristic acoustic impedance Z_a to another medium with acoustic impedance Z'_a the reflection coefficient is

$$\Gamma_a = \frac{Z'_a - Z_a}{Z'_a + Z_a}, \quad (2.27)$$



■ **Figure 2.12:** Wireless power transfer system overview.

where Γ_a is a complex number whose magnitude describes the fraction of the wave's amplitude that is reflected, and its phase describes the phase shift of the reflected wave relative to the phase of the incident wave [Ros14].

In the example from Fig. 2.11, we simplified the model by assuming that the transducer is perfectly coupled to the metal beam. However, the binding layer, which attaches the transducer to the surface, is never perfectly rigid, i. e., it will never fully transmit the mechanic wave in the underlying structure to the transducer. Furthermore, some of the wave energy dissipates during conversion between the mechanic and electric domain due to internal resistances in the transducer. Hence, the acoustic reflection coefficient can never fully achieve a magnitude of one. Detailed analysis and modeling of the piezoelectric reflection coefficient is beyond the scope of this thesis but can be found in [All21]. We experimentally analyze the characteristic impedances of the transducers and channels used throughout this thesis in Section 4.2.

2.4 Wireless Power Transfer

A Wireless Power Transfer (WPT) system consists of a power transmitter and one or more receivers. Different underlying mechanisms can be leveraged to transfer power wirelessly, e. g., electromagnetic waves, inductive coupling, capacitive coupling, and elastic waves. In this thesis, the discussion is limited to ultrasonic wireless power transfer. However, the building blocks depicted in Fig. 2.12 are similar, independent of the transmission medium and type of wave.

The *transmitter* transforms electric energy into a mechanical wave. Therefore, it generates an electric AC waveform and applies it to a transducer. The transducer excites a mechanic waveform in the coupled medium, which propagates on one or more paths to the receiver. A second transducer converts the mechanic waveform back to an electric waveform at the receiver. At this point, the transducer already provides an electric AC signal at the excited frequency. However, electronic devices such as sensors and Microcontroller Units (MCUs) require a DC voltage within a suitable range. High voltages must be avoided to prevent device damage, while too low voltages can not power the digital circuitry. Therefore, a harvesting circuit converts a wide range of

incoming AC waveforms into a regulated DC voltage that fulfills the requirements of digital electronics.

At every stage in this system, part of the energy is dissipated into the environment and is lost for the receiver. The challenge in a WPT system is to optimize the parts so that the largest possible fraction of the energy spent at the transmitter is available as regulated DC power at the receiver. However, certain system components are not under the designer's control, e.g., the wireless channel, where typically most energy is lost. Hence, the rest of the system must be tailored to best suit the channel's characteristics.

The targeted wave mode and frequency are the most influential design decisions in ultrasonic channels in solid structures [Tse+20]. Once these are known, the transducer must be chosen that is best suited to excite this kind of waveform. We discuss the options in Section 4.2, where we investigate the channel in more detail. The rest of this section discusses the electric building blocks of the WPT system and especially the interfaces between these building blocks. We first review the maximum power transfer theorem to understand the role of interfaces between circuits.

2.4.1 Maximum Efficiency and The Maximum Power Transfer Theorem

While circuits in WPT systems often comprise many elements, the Thevenin equivalent states that any system that contains only voltage sources, current sources, and passive elements can be replaced by a single voltage source in series with a single output impedance [Joh03]. This law can also be extended to networks containing active elements, such as transistors and diodes, as long as the circuit behaves linearly around the observed operation point.

Using the Thevenin equivalent, we can describe the transmitter as an ideal voltage source generating $V_{s,tx}$ with output impedance $Z_{s,tx}$, as shown in Fig. 2.13 (left). At the same time, the transducer driven by the transmitter is modeled as a single load impedance $Z_{t,tx}$. The goal is to match the source's parameters to the transducer to maximize the power transfer efficiency η_{tx} , i.e.,


$$\eta_{tx} = \frac{\operatorname{Re}\{P_t\}}{\operatorname{Re}\{P_s + P_t\}} = \frac{\operatorname{Re}\{I_{t,tx}^2 Z_{t,tx}\}}{\operatorname{Re}\{I_{t,tx}^2 (Z_{t,tx} + Z_{s,tx})\}}. \quad (2.28)$$

Efficiency in such a circuit is maximized by first matching the transducer's reactance by choosing a $Z'_{s,tx}$ with opposite reactance, i.e.,

$$\operatorname{Im}\{Z'_{s,tx}\} = -\operatorname{Im}\{Z_{t,tx}\}, \quad (2.29)$$

in that case, the efficiency calculation is simplified to

$$\eta_{tx} = \frac{\operatorname{Re}\{Z_{t,tx}\}}{Z_{t,tx} + Z'_{s,tx}}. \quad (2.30)$$



■ **Figure 2.13:** Impedance mismatch between a source and a load leads to a reduced power uptake of the load.

This equation shows that efficiency can finally approach unity by minimizing the source's output impedance as much as possible. Since the generator typically uses an amplifier with a small output impedance, e. g., a few Ohms to tens of Ohms, and piezoelectric transducers typically have much larger input impedances, high efficiencies of more than 90% can be achieved in the transmitter.

At the receiver depicted in Fig. 2.13 (right), a similar matching problem arises. However, in this case, the source represents a transducer excited by a mechanical wave, and $Z_{s,rx}$ is the transducer's output impedance. The voltage $V_{s,rx}$ is not under the designer's control but depends heavily on the transmission channel. The transducer's series resistance $Z_{s,rx}$ is usually assumed to be a fixed property of the transducer. At the receiver, the goal is slightly different: Instead of maximizing transfer efficiency, we want to adapt the harvester's input impedance to extract the largest possible amount of power from the source.

The maximum power transfer theorem [ASS16] states that this is achieved by choosing

$$Z_m + R_h = Z_{s,rx}^* \quad (2.31)$$

Therefore, the harvester typically uses a matching network with equivalent impedance Z_m to match the source's reactance. Furthermore, the harvester must control the current to mimic a resistive load with resistance R_h , hence


$$Z_m = -j \cdot \text{Im}\{Z_{s,rx}\} \quad \text{and} \quad R_h = \text{Re}\{Z_{s,rx}\} \quad (2.32)$$

Optimal matching results in a maximum harvestable power

$$P_h^{\max} = \frac{V_h^2}{R_h} \quad \text{with} \quad V_h = \bar{V}_{s,rx} \cdot \frac{R_h}{R_h + Z_m + Z_{s,rx}} \quad (2.33)$$

$$P_h^{\max} = \frac{\bar{V}_{s,rx}^2}{4 \text{Re}\{Z_{s,rx}\}},$$

where $\bar{V}_{s,rx}$ is the RMS voltage of the input waveform. Hence, even in a perfect match,



■ **Figure 2.14:** Popular rectifier topologies.

only a fraction of the available source energy can be harvested, and the absolute value depends on the open circuit voltage of the transducer and its series resistance.

In some harvesting applications, the source's series resistance or open circuit voltage change dynamically, e.g., in solar harvesting, series resistance changes with solar radiation intensity. This is relevant for transfer through metal channels, as we discuss in Section 4.2. In such cases, Maximum Power Point Tracking (MPPT) is a useful technique where the harvester dynamically adapts its resistance to the time-variant source.

2.4.2 Rectification and Voltage Multiplication

The harvester must rectify the AC signal and convert it to a suitable voltage level. Most commonly, diode-based rectifiers are used for this purpose, which can be arranged in several configurations: Most reported are the half-wave rectifier, full-wave-rectifier, full-bridge rectifier, and voltage quadrupler [TCPI7] shown in Fig. 2.14. The half-wave rectifier is the simplest, requiring only a single diode and a capacitor, but can only rectify power during the positive half of an input wave, not fully leveraging the available power of the waveform. The full-wave rectifier requires two diodes, leverages the power of both positive and negative half-wave, and can create a DC output voltage with twice the amplitude of the input wave. The full-bridge rectifier requires four diodes and leverages the whole waveform. However, it does not double the input voltage. Lastly, a voltage quadrupler is a rectifier constructed from four diodes, that can achieve a four times voltage gain. The choice of the rectifier depends on the application, e.g., the expected input amplitudes and required DC voltage.

When the input voltage is much lower than the required DC voltage, multiple full-wave rectifier stages can be chained to further multiply the output voltage, which comes at the cost of diminished efficiency. Due to the forward voltage drop of real diodes, the efficiencies of rectifiers vary with the input waveform's amplitude. When the input voltage is low compared to the diode forward voltage drop, the diode has a considerable resistance that reduces the efficiency of the rectifier to just a few percent. With high input voltages, however, the efficiencies can be well above 90%. Therefore, a diode with a low forward voltage can improve the harvesting efficiency at low voltages by up to 800% compared to a diode with a high forward voltage drop [YCL18].

2.4.3 Power Management

As the harvested power in WPT systems depends on parameters that are outside of the control of the system designer, e. g., the distance and the wireless channel gain between transmitter and receiver, the harvested power is often less than the consumed power of the load, e. g., a sensor node. However, even if the harvested power is insufficient, the sensor node can still provide some service when it performs energy-aware power management with a Power Management Unit (PMU).

A typical option is to run the sensor node in a *duty cycle*, i. e., to adapt the interval in which the sensor node performs its task to match the average consumption to the supplied power. Therefore, the sensor node remains in a sleep state—or even fully powers down—while the harvester charges an energy buffer up to a threshold voltage, such as a capacitor or supercapacitor. Once the threshold is reached, the PMU activates the sensor node, which then executes one task cycle, e. g., it reads a sensor and transmits the result. The energy buffer's capacity and threshold voltage should be large enough to supply at least one task cycle to avoid the node running out of energy before finishing a task.

A duty cycle could theoretically enable a sensor node to perform a task with any nonzero amount of harvested power, even if the task intervals become very large. A practical limitation, however, is the leakage current that exists in storage elements and harvester circuits. Hence, the harvested power must be larger than the leakage power to load the storage element up to the threshold voltage.

2.5 Elastic Waves in Solids

In solid media, multiple types of waves can propagate, having distinctive properties, such as wave speed and attenuation coefficient. These types are characterized by propagation direction and the direction of particle movement. Selecting the right wave mode for the application is crucial, and several factors influence how a specific mode can be excited.

Bulk waves are longitudinal waves, where the particle movement is in the same direction as the wave propagation. They also exist in fluid media and gases and propagate by alternating compression and expansion of the medium. In contrast, the transversal waves propagate perpendicular to their particle movement direction. Transversal waves require a shear coupling between neighboring particles in the material. Hence, they predominantly occur in solid media.

Specific superpositions of bulk and transversal waves in thin, plate-like structures form guided lamb waves [Su09, p.19]. Two basic modes can occur: symmetric and antisymmetric. Additionally, an infinite number of higher orders of these basic modes may occur in higher frequency ranges. Figure 2.15 illustrates the described wavetypes.

The properties of the wave types vary and impact their usefulness for ultrasonic power and data transfer. While bulk waves can efficiently be excited in a wide range of frequen-

2 Fundamentals




Figure 2.15: Non-exhaustant overview of wave types emerging in solid media. Every type has specific properties that can aid or harm the application.

cies, they experience more attenuation than lamb waves, making them unsuitable for longer distances. In contrast, lamb waves travel long distances with less attenuation and lower propagation speeds. However, not all lamb waves propagate equally. Attenuation also increases with frequency, rendering the basic lamb-wave modes more efficient than the higher-order modes [Sun+21]. Furthermore, the antisymmetrical modes are attenuated less than the symmetric modes. Hence, the basic antisymmetrical mode is the wave mode with the lowest attenuation.

The generation of a specific wave mode and order depends strongly on the excitation frequency. The wavelength should be larger than the structure's thickness to excite predominantly the most efficient guided lamb waves. When higher frequencies are chosen, more bulk and higher-order lamb waves occur [Ros14, p.3], which do not travel as far in the metal. Consequently, the carrier frequency is limited to a low frequency range imposed by the structure's thickness for efficient power transfer.

State of the Art


This chapter first summarizes the targeted applications for acoustic power transfer and communication and analyzes the design requirements for these applications. We then briefly summarize the history of backscatter communication and introduce the fundamental principles and nomenclature. While electromagnetic backscatter is well established, acoustic backscatter is a relatively new field and has not been studied intensively. However, it adopts many techniques from Radio-Frequency Identification (RFID). Therefore, we start by summarizing the current state of RFID, then progress to more recent developments and adaptations in acoustic backscatter. Finally, we provide an overview of acoustic through-metal communication, focussing on active communication and power transfer required for fully batteryless embedded sensor nodes.

3.1 Applications and Requirements

Acoustic power transfer and communication in metals can supply and read out devices in places that are out of reach for conventional, electromagnetic-wave-based techniques. Examples are sensors embedded in a pipeline [Tse+22; SN22; ACG2]; TSM18; YS16; GNS11], in pressurized tanks and nuclear-waste containers [Zha+19; All21], into the reinforced concrete reactor chambers in nuclear power plants [Hei+19], and cargo containers [Alg+19; Hos+10; RMK09]. Others mention also reading out sensors through thick steel structures, such as ship hulls and bulkheads [Pri+07]. Using ultrasonic waves in these applications is an alternative to drilling holes and feeding cables directly through the metal, which compromises the structural integrity and may be costly and difficult to seal in pressurized tanks.

A second class of applications is Structural Health Monitoring (SHM) of bridges and wind turbines [TMÖ22; Sun+21; STM20], airplane hulls [Xu+22; Tse+22; STM20; Mäi+19; TBL18; TSM18], and similar structures, where conventional electromagnetic communication may be a viable alternative. However, the use of acoustic waves has advantages:

- **Hardware reuse:** Many SHM techniques use ultrasonic-guided waves to detect cracks and other damages in the structure [Ric+22; Lam+17]. The same transducers,



■ **Figure 3.1:** Ultrasonic power transfer and communication is promising for monitoring diverse structures, e.g., nuclear power plants, pipelines, aircraft and ship hulls, and wind turbine blades.

measurement circuits, and signal generators can also serve for power transfer and communication, saving costs and space for additional communication hardware.

- **Long-range power transfer:** Wireless power transfer with electromagnetic waves is inefficient over long distances and usually covers only a few meters. Metal structures can serve as a waveguide for ultrasonic waves, transmitting power more efficiently over longer distances.
- **Privacy and security:** Attackers can overhear electromagnetic communication, insert false data, or jam the medium from the outside. Physical access to the structure is required to achieve the same attacks in an ultrasonic through-metal communication scheme, making it easier to protect critical systems.

Monitoring applications often require large amounts of sensors distributed over the structures. Hence, economic viability introduces the following requirements.

1. Sensor nodes must be *cheap and small* to be distributed flexibly and in large quantities.
2. In many applications, sensor nodes are placed in hard-to-reach locations, making regular battery replacement difficult, if not infeasible. Hence, *wirelessly powered* nodes are desirable.
3. **WPT** requires a power emitter that is more power-consuming, complex, and costly than the sensor nodes themselves. Establishing *long-range transfer* enables a

single emitter to power several sensors and reduce the number of required devices significantly.

4. As the harvestable power shrinks with the distance from the emitter, a sensor device with less power consumption can power up over a longer range. Hence, the sensor devices should also be *ultra-low-power consuming*.

These requirements motivate the investigation of backscatter communication, which helps to achieve two goals at once: It enables ultra-low-power consumption at the transmitting sensor node, since it does not actively generate a waveform for communication, and it enables simple and cheap circuitry by omitting the hardware for active signal generation.

3.2 Backscatter Communication

Conventional wireless communication relies on a sender actively generating a signal to embed information. To propagate well in the targeted medium, e.g., in air, these signals must be in a suitable frequency range and have sufficient power to be detectable at the receiver. Consequently, the sender requires active circuitry and a power source to generate the high-frequency carrier wave. Opposed to conventional communication, backscatter is passive, i.e., the sender does not generate the carrier wave, but relies on an existing carrier generated by a different entity. To transmit information, the sender modulates the reflection of the incident wave. The receiver can then detect the modulation in the reflected signal and decode the embedded information.

The concepts of sender, transmitter, and receiver are somewhat ambiguous in the context of backscatter communication, as the backscattering device is simultaneously receiving a carrier wave, but sending data. Therefore, different terminology has evolved in the community. The active device, which generates the carrier, is termed the *signal source*, and the backscattering device is termed a *tag*. The *reader* receives the backscattered signal from the tag. The reader often includes the signal source and even shares the same antenna (or transducer) for carrier generation and signal reception, which is called a monostatic backscatter configuration. Communication may occur in both directions, where the communication from reader to tag is called *downlink*, and the tag to reader communication is the *uplink*. The setup and terminology are illustrated in Fig. 3.2.

Backscatter is often referred to as an ultra-low-power communication paradigm. However, it is crucial to understand that only the passively communicating tag nodes fulfill this promise. When comparing the energy per transmitted information bit, the most energy-efficient backscatter schemes communicate at less than 10 pJ bit^{-1} —more than two orders of magnitude less than Bluetooth Low-Energy (BLE) or WiFi [Ensl6]. At the same time, typical backscatter readers transmit carriers with up to 1W power, about three orders of magnitude more than WiFi and BLE transmitters. This imbalance in




Figure 3.2: General setup and terminology of a monostatic backscatter configuration. The signal source and receiver can also be separate devices (bistatic configuration).

power consumption stems from the low efficiency of wireless power transfer, and the total power consumption of readers and tags will be much higher than in conventional wireless sensor networks. Hence, backscatter networks are advantageous when a) a single reader serves many tags, b) tags are located at inaccessible places where regular battery replacement is infeasible, or c) tags must be cheap, disposable devices.

The earliest reports of backscatter communication date back to the 1930s, when fighter planes intentionally flew specific maneuvers to modify their reflection of radar signals. Allegedly, the radar ground stations evaluated the systematic modulation to discriminate the planes from their own side from others [Dob08]. A more formal foundation for backscatter communication was given in 1948 by Harry Stockman, who already foresaw various applications and backscattering of different types of waves, i.e., RF, microwaves, visible light, and sound waves [Sto48].

3.2.1 Radio-Frequency Identification

Today, the most prominent and widely adopted backscatter system is RFID, which was designed to identify objects wirelessly. In this sense, RFID tags replaced the conventional bar code, which is reflected in the name of the dominating RFID standard, called *Electronic Product Code* (EPC global) [EPC18]. Identifying applications include warehousing, where cheap RFID tags are attached to goods in stores and warehouses, which can quickly be accounted for with a reader. Other applications are toll collection—where tags are attached to vehicles and identify themselves when passing a reader at a toll station—and mobile payment, keycards, and many more. Typical RFID tags have a simple structure consisting of an antenna, a rectifier to harvest power from the incident carrier, and a load modulator to alter the reflection coefficient of its antenna. Only simple logic is required, e.g., to store their ID and execute the communication protocol, reducing the cost of such tags to several eurocents [Her+19].

With the rising interest in wireless sensor networks, RFID has been repurposed to fulfill wireless sensing tasks [KMI21; Sch+22]. The tag antenna is directly connected to a sensing

element that changes its impedance or resonance frequency with the observed quantity, e.g., a thermistor. The observed quantity thereby modifies the antenna's reflection coefficient and allows the reader to estimate the quantity. More relevant for the goal of this dissertation, *computational RFID* introduces a new type of wireless sensor node that uses the RFID physical layer to extend battery life by using backscatter for ultra-low-power communication or harvest incident energy to ditch batteries altogether. These nodes conform to the EPC global standard and are readable by commodity RFID readers. In contrast to classical RFID tags, however, they employ logic that exceeds the original purpose of identification but enables complex sensing tasks, including the readout of digital sensor ICs and processing samples before transmission.

A popular example of computational RFID is the open-source Wireless Identification and Sensing Platform (WISP), a programmable, microcontroller-based tag constructed from discrete commercially available components, extendable with arbitrary sensors. In contrast to most other RFID tags, the specifications and internal design of the WISP nodes are publicly available, and many researchers have reported on the capabilities and limitations of the platform. As a consequence, the platform has become hugely popular and was investigated or used in more than 130 publications in renowned conferences and journals [23a]. WISP's flexibility, however, comes at a cost: The more complex logic and general-purpose hardware require significantly more power than application-specific circuits (ASICs). The higher power consumption translates to a shorter readout range, as more power must be available over WPT to turn on the device. Table 3.1 compares the WISP platform with a System-on-Chip (SoC) based research prototype developed to wirelessly read out analog sensor values, e.g., a temperature sensor, and a commercially available sensor logger tag IC. The general-purpose MCU in the WISP requires roughly 850 μW when active, up to 25 times more than the ASIC-based platform requires. Consequently, read ranges are much shorter. As the required power is rarely available from the source, the WISP usually works in a duty cycle, which may cause missing packets from the reader, complicating communication.

3.2.2 RFID Physical Layer

The vastly different properties of reader and tag require different communication schemes for uplink and downlink communication. For brevity, we limit the discussion to the solutions adopted in the EPCglobal Class 1 Generation 1 (C1G1) and Class 1 Generation 2 (C1G2) standards. Analogously, we use the WISP platform as a case study for an implementation that fulfills this standard.

Downlink

Two requirements govern the downlink communication scheme's design. First, the tag must be able to demodulate and decode the message with minimal resources, i.e., memory, computational complexity, and energy. Second, the transmitted signal must

	WISP [332]	SOCWISP [100+3]	Commercial [54+10]
Application	Flexible	Analog Sensor readout	A/D Sensor Logging
Processing Hardware	MCU (MSP430)	130nm CMOS SoC	ASIC (SL900A)
Min. Startup power	-0.5 dBm	-12 dBm	n.a.
Read range	2.2 m	>3 m	up to 4 m
Active Power	850 μ W	33 μ W	150 μ W
Sleep Power	2 μ W	n.a.	< 1 μ W
Duty Cycling	✓	✗	✗

Table 3.1: Comparison of SoA computational RFID tags based on different hardware platforms.




Figure 3.3: The downlink in CIG1 and CIG2 employs ASK-based modulation schemes with different encodings. The tag demodulates the signal with a diode-based envelope detector.

have sufficient power to supply the batteryless tag with energy during reception. Existing RFID protocols usually only modulate the envelope of the single-frequency carrier signal using binary ASK to conform with the first requirement.

Different forms of encodings are possible with ASK modulation (see Fig. 3.3). The CIG2 standard, for example, uses Pulse-Interval Encoding (PIE), in which the interval between two pulses with low-amplitude level encodes a binary one or zero, and the C0G1 standard uses a pulse duration modulation (PDM) where the duration of a low-level pulse encodes the bit. Such pulse-based encodings are helpful to the second requirement, as they ensure that every symbol contains a high pulse, where the tag can harvest some energy. A pulse in every symbol interval guarantees a minimum average signal power throughout the message, considerably reducing the data dependency of the received energy at the tag compared to pure OOK. The maximum symbol rates specified in the standards go up to 40 kbit s⁻¹ (CIG2), and 70 kbit s⁻¹ (CIG1).

ISI is insignificant at data rates used in the EPC protocols because the symbol duration is much longer than the channel delay spread. Experimental studies report RMS delay

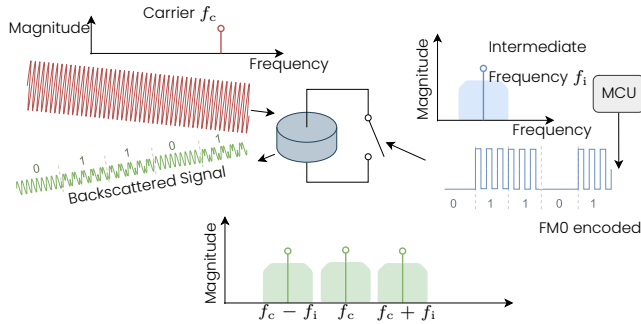


Figure 3.4: The tag uses FM0 modulation, shifting the transmission band away from the carrier. The receiver can then isolate the backscattered signal by filtering out self-interference.

spreads for typical backscatter channels of 10 – 80 ns [AMW12]. Hence, even at the maximum data rate of 70 kbit s^{-1} , the symbol duration is about 180 times longer than the RMS delay spread. The low ISI enables tags to employ simple and efficient demodulation schemes based on analog comparators.

Uplink

The uplink communication is subject to different constraints. The tag can only modulate its reflection coefficient, and most implementations switch between two different reflection coefficients to modulate the backscattered signal. The reader, however, has many more resources at its disposal for demodulation and decoding. The main challenge is that the backscattered signal from the tag is typically weak, and the reader's *self-interference* is often orders of magnitude stronger. Self-interference describes the remnants of the generated carrier, which spill over to the receiver circuitry. Additionally, the reflections of the carrier that are not reflected from the tag but from the environment also contribute to self-interference.

The tag in C1G1 and C1G2 uses FM0—a special type of Frequency Shift Keying (FSK)—to modulate its reflection coefficient. In FM0, the tag varies the frequency at which it switches between its two distinct reflection coefficients so that the switching frequency is zero in one state, and non-zero in the other state. In the latter state, the incident carrier signal is mixed with the switching frequency, shifting the passband away from the carrier frequency, which enables the reader to remove its self-interference and isolate the backscattered signal through filtering. Data rates from tag to reader reach up to 640 kbit s^{-1} in C1G2 and 140 kbit s^{-1} in C1G1. Since the reader has much more resources, it can perform active filtering and amplification. Therefore, the reader can receive backscattered signals with lower SNR than the tag at the downlink. The limiting factor

for passive tags is not the communication range but insufficient transmitted power to support the tag's circuitry. Battery-assisted tags, however, may communicate passively with the reader over much longer distances.

Power Transfer

The standard RFID tag is entirely passive, i.e., it harvests the power to run their circuits purely from the received carrier. Therefore, the carrier is rectified, and the voltage is regulated to be suitable to supply the integrated circuits. For power harvesting, the input impedance of the harvester plays a significant role because large mismatches with the source's output impedance result in inefficient harvesting and reduce the read range of the device.

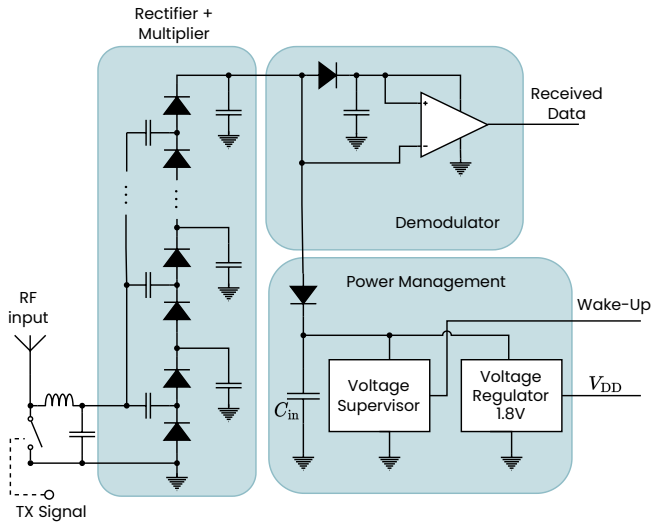
Furthermore, the input capacitance must be chosen carefully. A large capacitor requires a long time to charge up, during which the tag is unavailable to receive downlink messages. However, the capacitor has to be large enough to buffer periods during which no energy is harvested. These periods happen during the downlink, when the reader modulates the carrier's envelope, and during the uplink, when the tag short-circuits the antenna to change its reflection coefficient. Hence, the RFID modulation schemes are designed to ensure the availability of at least 50% of the power of an unmodulated carrier during reception and transmission. Furthermore, message durations in up- and downlink communication are short ($< 100 \mu\text{s}$), requiring only small buffers.

The power transfer efficiency of typical RFID systems depends on many factors, e.g., the distance between reader and tag, antenna gains, and harvester efficiency. Hence, quantifying a typical efficiency in RFID is not possible. To estimate the minimum efficiencies from the literature, we assume that the reader uses the maximum transmit power allowed by regulations, which is 30 dBm. When the WISP requires a minimum startup power of -0.5 dBm , the transfer efficiency at maximum read range is -30.5 dB . For the SOCWISP, the minimum power transmission efficiency can even reach -42 dBm , i.e., less than 10^{-4} .

WISP Case-Study

Since the fundamental tasks of an acoustic backscatter tag are similar to RFID, we now closely examine how WISP implements the tag to receive downlink messages, transmit uplink messages, and harvest energy from the carrier. Figure 3.5 shows the analog frontend of a WISP tag. Directly at the antenna, an RF-switch—a switch integrated circuit designed to work at radio frequencies—can short-circuit the antenna. When closed, this sets the reflection coefficient of the antenna to -1 . The TX signal controlling it comes directly from the MCU on the WISP. If the switch is not closed, the signal from the antenna passes through a matching network constructed from an inductor and a capacitor to match the antenna's output impedance to the harvester's input impedance.


After impedance matching, the received carrier amplitude is rectified and its amplitude is multiplied with five consecutive stages of a full-wave rectifier. The rectified carrier



■ **Figure 3.5:** Schematic of WISPs analog frontend, that includes voltage multiplier, harvester, and demodulator (adapted from [SS13]).

then charges an input capacitor in the power management unit. In the PMU, a regulator generates the constant 1.8 V supply voltage for the MCU. However, the MCU directly goes into a deep sleep mode with ultra-low power consumption and waits for a wakeup signal from the voltage supervisor which observes the capacitor voltage and toggles a pin as soon as the input capacitor reaches 1.9 V. Since the MCU often consumes more power than the carrier can supply, it drains the capacitor until the MCU switches back into deep sleep at about 1.8 V. As a result of this behavior, the voltage on the input capacitor fluctuates between 1.8 – 1.9 V. Hence, the harvester's input impedance varies only slightly throughout a power cycle.

The demodulator in WISP uses the rectifier's output as envelope of the signal. A passive lowpass filter also tracks the average envelope over many symbol periods on one branch. Both signals—the envelope and the multi-symbol average—are compared with an analog comparator. The MCU can then detect the edges in the comparator output instead of continuously sampling the envelope. This way, no ADC is required, and the MCU can sleep and wait for pin interrupts most of the time, significantly reducing its active power consumption. Similar techniques are described for other state-of-the-art RFID tags based on ASICs and SoCs [Cur+07].



■ **Figure 3.6:** To achieve higher-order modulation, the backscatter tag switches between different load impedances, which produce reflection coefficients in a desired constellation.

3.2.3 Beyond the RFID Standards

The previous section described the widely adopted RFID standard, which achieves uplink data rates up to 640 kbit s^{-1} . While these data rates are sufficient for most applications, the research community investigated ways to increase achievable data rates to enable future applications. We briefly discuss two of these approaches—higher-order modulation and equalization—which we adopt for acoustic backscatter in this dissertation.

One way to increase the achievable data rate is to increase the modulation order, i. e., transmit more than one bit per symbol. Thomas et al. demonstrated a semi-passive backscattering tag that employed 16-QAM modulation, reporting a data rate of up to 96 Mbit s^{-1} [TR12]. Instead of switching only between two impedance states, the tag implements higher-order modulation by switching between up to 16 load impedances, thereby producing 16 distinct reflection coefficients. Carefully choosing the load impedances in correspondence to the antenna's characteristic impedance, the reflection coefficients can form a regular QAM constellation as in conventional active communication (see Fig. 3.6). Thomas reports that symbol rates were limited to less than 12 MHz because ISI becomes relevant at such rates in typical indoor environments, and the RF switches employed to switch between load impedances have finite transition times in the range of 25 ns.

To further increase symbol rates beyond a couple of Mbit s^{-1} —at which ISI becomes relevant—equalization has also been considered for RFID systems, similar to active high-speed communication networks. However, only very few studies tackled equalization in RFID, and those relied purely on simulation results. For example, Fawky et al. [Faw+14] employed the well-known zero-forcing equalizer to increase the robustness of tag detection in environments with multipath characteristics, while Lunglmayr et al. [LH10] showed that a least squares equalizer reduces bit error rates while backscattering at 6.7 Mbit s^{-1} .

3.2.4 Piezoacoustic Backscatter

In environments where electromagnetic waves are unfavorable, communication with acoustic waves may be an alternative. For example, acoustic communication is an established technique under water [CSR23]. Accordingly, backscatter schemes similar to RFID have been developed recently with acoustic waves in different environments. Instead of an antenna, piezoelectric transducers, called hydrophones in underwater applications, transform electric signals into acoustic waves and vice versa. However, despite the different physical wave and transducer types, the same fundamental principle holds: A mismatch between the transducer's characteristic impedance and the connected circuits causes a reflection of the incident wave. Hence, similar circuits have been adopted for wireless energy and data transfer.

Underwater

Underwater backscatter was first demonstrated in [JA19] to wirelessly power and read out MCU-based sensors submerged in shallow waters. The most prominent difference to RF backscatter is the narrow bandwidth in underwater communication. Only low frequencies, often in the tens of kilohertz, can be used because higher frequencies experience severe attenuation in water. Furthermore, the piezo transducers have a narrow usable frequency range around their resonance frequency. In the discussed study, only a frequency range between 15 and 18 kHz was employed. Because of the limited bandwidth, the achieved uplink data rates using frequency modulation were below 3 kbit s^{-1} .

Follow-up work mitigated the bandwidth limitation in two ways: First, Afzal et al. [Afz+20] increased the data rate despite the narrow bandwidth with higher-order modulation. The different constellation points were achieved with the same technique that Thomas et al. implemented for RF backscatter, which we discussed in Section 3.2.3. The demonstration was limited to a 4-QAM modulation, effectively doubling the data rate compared to binary modulation. A different approach presented in [Gha+20] widened the narrow transducer bandwidth by combining several piezo crystals with different resonance frequencies joined by a coupling material. The presented technique successfully increased UW backscatter bitrates to up to 20 kbit s^{-1} with conventional binary modulation.

The power transfer capabilities of UW backscatter systems are limited by the strong attenuation of acoustic waves underwater. This is well-described in [Afz+22], where a transmit power of 44 dBm—25 times more power than the maximum allowed transmit power of RFID readers—was required to power a wireless sensor node over several meters. The authors do not mention the power transfer efficiency directly, but the study suggests efficiencies as low as -50 dB .

Other Propagation Media

The properties of backscatter communication and wireless power transfer are also favorable for sensors implanted deep in the human body, which consists of more than 60 % of water and therefore strongly attenuates electromagnetic waves. The attenuation of a 2.45 GHz RF link in the human body is up to 70 dB higher compared with a 1 MHz ultrasonic link [Gui+19]. Furthermore, battery replacement requires surgery, which comes at high costs and risks for the patient. Instead, a growing field of research of powering and reading out such sensors with piezoacoustic backscatter has evolved [GM20; Web+18]. In contrast to the underwater environments, higher carrier frequencies and larger bandwidths can be achieved: Attenuation is less prohibitive because the transmission distances are only about ten centimeter. For example, [Web+18] reports a link loss between $0.5 - 1 \text{ dB cm}^{-1}$ at a carrier frequency of 790 kHz, which is in line with the results reported by [Gui+19]. Other than the transmission channel, however, the circuits and methods for WPT and backscatter communication are similar to the underwater backscatter systems. ISI was not reported to be a problem, reaching data rates up to 95 kbit s^{-1} for the uplink and total power transfer efficiencies in the range of 1 – 2 %.

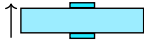

3.3 Through-Metal Backscatter Systems

Acoustic wave propagation in solids is—in contrast to underwater—more complex. Waves in solid media propagate not only as longitudinal (pressure) waves but also as transversal and surface waves. Accordingly, reviewing the literature about acoustic through-metal communication and power transfer demands separating existing studies into two categories. We distinguish between *Sandwich-Plated* (SWP) channels and Guided-Wave (GW) channels. SWP channels consist of two axially aligned transducers on opposing sides of a metal barrier and primarily rely on longitudinal waves, while GW channels consist of two transducers at arbitrary locations on a thin metal structure (see Fig. 3.7) and rely primarily on transversal waves. We provide an in-depth experimental analysis of the two fundamental channel types and their properties in Section 4.2.

To the best of our knowledge, no prior work employs backscatter communication over GW channels. We, therefore, focus our review on the two building blocks necessary for our purpose, i.e., active communication and wireless power transfer in GW channels. Furthermore, we review existing passive communication in SWP channels as well.

3.3.1 Active Communication through GW Channels

A rich background of literature about active communication through GW channels exists. However, before we can contextualize the different results of existing approaches, we must first discuss three fundamental characteristics of the GW channel: Attenuation, frequency ranges, and multipath propagation. If excited in the suitable frequency range, so-called lamb waves emerge, which are a superposition of several fundamental wave

$\approx \text{Mbit s}^{-1}$ High-Order Modulation	$\leq 50 \text{ kbit s}^{-1}$ Binary Load Modulation	SWP Bulk waves Moderate ISI	
$\approx \text{kbit s}^{-1}$ Binary Modulation	Not explored Focus of this dissertation narrowband-only	GW Lamb waves High ISI	
Active	Passive		

■ **Figure 3.7:** The sandwich-plated channel primarily uses bulk waves for power and data transmission, while the guided wave channel employs lamb waves. No research in passive backscatter communication yet exists for GW channels.

types. The GW channel enables efficient transmission of such lamb waves over longer distances in thin structures. For example, the authors of [Su09, p.42] state wave attenuation of guided wave amplitude in a 5 mm thick steel plate of only 1 dB m^{-1} . In general, the inverse square root law for wave decay applies, stating that the magnitude of a lamb wave decays with distance d proportional to the inverse square root of the distance, i.e., if measuring the amplitude $A(d)$ at two distinct distances d_1 and d_2 , we obtain

$$\frac{A(d_1)}{A(d_2)} = \frac{\sqrt{d_2}}{\sqrt{d_1}}, \quad (3.1)$$

equivalent to an inversely proportional decay of power.

To achieve guided waves with such favorable properties, a suitable excitation frequency range must be chosen, which is subject to the material parameters and the thickness of the metal structure in which the wave propagates. Several sources state that these lamb waves are best excited if the wavelength is in the same order of magnitude as the plate thickness [Ros14; Lam+17]. In many practical applications, that results in frequencies in the range of several hundred kilohertz. However, suitable frequencies may drop to tens of kilohertz for structures with several centimeter-thick walls. While low attenuation favors long communication distances, it also poses a challenge to communication in constrained structures, as waves reflect nearly entirely at the structure's boundaries (the admission coefficient of a steel-air barrier has a magnitude of roughly 10^{-5}) and lead to significant multipath propagation [Hua+18; Hei+2; JZY11]. Furthermore, lamb waves are dispersive, which means that their propagation velocity is a function of frequency, which additionally distorts a transmitted signal [SWH18; JYZ13].

These fundamentals help to contextualize the results of previous studies on active communication through GW channels, listed in Table 3.2. The range of achievable data rates is large, ranging from as little as 20 bit s^{-1} up to one megabit per second. Different assumptions and requirements in the designed systems cause these large variations. Communication schemes that can be received with simple ultra-low-power devices—as targeted for batteryless and backscattering sensors—have achieved data rates in the

range of a hundred bit s⁻¹. Tang et al. [TMÖ22] designed a custom CMOS integrated circuit for a batteryless SHM sensor node. The presented system consisted of a high-power reader that generates a waveform for power transfer and low-power batteryless sensor nodes. Because of the different limitations of the two device types, the up- and downlink of both devices use different schemes, where the uplink achieves much higher data rates than the downlink because of the limited decoding capabilities of the batteryless nodes. However, no backscatter modulation was used.

All sources attribute the low data rates to the strong ISI caused by multipath propagation and dispersion. Therefore, many researchers tried to mitigate ISI with various techniques, partly achieving speed-ups of four orders of magnitude. For example, Huang et al. [HS22] applied OFDM modulation to achieve up to 2 Mbit s⁻¹ data rates over 70 cm through an aluminum plate. However, receiving OFDM is infeasible for batteryless sensor nodes, requiring complex filters and high sampling rates to distinguish sub-channels at multiple carrier frequencies. Furthermore, OFDM requires the tag to generate and mix multiple sub-carriers, demanding power-hungry active hardware.

Other approaches record the impulse response of the specific channel and correlate the received signal with the expected impulse response to detect a distorted pulse or pre-distort the transmitted signal with an impulse response tailored to the specific channel, so that the the received signal is distortion-free at. These techniques, however, were only demonstrated with offline processing using expensive, bulky, and power-hungry lab equipment to generate and acquire signals. Transferring the correlation approach to the resource-constrained tag devices is not straightforward, as sampling rates and filter complexity are strictly limited by the device's capabilities. Furthermore, a pre-distortion filter would require the tag 1) to know the impulse response of the channel from tag to reader, and 2) to modulate the backscattered signal arbitrarily, which is not feasible with known backscatter techniques.

Tang et al. showed promising rates up to 15 kbit s⁻¹ in the uplink without complex ISI mitigation methods. In this study, the batteryless node generates short rectangular pulses to encode the data, resulting in a wideband transmission signal that proves to be more resilient against ISI. Even higher data rates up to 2 Mbit s⁻¹ were achieved by Shaik et al. [STM20] with wideband pulses. Unfortunately, the study reports no details about the receiver implementation and its complexity. The much higher data rates in these wideband studies are also possible, because the presented reader devices use a piezo array to isolate the line of sight path from other reflections, mitigating ISI significantly.

In total, communication distances in most studies were small (below 1 m), but Mijarez et al. demonstrated that distances up to 30 m are achievable, although only with a very low data rate of 20 bit s⁻¹.

3.3.2 Power Transfer through GW Channels

The general feasibility of power transmission with acoustic waves through GW channels was proven in numerous studies. We summarize the results in Table 3.3 and discuss them

Study	Spec.	Thickn.	Distance	Freq.	Modul.	Datarate	ISI Mitigation ¹	Impl. Level
		mm	m	kHz		kbits ⁻¹		
High complexity receiver								
Tang, 2022 (Uplink) [TMO22]	Plate (CRFP)	2	< 0.3	WB ³	Pulse-OOK	15	✘	CMOS IC
Huang, 2022 [HS22]	Plate (Alum.)	8	< 0.7	2000	OFDM	1000	CCIR	Lab Equipment
Bahouth, 2022 [Bah+22]	Plate (Alum.)	6	0.08	200	BPSK	350	CCIR	Lab Equipment
Sun, 2021 [Sun+21]	Plate (Alum.)	16	0.4	500	OOK	100	CCIR & MD	Lab Equipment
Shaik, 2020 (Uplink) [STM20]	Plate (Alum.)	1.5	0.5	WB ³	Pulse	2000	✘	AF ² + MCU
Jin, 2013 [JYZ13]	Pipe (Steel)	4	1.5	250	PPM	10	TR	Lab Equipment
Low complexity receiver								
Tang, 2022 (Downlink) [TMO22]	Plate (CRFP)	2	< 0.3	400	BFSK	0.2	✘	CMOS IC
Shaik, 2020 (Downlink) [STM20]	Plate (Alum.)	1.5	0.5	250	BFSK	0.08	✘	AF ² + MCU
Mjarez, 2013 [MG13]	Pipe (Steel)	4	30	40	PPM	0.02	✘	DSP
Chakraborty, 2015 [Cha+15]	Pipe (Steel)	18	4.8	455	Chirp-OOK	0.1	✘	AF ² + MCU

¹ **CCIR**: cross-correlation with impulse response, **MD**: mechanical damping, **TR**: time reversal

² **AF**: Custom Analog Frontend

³ **WB**: Wideband

Table 3.2: Summary of existing studies on lamb-wave-based active communication studies.

subsequently. Unfortunately, these studies cannot be compared (among each other) because they used vastly different setups—i. e., transducers, frequencies, impedance matching circuits, and evaluation scopes—yielding output powers ranging from less than 100 μW to 470 mW. In the following paragraphs, we briefly explain the critical design considerations in which the existing studies differ.

The chosen transducer impacts the performance of the power transfer significantly. While the dimension and material determine the resonance frequency, its shape influences the propagation direction. This has several implications:

- **Directionality:** Disk-shaped piezoelectric transducers generate a circular wave, propagating in all directions uniformly, whereas rectangular transducers may generate directed waves. Directed propagation can increase efficiency [Tse+20] but requires careful transducer alignment and severely complicates applications in which multiple tags are powered by a single reader. Beamforming with transducer arrays enables directionality and selective charging of multiple transducers [STM20]. However, transducer arrays come with increased size and add implementation complexity.
- **Resonance:** The carrier frequency choice depends on the transducers and the metal structure. Transducer bandwidth limits the range of frequencies for efficient power transmission. At the same time, certain frequencies excite specific types of waves in the metal structure, of which asynchronous lamb waves are the most efficient for power transfer in thin metal structures, especially over long distances [Xu+22]. In any case, the transducers must be chosen appropriately so that their resonance matches the structure's resonance.

Impedance matching at the transmitter and receiver is critical to achieving high efficiency; e. g., Xu [Xu+22] reports a 6.75-time increase in output power when using an ideal complex impedance matching circuit at the receiver's transducer. Unfortunately, in GWCs, the transducer's input and output impedances are coupled to the structure's mechanical impedance and change strongly even when the carrier frequency changes only by a few hundred Hertz. As the impedance is not known before transducer deployment, a dedicated impedance matching circuit would require measuring the transducer impedances after deployment and adapting the circuits appropriately, therefore limiting the ease and practicability of deploying such systems. Therefore, flexible MPPT techniques are useful to dynamically match the harvester's input impedance to the source [TMÖ22; STM20]. MPPT is a significant addition in through-metal power transfer in contrast to established RFID and underwater techniques, where the coupling between the transducer and its environment is much weaker, and the designer assumes a fixed transducer (or antenna) impedance to match the harvester impedance with.

Finally, the total output power and efficiency reported in the studies encompass different scopes. Some studies measured the AC output power at the receiver's terminals, considering only the losses within the metal channel and transducers. Others also include

Study	Dist. [m]	Freq. [kHz]	Output Power [mW / dBm]	Efficiency [% / dB]	Impedance Match	Scope ¹
Xu, 2022 [Xu+22]	0.4	150	1.537 / 1.8	0.217 / -26.6	Ideal	DC
Tang, 2022 [TMÖ22]	< 0.3	400	0.075 / -11.2	n.a.	MPPT	DC
Sun, 2021 [Sun+21]	0.4	150	3.18 / 5.0	n.a.	Ideal	AC
Tseng, 2020 [Tse+20]	0.24	25	470 / 26.7	56 / -2.5	Ideal	AC
Shaik, 2020 [STM20]	0.5	250	0.181 / -7.4	n.a.	MPPT	DC
Kızıroglu, 2017 [Kiz+17]	1	47.5	18 / 12.6	0.1 / -30	Resistive	AC
Kural, 2014 [Kur14]	0.54	220	17 / 12.3	30 / -5.2	n.a.	AC

¹ DC: Efficiency was measured, including rectification and voltage regulation.

Table 3.3: Summary of existing studies on lamb-wave-based wireless power transfer through metal channels.

the rectifier and energy-harvester circuit, which are required to supply a sensor node. Especially when the available power at the source only provides several hundreds of microwatts, rectifier and regulator efficiency is poor and limits the total available power critically. Therefore, the investigation must consider the entire transfer chain from the generator to regulated DC power at the receiver. Accordingly, studies that reported very high transfer efficiencies, e.g., 56% [Tse+20] and 30% [Kur14], only considered the efficiency of the AC power transmission. However, considering only the channel loss, the potential for energy transfer in GW channels is large compared to RFID, where efficiencies in the range of -30 dB, and underwater power transfer down to -50 dB, are common.

Existing studies only investigated short channels of less than 1m length. However, since the power of guided ultrasonic lamb waves decays only inversely proportional to distance (see Eq. (3.1)), we can expect that much longer transfer distances are achievable. Furthermore, in contrast to the tight regulations of the electromagnetic spectrum, no regulations apply to the transmit power of lamb waves in metal structures, enabling to scale up transmit power and achieve even further transmission distances.

Electromagnetic Through-Metal Power Transfer

Electromagnetic wave-based wireless communication and power transfer techniques are rarely applicable, e.g., only for non-ferrous metals [Yan+15]—excluding ordinary steel used in many structures. Furthermore, successful prototypes using inductive or magnetic resonance coupling to transmit power through non-ferrous metal barriers had to apply frequencies below 100 Hz to achieve acceptable efficiencies. Accordingly, the required coils were large and heavy to resonate at such low frequencies, e.g., in [Yam+14], the coils on the transmitter and receiver weighed roughly 3 kg each and were more than 40 cm long.

Alternatively, some promising approaches have recently demonstrated the feasibility of using electromagnetic guided waves propagating at the surface of a metal plane to transmit power with high efficiency [Bie+21]. Due to the high conductivity of most metals, electromagnetic waves enable high efficiencies. However, the reported approaches required large and bulky devices in the order of 40 x 30 cm to excite frequencies with resonate wavelengths.

3.3.3 Sandwich Plated Channel

Sandwich-Plated (SWP) channels use bulk waves for data and power transmission, which have different characteristics compared to lamb waves. They can be excited efficiently in higher frequency ranges, therefore providing wider bandwidths for communication. Furthermore, bulk waves are attenuated much stronger, reducing the impact of multipath propagation. Hence, the channel gain and transducer impedances are more predictable and depend primarily on the material parameters and the structure's thickness.

Consequently, systems with several Mbit s⁻¹ data rates [Cun+16b; Din+20; Sho+07], and power transfer systems with DC-to-DC transfer efficiencies of up to 68% [Ali+23] have been demonstrated. More relevant for this dissertation, batteryless systems integrating backscatter communication have also been investigated for SWP channels. In this community, backscatter communication was termed *reflected power*, but the underlying principle of modulating the reflection coefficient on the transducer is the same [Ash+13; Din+20].

While these advantages make designing highly performant and reliable batteryless sensor nodes easier, the required coaxial alignment and the strong attenuation in SWP channels prohibit transmission distances exceeding several centimeters, which limits the range of applications where these systems can be used.

Goals and Preliminary Analysis

In light of the state-of-the-art, this chapter first determines the goal of this dissertation. We identify key challenges and summarize how our contributions help to meet these challenges. We then present a preliminary analysis of the characteristics of the through-metal channel used for evaluation in the following chapters.

4.1 Dissertation Goals

Our overarching goal is to enable sensor networks to extend into difficult-to-reach regions shielded by metal, as required for applications presented in Section 3.1. The state-of-the-art overview has shown that we can distinguish two general channel types, *GW* and *SWP* channels, and *SWP* channels were investigated more thoroughly and pose less of a challenge. However, *GW* channels are favorable for many applications, as they allow supplying several tags at different locations from a single reader and offer more flexibility in the placement of reader and tag in general.

Especially in the monitoring of infrastructure, which is often large and requires the distribution of many sensors over the structure, the economic costs of such a sensor network are critical. The sensing devices are, hence, required to be cheap, small, easy to deploy, and easy to maintain. The evolution of RFID demonstrated that wirelessly powered backscattering devices meet these requirements well: Backscattering requires minimal additional communication hardware, no active signal generation, and wireless powering enables batteryless devices that neither require costly cabling nor regular battery replacement. However, no investigation of backscatter communication in *GW* channels exists, leading to the first set of goals of this thesis.

- **Goal 1:** Investigate the feasibility and explore the limitations of backscatter (uplink) communication in *GW* channels. Limitations include the usable bandwidths, data rates, and communication distances.
- **Goal 2:** Analyze existing active (downlink) communication schemes for suitability with a resource-constrained tag and improve the achievable data rates.

4 Goals and Preliminary Analysis

- **Goal 3:** Characterize the potential of acoustic wireless power transfer and explore the tradeoffs between systems optimized for communication and power transfer.

The existing literature includes investigations of different components we require to achieve these goals, e.g., active communication schemes through GW channels and wireless power transfer. However, these investigations had several shortcomings.

- **Generality:** The existing studies often only considered a single physical transmission channel for their evaluation and tailored their experiment setup to perform optimally in this particular channel. However, to achieve easy deployment, we evaluate our approaches in multiple channels to enable a plug-and-play-like deployment without tedious manual setups and adaptations.
- **Realizability:** Many existing studies used bulky, expensive, and power-hungry lab equipment (signal generators, oscilloscopes, power amplifiers) to evaluate their approaches. More realistic low-cost hardware comes with additional challenges, such as inaccurate and unstable oscillators, which must be considered for practical implementations.
- **Comparability:** The existing studies are heterogeneous in the employed hardware, methods, and physical structures (size, geometry, material) they used for evaluation. Therefore, it is difficult to conclude which methods are superior to others and in what circumstances.

These shortcomings lead to the goals.

- **Goal 4:** Investigate the influence of non-idealities in practical implementations, e.g., clock mismatches and time-variant channels, and explore mitigation strategies.
- **Goal 5:** Construct a hardware prototype made from commercially available, cheap, and low-power components and evaluate different approaches with this hardware.

Lastly, we made parts of the constructed hardware, firmware, and software available for other researchers under an open-source license and plan to publish the remaining parts with this dissertation. Thereby, we hope to motivate further research in this field, similar to WISP, which inspired researchers to investigate computational RFID.

4.1.1 Contributions

In the following paragraphs, we outline the contributions this dissertation makes to achieve the formulated goals. A more detailed analysis of the addressed problems and proposed solutions follows at the beginning of each subsequent chapter.

Chapter 5 investigates the use of higher-order modulation to increase the achievable data rates despite the limited coherence bandwidth. Choosing suitable load impedances for the tag is not straightforward, so we construct a dynamic tag and

analyze the backscatter channel's transfer function. With this channel model, we develop and test an automated setup procedure to realize higher-order modulation more efficiently and without manual setup.

The backscatter channel is generally nonlinear, prohibiting the use of common equalization techniques to mitigate ISI. Chapter 6 analyzes the transient backscatter channel and develops a channel linearization approach that enables the reader to apply standard linear and decision-feedback equalizers. We thoroughly evaluate the speedups an equalizer can achieve in GW channels in combination with higher-order modulation and develop tracking mechanisms to adapt the equalizer to the time-variance found in the acoustic backscatter channel.

To maximize the efficiency of wireless power transfer, impedance matching between transducer and harvester is critical, which is not achievable before deployment in GW channels. Chapter 7 investigates WPT and the use of MPPT techniques to adapt the harvester impedance dynamically. Furthermore, it implements the tag on an ultra-low-power platform and incorporates dynamic compensation of clock offsets to mitigate the non-ideal oscillators on such low-power devices.

Finally, we investigate the downlink communication in Chapter 8, where wideband pulses are used to achieve higher data rates despite the constrained resources on the tag. We show that wideband pulses enable higher data rates with the same simple receivers as the conventionally used narrowband communication schemes. To mitigate the effects of clock offsets without adding further computational burden to the tag, we developed a tracking and compensation mechanism working on the reader.

4.2 Channel Analysis

The transmission channel is of utmost importance for wireless communication and power transfer. Hence, a detailed investigation of the properties of a wireless acoustic through-metal channel is fundamental to making appropriate design choices during prototype development and protocol definition. We consider both the propagation medium for the acoustic waves as the channel, e.g., the metal structure, and the piezoelectric transducers. Furthermore, intermediate layers, e.g., a layer of epoxy resin used to attach the transducer to the metal surface, affect channel characteristics and must be included in the investigation.

This section analyzes the transducers and metal structures used throughout this thesis. We look into the characteristics of the selected piezoelectric disks that we utilize as transducers. With a series of experiments on the primary specimen, on which we validate our solutions in the later chapters, we then illustrate the differences between *GW* and *SWP* channels and show why the former is more challenging than the latter. Furthermore, this section clarifies why we can not simply adopt existing techniques from RFID or other piezoacoustic backscatter prototypes in GW channels.




Figure 4.1: Voltage gain over two directly attached transducers.

4.2.1 Piezoelectric Transducer

Selecting a suitable transducer is a crucial step. The transducer geometry determines the directionality of the excited waveform. In this dissertation, we selected disk-shaped transducers because they do not require any alignment of reader and tag. The transducers have a thickness of 2 mm and a diameter of 10 mm. According to the manufacturer specifications [Pl 23], the disk has a radial resonance at 200 kHz and a thickness resonance at 1 MHz. The transducer’s resonance ranges must match the structural resonances to achieve optimal efficiencies, i.e., excite strong vibrations in the structure. Therefore, different structures require the choice of different transducers. We employed transducers with a wraparound electrode to enable plain contact with the metal surface.

Transducer Gain

To characterize the gain of the selected transducers in isolation, i.e., without a metal channel in between, we attach two transducers of the same kind directly with epoxy resin as depicted in Fig. 4.1 (left). We then excite one transducer with a 1V amplitude sinusoid and measure the emerging amplitude at the other transducer after convergence to a steady state. The excitation frequency sweeps over a range from 50 kHz to 10 MHz.

The channel’s voltage gain is shown in Fig. 4.1 (right). The specified radial and thickness resonances are well observable. The most pronounced resonance peak occurs in the radial resonance, with over 14 dB more than the thickness resonance. As typical for piezoelectric transducers, we observe narrow resonance peaks around some frequencies, making broadband applications challenging to realize. The transducer’s frequency response shows many minor resonance peaks in addition to the fundamental resonances. Partly, these are higher-order resonances at multiples of the fundamental frequencies. However, the wraparound electrode also complicates the frequency response because the active area of the transducer is no longer disk-shaped but has a non-symmetric shape.

Wireless communication can usually tolerate gains of -40 dB or less because receivers can filter and amplify small signals. Hence, the transducers are suitable for communication over a wide bandwidth. However, efficient wireless power transfer relies on high transducer gains. Hence, the usable frequency range for WPT is limited to narrow bandwidths around the resonances. Additionally, for power transfer and backscattering, the transducer's impedance is a crucial property we investigate next.

Transducer Impedance

The electrical impedance of a transducer depends on its electrical and mechanical behavior. Electrically, a transducer impedance combines the capacitor formed by the electrodes and the dielectric piezoceramic in between and the coupled acoustic impedance. Hence, the mechanical strain field influences the transducer's electrical impedance. At resonance, the magnitude of the impedance is at a local minimum, and the phase is zero. An antiresonance follows at a slightly higher frequency, where the impedance magnitude approaches a local maximum. In resonance and antiresonance, the transducer impedance is purely resistive.

We experimentally measure the current through a single transducer during excitement with a scope and signal generator. For this experiment, the transducer is not attached to any structure but is surrounded by air. With known excitation voltage, we can derive the transducer's impedance Z_T shown in Fig. 4.2. The results reveal that the impedance magnitude drops significantly at the radial resonance frequency to about $|Z_T| = 60 \Omega$. As discussed in Section 2.4, the lower the impedance magnitude, the higher the maximum extractable power of a source at a given amplitude. Therefore, the combination of the highest voltage gain and lowest transducer impedance magnitude makes the radial resonance frequency the best choice for power transfer. However, the results also show little leeway for varying the frequency because the impedance magnitude increases more than 10-fold when the excitation frequency deviates less than 5% from the resonance.

4.2.2 Metal Channel

When attaching a disk-shaped transducer to a thin metal structure, the vibration in the transducer excites bulk and lamb waves in the structure. The lamb waves propagate omnidirectionally along the structure, and reflections occur at the structure's boundaries. Caused by the large acoustic impedance mismatch between metal and the surrounding air, close to none of the wave energy is dissipated into the environment around the structure. Hence, upon continuous excitation, multiple reflections of the wave superimpose and form a standing wave, i.e., an interference pattern that can be constructive or destructive. The exact pattern depends on the material parameters and geometry of the structure and is highly sensitive to frequency and the exact location of the exciting transducer.




Figure 4.2: Impedance of a piezoelectric disc with wraparound electrode. At resonance, the impedance magnitude is at a local minimum, while the phase is zero, i.e., the impedance in resonance is purely resistive. The shunt resistor for current measurement was 1 k Ω .





Figure 4.3: A large specimen with five transducers is used to evaluate methods throughout this thesis (not to scale).

The test structure we used to validate our approaches in the remainder of this thesis is a flat steel bar, a common part in supporting structures, and a component of more complex elements such as I-beams. We attached five piezo disks to the bar to form different channels, of which one transducer pair creates an SWP channel as a reference. The dimensions of the bar and transducer locations are shown in Fig. 4.3. We use the following convention to conveniently reference the channels formed by a transducer pair in this test structure: The label $CH_{a \rightarrow b}$ references the channel from transducer a to transducer b .

Guided-Wave Channel

This section repeats the gain and impedance measurements for one of the GW channels on our test specimen. For brevity, we only show the measurements for channel $CH_{3 \rightarrow 2}$, which covers a distance of one meter from transmitter to receiver. However, the qualitative characteristics of all measured channels were similar. Figure 4.4 (left) displays the gain over a wide frequency range. We can also observe the two major resonance peaks around 200 kHz and 1MHz. However, the absolute gain is significantly lower compared to the pure transducer channel, i.e., the maximum gain around the radial resonance is about 15 dB smaller. The thickness mode resonance is attenuated



■ **Figure 4.4:** Gain of the GW channel $CH_3 \rightarrow 2$. The frequency band around 200 kHz (marked in the plot) excites guided waves best (left), but interference patterns create a strong frequency dependence (right).

more than the radial mode, leading to a larger difference between the two modes of more than 24 dB compared to the 14 dB we observed with only the transducers attached.

Furthermore, we observe a fissured frequency response without the pronounced peaks observed earlier. A close look at the radial resonance region in Fig. 4.4 (right) reveals a highly frequency-selective channel, with deep spectral nulls that fall below -40 dB gain. The frequency dependency is a consequence of the interference patterns in the GW channel and is sensitive to the exact position of the transmitting and receiving transducer.

To quantify the multipath intensity, we measure the RMS delay spread τ_{RMS} experimentally. We derive the frequency responses of all narrowband channels in the resonance region by downconverting the impulse response to the baseband with a selected carrier and use a fourth-order lowpass filter to limit the bandwidth of the resulting narrowband channel to 20 kHz. We then compute the RMS delay spread of each narrowband channel. The results depicted in Fig. 4.5 show that for the three tested GW channels, the delay spread was between 1.2 and 2 ms for all carriers. That relates to a coherence bandwidth of 500 – 833 Hz. Hence, the results show that ISI severely limits the data rate with conventional narrowband modulation schemes.

Finally, the transducer impedance plays a significant role in the design of a prototype. Attaching the transducer firmly to a metal structure couples the transducer's electric impedance to the structure's acoustic impedance. Due to the interference patterns, the structure's acoustic impedance is—similar to the gain—frequency-selective. We also observe this in the measured electrical transducer impedance (see Fig. 4.6). The impedance magnitude varies from less than 400 Ω to more than 1500 Ω within the observed range, often changing more than 50% within a few hundred Hertz. Since the acoustic impedance depends on geometry, material parameters, and exact placement of the transducers, the resulting impedance response differs between individual deployments. We cannot predict exact impedance values before deployment, e. g., during design time.




Figure 4.5: The channel RMS delay spread τ_{rms} and coherence bandwidth for narrowband channels with different carrier frequencies near the resonance.





Figure 4.6: Transducer Impedance when attached to a thin metal structure.

Sandwich Plated Channel

Finally, for comparison, we also investigate the gain and impedance of a SWP channel. Both bulk and guided waves can propagate in an SWP channel. If primarily guided waves are excited, the gain and impedance of the SWP channel behave similarly to the results presented in the previous section. However, Section 2.5 discussed that guided waves are only excited to a significant extent in a certain frequency range. Higher frequency ranges primarily excite bulk waves. In Fig. 4.7 (left), we show the gain of $\text{CH}_3 \rightarrow 5$ in the thickness resonance range around 1.2 MHz. The channel characteristics are much more favorable for communication, as the coherence bandwidth exceeds 45 kHz—more than 50 times larger than in the GW channels—showing that ISI is considerably reduced compared to GW channels.

Furthermore, Fig. 4.7 (right) shows the transducer’s impedance in the thickness resonance mode. The impedance is much more stable, with the gain varying by only 5% within this frequency band.



■ **Figure 4.7:** Gain and Impedance of an *SWP* channel primarily driven with bulk waves. Large bandwidth with little distortion and a more stable impedance allows higher data rates in the *SWP* channel.

4.3 Summary and Discussion

This chapter defined and explained the goals and challenges of this dissertation. We presented the test structure we evaluated our methods on and investigated the channel characteristics, i.e., the channel gains and measurements of the transducers' electrical impedance.

The transducers have a narrow resonance bandwidth determined by their geometry and material parameters. When attaching the transducers to a metal structure, their impedances are strongly coupled with the structure's frequency-selective acoustic impedance. Only low frequency ranges are appropriate to excite lamb waves in *GW* channels, and care must be taken that the transducer's resonance range matches this. The channel is frequency-dependent in this range, with deep spectral nulls and a narrow coherence bandwidth of only a few hundred Hertz. At the same time, the transducer impedance varies with frequency and can not be predicted accurately before transducer deployment.

In contrast, *SWP* channels can operate with bulk waves in higher frequency ranges, having a larger bandwidth. The *SWP* channel is much less frequency-dependent, and the transducer impedance changes only lightly with frequency. It is, therefore, much better suited for high-bitrate communication. However, the *SWP* channel can only be used over short distances because of the higher attenuation of bulk waves.

To put the results into perspective, we can compare the characteristics with typical RFID channels. Several sources report RMS delay spreads of at most $\tau_{\text{rms}} \approx 80$ ns (see Section 3.2.2). At the same time, RFID protocols specify data rates of less than 70 kbit s⁻¹. Consequently, the symbol duration is about 170 times larger than τ_{rms} . To achieve a similar relation between symbol duration and delay spread in the selected *SWP* channel, data rates would be limited to three bits per second. In the investigated *SWP* channel, the same relation holds at a data rate of roughly 250 bit s⁻¹. Therefore, ISI reduction mechanisms are more important in *GW* than *SWP* channels.

Increasing the Modulation Order


5.1 Overview

Section 4.2 reviewed the challenges of the acoustic channel, i. e., the strong reverberation due to multipath propagation and the bandwidth of the piezoelectric transducers. The achievable data rates of acoustic communication are severely limited when no countermeasures are implemented. At the same time, high data rates are favorable for several reasons. First, higher data rates enable novel applications such as live video transmission over the acoustic channel, proposed in various scenarios for through-metal applications [Hei+21; HS22]. Even if less demanding sensing applications are targeted, higher data rates can reduce active times of the MCU and reader, reducing the system's overall energy consumption. Furthermore, if multiple tags are employed, a shorter message duration increases the number of tags a reader can serve simultaneously.

Symbol rate and occupied bandwidth are inversely proportional. Hence, the narrow coherence bandwidth directly limits the maximum symbol rate in acoustic backscatter. One way to increase the data rates without requiring more bandwidth is to employ higher-order modulation, i. e., transmitting more than one bit per symbol. Higher-order modulation uses the available bandwidth more efficiently and increases the total data rates at the same symbol duration. In this way, the symbol duration can be chosen long enough so that ISI does not reduce the likelihood of successful decoding at the receiver. As discussed in Section 3.2.3, higher-order modulation significantly increased achievable data rates in underwater acoustic and electromagnetic backscatter. This chapter demonstrates a similar improvement for uplink communication over the through-metal backscatter channel.

5.1.1 Problem Analysis

Section 2.3.2 discussed the complex reflection coefficient, which is determined by the transducer's internal and load impedance. Conventional backscatter systems only use a binary modulation, i. e., they employ a transistor to switch between two states,



■ **Figure 5.1:** Load impedances are selected to yield a desired constellation for a specific transducer impedance. The constellation is distorted when the transducer impedance deviates from the expected value. In the received signal at the tag, self-interference poses an additional challenge.

$z_1 = 0$, and $z_1 \rightarrow \infty$. Instead, switching between different complex impedances to select z_1 enables the generation of an arbitrary amount of different constellation points for a higher-order modulation scheme. Hence, the backscattered signal can use any symbol constellation to modulate its data, paving the way for more bandwidth-efficient modulation schemes such as QAM.

Implementing such a modulator was straightforward in the prior studies of higher-order modulation, e.g., in underwater backscatter, where the characteristic impedance of the backscattering transducer z_0 is known. It is only weakly coupled to the specific channel and can be assumed as a constant property of the transducer or the antenna, respectively. We can then pick favorable constellation points from within the unit circle and compute the complex impedances z_1 that produce the selected reflection coefficient.

The acoustic GW channel differs from the underwater or general RF channel. We observed in Section 4.2 that the transducer's impedance z_0 is unknown before deployment due to its strong coupling with the structure's acoustic impedance. The acoustic impedance depends on the structure's geometry, material parameters, and exact placement of the transducers. Therefore, z_0 is unpredictable, and no fixed set of load impedances can be chosen to generate a desired constellation. Figure 5.1 visualizes this, where loads for an ideal 4-QAM constellation are picked for a specific transducer impedance. With a deviating transducer impedance, the resulting constellation is distorted. A further complication arises when inspecting the received signal at the reader, where the distorted constellation is additionally scaled, rotated, and shifted due to self-interference. The straightforward solution—measuring z_0 after deployment and manually adapting every tag's loads—is impractical, cumbersome, and vastly increases deployment costs. A more dynamic solution that adapts automatically to the specific channel is desirable.

As a second complication, we observed that the GW channel is highly frequency

dependent. A specific carrier frequency can work great in some channels, but communication may fail with the same carrier in other channels. Ultimately, chances are high that a carrier frequency with a much better SNR exists in any specific channel. A third difficulty is the time-variant nature of the channel. Once a device has adapted to the channel and chosen carrier frequency and load impedances, the channel may change due to varying environmental conditions, e.g., temperature and boundary conditions. Ideally, these estimates should be tracked continuously and adapted during communication.

5.1.2 General Approach

To achieve higher-order modulation despite the challenge of unknown transducer impedance, we envision a tag that dynamically adapts load impedances to the specific channel and the selected carrier frequency. Furthermore, we aim to let the reader dynamically select a suitable carrier frequency in the channel. To achieve this, we model the steady-state behavior of the backscatter channel mathematically. With the model, we analyze the potential benefits of dynamic load and frequency selection over conventional static approaches. For practical reasons, however, the dynamically selected constellations can not be perfectly symmetrical as in conventional communication systems. Therefore, we also investigate how channel coding can increase the robustness in the presence of interference and noise in such an unconventional constellation.


Since the resource-constrained tag can not detect the optimal configuration by itself, we derive a setup protocol to negotiate a suitable carrier frequency and the best load impedances for a desired constellation between a reader and a tag. The model parameters can be estimated at the reader with only three measurement points, enabling fast channel characterization. To assess the approach's practicability, we construct a prototype from commercially available off-the-shelf components and test our method in a study over multiple GW channels and an SWP channel. We further implement and evaluate a feature to track the channel's time-variant behavior and continuously adapt the reader's estimate of the constellation.

In this chapter, we only investigate the uplink communication. A functioning downlink, however, is required to perform the presented setup procedures, which we investigate in Chapter 8.

5.1.3 Structure

The chapter is structured as follows. We first describe the envisioned communication system and notation in Section 5.2, and then successively derive the channel model and the required setup protocol to estimate the model parameters on the fly. We then describe the implementation details of our tag and reader prototypes. In Section 5.4 we first describe the setup we used to test our prototypes and communication system. Then, we provide a detailed evaluation of the gain of dynamic selection of carrier frequency and tag loads over conventional backscatter systems. We then characterize the noise and

5 Increasing the Modulation Order



■ **Figure 5.2:** Overview of the communication system. As the backscattering tag works differently from a conventional transmitter, this system deviates from Fig. 2.1.

interference in the investigated channels and measure the maximum achievable data rates with higher modulation orders. Finally, we also show that the channel varies over time and that continuous adaption of the channel estimation improves the robustness of the communication system. Section 5.5 summarizes the results, discusses remaining challenges, and points to open research questions.

5.2 Communication System

This section provides an overview of the building blocks of the designed wireless communication system and clarifies the notation of signals at various steps in the transmission path. Figure 5.2 depicts the signal flow from an initial message at the tag to a decoded message at the reader.

The presented method is a packet-based communication scheme. The tag initially has a message $m[k]$ to transmit. Channel coding adds redundant bits to the message to allow the receiver to detect and correct potential transmission errors. The encoded bitstream is $m_c[l]$. The next step is symbol mapping, in which a group of $\log_2(S)$ bits is uniquely mapped to one of S symbols, resulting in a symbol stream $c[i]$. In contrast to conventional wireless communication described in Section 2.2, the transmitter does not generate a passband signal. Instead, it selects a load impedance at its tag for every symbol. We denote the stream of load impedances as $z_1(t)$.

The acoustic wave passes through the channel while the reader measures the transducer's current $i_R(t)$ and voltage $v_R(t)$. The demodulator recovers the complex impedance of the transducer $z_r(t)$, which is influenced by $z_1(t)$. The relation between load impedance at the channel input and observed impedance at the reader is characterized in more detail in Section 5.2.1. After demodulation, the reader computes probabilities for the individual symbols given the observed impedances. Finally, a soft decoder uses these probabilities and the channel code to reconstruct the received message $\hat{m}[k]$.

5.2.1 Steady-State Channel Model

The steady-state channel model describes the relation between channel input and channel output. The input is the selected load impedance at the tag's transducer, z_1 , and the output is the observed impedance at the reader's transducer, z_r . To avoid confusion, we want to point out that our understanding of the steady-state channel model in this chapter deviates from the common notion of a channel model in wireless communication. Typically, channel models describe the transient behavior of a channel, often characterized by a channel impulse response.

In contrast, our model describes the channel in steady-state, i. e., the channel output when the tag applies a constant load impedance for a sufficiently long time. In that case, all multipath effects have converged, and the channel output is also constant. The main goal of this channel model is to assist the reader and tag in negotiating a favorable set of load impedances for a specific structure. We are, therefore, interested in deriving a mapping $\mathcal{C} : z_1 \rightarrow z_r$.

We want to clarify at this point that we do not try to model the channel physically. Section 4.2 showed that multiple components contribute to the channel output, e.g., the transducer's transfer function, the mechanic waves propagation through the metal structure, including reflections, attenuation, dispersion, and complex interference patterns. Moreover, binding layers, such as epoxy resin, between the transducers and the structure also introduce dampings and phase shifts. A physical model that considers all of these factors is doomed to fail.

However, the electromechanical system in steady-state can be represented by an equivalent circuit [Roel15]. Therefore, the system's mechanical components are represented as complex impedances. The *acoustic impedance* of such a component is analog to the electrical impedance, where the pressure acting on a mass corresponds to the voltage over an electric component, and the particle velocity corresponds to the electric current. Using such an equivalent circuit model has a key advantage: It allows electrical engineers to use established tools to analyze the behavior of electromechanical systems.

Figure 5.3 (left) shows our proposed equivalent circuit model. We employ the widely used Butterworth-Van Dyke equivalent circuit model to represent the piezoelectric transducers. The transducers at reader and tag are each represented by a resistance and a capacitance (R_T^e , C_T^e , and R_R^e , C_R^e). The wave propagation from reader to tag is represented by many parallel branches, where every branch i corresponds to a resonant mode of the acoustic channel. Every path is modeled with a resistor R_m^i , an inductor L_m^i and a capacitor C_m^i . In resonance, the reactances cancel each other, and only the resistance remains, representing the gain in this specific resonant mode. Finally, across the reader's transducer, an AC voltage source generates the carrier signal, and across the tag's transducer, a load impedance z_1 is placed.

The model still has an infinite number of parameters, as there are an infinite number of resonant modes on the structure. However, we do not suggest estimating the parameters




Figure 5.3: Butterworth Van-Dyke lumped element model for the through-metal channel. Adapted from [Tse+20] (left), and simplified for only a single frequency (right). The components corresponding to the reader (blue), channel (orange), and tag (purple).

for all components. For plug-and-play-like deployability, we have to be able to estimate the channel with our low-cost reader and tag hardware. In steady-state, only a single frequency component travels through the channel. In that case, we can combine multiple components to yield a single, frequency-dependent complex impedance. The simplified diagram is shown in Fig. 5.3 (right).

The reader must acquire the transducer impedance z_r . Therefore, the reader applies the known sinusoidal voltage v_R at the transducer's terminals and measures the resulting current i_R . Using the equivalent circuit, we can analytically derive the reader's impedance depending on the load z_1 at the tag as

$$z_r = \frac{v_R}{i_R} = \frac{z_1 \cdot (Z_r^e Z_m + Z_r^e Z_t^e) + Z_r^e Z_m Z_t^e}{z_1 \cdot (Z_r^e + Z_m + Z_t^e) + Z_t^e (Z_r^e + Z_m)}. \quad (5.1)$$

The resulting transform from z_1 to z_r is a Moebius transform, a rational function in the form

$$z_r = \frac{z_1 \cdot a + b}{z_1 \cdot c + d}, \quad (5.2)$$

where $a, b, c, d \in \mathbb{C}$. The Moebius transform is always bijective, i.e., it guarantees a unique z_r for every selected z_1 .

The parameters a, b, c and d can be derived analytically with only three known points $(z_r^i, z_1^i), i \in \{1, 2, 3\}$. For a specific carrier frequency, we can determine the parameters by observing the impedances z_r with three different loads. Knowledge about the transformation will allow us to select the best set of loads to approximate a desired constellation. The channel estimation procedure described in Section 5.2.2 uses that model to determine a set of load impedances.

5.2.2 Setup Protocol

The tag and reader have to calibrate before they can communicate. They first negotiate a suitable carrier frequency, and then, the reader picks the best set of impedances for the tag. Lastly, the reader has to *learn* the exact mapping between the set of symbols and the expected impedances. This section illustrates the described procedure with exemplary results from a real channel. The detailed hardware and experiment setup used to generate these examples is described later in Section 5.3.1 and Section 5.4.1.

Carrier Frequency Selection

The reader aims to find a frequency at which the modulation of the load z_1 produces a large variation of the observed impedance z_r . This property depends mainly on the channel gain between the reader and the tag. We discussed in Section 4.2 that the channel's frequency response itself depends on many factors and is very sensitive to small changes—e.g., slightly different placement of transducers. Hence, it is not practical to determine the frequency response with simulation, as some existing studies about acoustic communication attempted.

Instead, we chose an experimental approach for frequency determination. The tag switches through a set of load impedances while the reader performs a frequency sweep. For every frequency, the reader measures the transducer's impedance long enough to observe the channel's response to all the generated load states. The standard deviation of the collected samples for this frequency serves as a metric for the channel gain. Combined with a noise power estimate in the channel, the reader calculates the SNR for each frequency. Finally, the reader picks the frequency with the highest SNR.

The experimental determination may take a long time, depending on the frequency resolution and the number of different loads through which the tag cycles. Furthermore, each load must be applied sufficiently long for a steady state to emerge, e.g. several milliseconds. The process may take up to 30 s for a high frequency resolution of 10 Hz. However, the variability of the environment is usually very low, as transducers are fixed to the metal structure, and structures do not change frequently. Boundary conditions and temperature may slightly change the channel, but usually not within short time spans. Therefore, the reader does not have to repeat this estimation often. Instead, it could even avoid periodic reestimation completely and wait for triggers such as repeating communication failures or very low SNR. Figure 5.4 shows the result for an exemplary channel.

Choosing load impedances

In the next step, the reader determines the mapping between symbols and load impedances. Testing all possible load impedances is time and energy-consuming, as every symbol requires up to 10 ms for a steady state to emerge. Alternatively, we leverage that the mapping between the load impedance and the observed impedance

5 Increasing the Modulation Order




Figure 5.4: The result of the frequency calibration procedure in channel $\text{CH}_{3 \rightarrow 2}$. The SNR is computed for every frequency in the sweep, showing large variations between carriers.

is a Moebius transformation. The reader can determine the transform's parameters from three pilot symbols, reducing the required time for load selection to less than 200 ms.

The tag transmits a pilot sequence with three load impedances while the reader observes the resulting impedances. From these three observed points, the reader computes the transform and then estimates z_r for any possible z_1 . Figure 5.5 (left) shows the estimated transformations for a specific channel. The reader picks the loads

$$z_1^i, \quad i \in \{1, 2, \dots, S\}, \quad (5.3)$$

such that the constellation's minimum distance between any two constellation points is maximized, i.e.,

$$\max \left(\min_{i,j:i \neq j} \|z_r^i - z_r^j\| \right). \quad (5.4)$$

After determining the best set of loads, the reader informs the tag about the chosen set Z^L for each symbol via the downlink. Figure 5.5 (right) shows the selected loads based on the estimated transformation. For evaluation purposes, we measured the actual z_r^i for every selected z_1^i and compared it to the prediction from the channel model. The tag is limited in the load impedances it can choose from: The hardware, which we discuss in Section 5.3.1, only enables selecting from a set of reactances, and modifying a series resistance.

Channel Estimation

The estimate of the Moebius transform is accurate enough to pick well-distributed loads, but Fig. 5.5 (right) shows that they do not perfectly match the observed impedances. Furthermore, the exact values may vary over time due to changes in ambient conditions, such as temperature. Inaccurate estimates of the constellation points, however,




Figure 5.5: The estimated transformations are based on three measurements per reactance (left). The branches show the resulting z_r with constant reactance and varying resistance. In this example, we pick 16 loads such that we achieve the largest possible distance between adjacent constellation points (right), and compare the predictions from the channel model with actually observed reader impedances.

decrease communication reliability. Hence, the reader regularly repeats a channel estimation, updating its estimate of the constellation points.

In the channel estimation step, the tag sends a pilot message switching through the selected constellation, i.e., all z_r^i . Using this pilot, the reader can obtain the z_r^i very accurately. Because of the long delay spread in the metal channel, the pilot uses a long symbol duration of ~ 20 ms, allowing convergence to a steady state within each symbol period. Consequently, the pilot message duration in this step depends on the number of symbols. For a 16-QAM-like modulation scheme, a pilot message takes 320 ms.

5.2.3 Message Transmission

After performing the setup protocol, the tag has a suitable set of load impedances and the reader established the constellation points for each load. Reader and tag are now ready to exchange the actual payload. This section describes the packet structure and the message encoding performed at the tag. We then explain the receiver structure, including synchronization, a novel channel adaptation, and decoding.

Sending

To transmit a message, the tag encodes the payload bits \mathbf{m} with an LDPC code, adding redundant parity bits. The code rate R describes the ratio of payload bits to total transmitted bits. The encoded bitstream $m[k]$ is then converted to a stream of symbols. We use a grey code to minimize the expected number of bit errors that result from a wrong symbol decision. The grey code ensures that adjacent symbols, i.e., symbols with

5 Increasing the Modulation Order

neighboring constellation points, only differ in a single bit. In case of an error, e.g., due to a noisy channel, a symbol is most likely confused with its direct neighbor. Hence, the grey code minimizes the amount of bit errors resulting from wrong symbol decisions.

Finally, the tag prepends a frame synchronization preamble consisting of three symbols. The tag then modulates the backscattered signal by applying the corresponding loads for each symbol to its transducer. Each symbol lasts for a symbol period T_{in} .


Synchronization and Channel Adaption

The reader continuously samples the transducer impedance and correlates it with the expected preamble sequence. The preamble serves two purposes. First, the reader uses it for frame synchronization via cross-correlation. Second, the reader further utilizes the observed preamble to refine its channel estimates. Small changes in the channel occur due to temperature variations, changing boundary conditions, or similar environmental influences. These variations may significantly change the mapping $z_1 \rightarrow z_r$. Instead of retransmitting the channel estimation message, which is time-consuming, we leverage the observed preamble sequence to adapt to the new channel parameters.

Assume that the Moebius transform w described the channel at the time of the initial channel estimation. In the meantime, the channel has changed slightly and is now described by a different transform \tilde{w} . This implies that a third Moebius transform exists, which maps $w \rightarrow \tilde{w}$, and this transform can again be determined by three known points. Based on this observation, we designed the preamble to contain three distinct symbols, so that we can leverage the received preamble to estimate the mapping from the original w to the adapted transform \tilde{w} . With the resulting transform, we update our estimate of all symbol constellation points at the reader accordingly. The effects of this adaption are evaluated in Section 5.4.6.

Sampling and Decoding

After preamble detection, the reader samples the payload impedances. Following each symbol transition, the channel shows transient behavior due to multipath propagation before converging to a new steady state. We chose a sampling rate several-fold higher than the symbol rate. Once the cross-correlation with the preamble yielded the symbol boundaries, the ideal sampling times for each symbol are chosen at the end of each symbol period—when the channel has converged to a steady state. This yields the signal $\tilde{z}_r[i]$, which is sampled at symbol rate. As an example, Fig. 5.6 shows the samples in a constellation diagram when taken at 2 ms and 6 ms after symbol transition. The influence of ISI is obvious.



■ **Figure 5.6:** Samples and respective constellation points. Shorter symbol durations increase ISI, leading to stronger deviation from constellation points.

The belief propagation LDPC decoder requires soft information, i.e., the probabilities that a certain symbol was transmitted given the acquired samples. We represent these probabilities as log-likelihood ratios (LLRs) during decoding. The LLR [PF14] is defined as

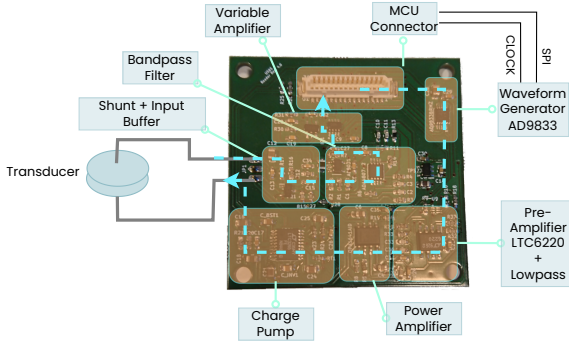
$$\text{LLR}(\tilde{z}_r[i]) = \log \frac{\mathbb{P}(\tilde{z}_r[i] \mid m_c[l] = 0)}{\mathbb{P}(\tilde{z}_r[i] \mid m_c[l] = 1)} - \log \frac{\mathbb{P}(m_c[l] = 0)}{\mathbb{P}(m_c[l] = 1)}. \quad (5.5)$$

Without prior knowledge of the transmitted message, the a priori probabilities of a bit being one or zero are assumed equally likely. Consequently, the second term resolves to zero. The LDPC code works bitwise, but every symbol represents multiple bits. The probability of an observed impedance given a single transmitted bit can be recovered by summing the probability of every symbol, where this bit is zero, or one, respectively. More formally, this is

$$\begin{aligned} \mathbb{P}(\tilde{z}_r[i] \mid m_c[l] = 0) &= \frac{\mathbb{P}(\tilde{z}_r[i], m_c[l] = 0)}{\mathbb{P}(m_c[l] = 0)} = \frac{\sum_k \mathbb{P}(\tilde{z}_r[i], c[i] = k)}{\mathbb{P}(m_c[l] = 0)} \\ &= \frac{\sum_k \mathbb{P}(\tilde{z}_r[i] \mid c[i] = k) \mathbb{P}(c[i] = k)}{\mathbb{P}(m_c[l] = 0)} \\ &= \frac{2}{S} \sum_k \mathbb{P}(\tilde{z}_r[i] \mid c[i] = k), \end{aligned} \quad (5.6)$$

where k loops over all symbols in which the l -th bit is zero. The last equality follows from the assumption of equally likely symbols, i.e., we assume $\mathbb{P}(c[i] = k) = 1/S$ and $\mathbb{P}(m_c = 0) = 1/2$. The probability $\mathbb{P}(\tilde{z}_r[i] \mid m_c[l] = 1)$ is constructed analogously.

5 Increasing the Modulation Order



■ **Figure 5.7:** The reader circuit performs the analog signal processing and carrier generation.

Assuming Gaussian noise with zero-mean and standard deviation σ ,

$$p(\tilde{z}_r[i] | c[i]) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\|\tilde{z}_r[i] - z_{R,k}\|^2}{2\sigma^2}\right), \quad (5.7)$$

where $z_{R,k}$ is the estimate of the constellation point related to the k -th symbol.

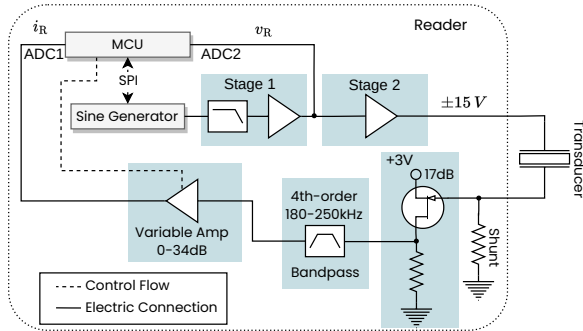
Using these probabilities, the reader implements a belief-propagation decoding algorithm using the LLRs for each bit. Finally, the redundant parity bits are removed, leaving only the decoded information bits.

5.3 Implementation

This section describes the prototype implementation, including the reader and tag hardware, and the implemented firmware.

5.3.1 Reader

We designed the reader as an extension to the *ahoi* modem [RHS20], an underwater acoustic modem with a modular design. By extending pre-existing hardware, we only need to redesign the analog circuitry, shown in Fig. 5.7, while we reuse the MCU and interfacing circuitry on the mainboard module of the *ahoi* modem. The mainboard features a performant 180 MHz STM32 Cortex-M4 microcontroller with 128 kB memory. The firmware on the MCU samples and demodulates the received samples.



■ **Figure 5.8:** Simplified schematic of the reader circuit.

Hardware

Our reader board serves as an analog frontend with two purposes. It generates a carrier wave—a sine wave with a selectable frequency—and it detects changes in the impedance of the reader’s transducer. A simplified schematic of the circuit is shown in Fig. 5.8. Sine generation leverages a dedicated waveform generator Integrated Circuit (IC), which has two advantages: It frees the MCU of additional workload for computing samples for digital-to-analog conversion, hence, all the computational power is available for reception and decoding. Moreover, it allows selecting carrier frequencies with a very high resolution. Using direct digital synthesis, the AD9833 from Analog Devices can generate sinusoidal waves from an input clock source. The output frequency f_c can thereby be configured as a fraction of the input clock frequency f_{mclk} by setting the value of the IC’s frequency register N_{freq} , such that

$$f_c = N_{\text{freq}} \cdot \frac{f_{\text{mclk}}}{2^{28}}, \quad N_{\text{freq}} \in \{1, 2, \dots, 2^{28}\}. \quad (5.8)$$

We configured the MCU to provide the sine generator with a $f_{\text{mclk}} = 5 \text{ MHz}$ input clock, yielding a resolution of the selectable carrier frequency of less than 20 mHz.

The waveform generator has a high-ohmic $0.3 V_{\text{pp}}$ sine wave output, which does not provide sufficient output power to drive the transducer. Therefore, this signal is passed through two operational amplifier stages, where the first amplifier also implements a lowpass filter to remove higher-order harmonics and noise that occur in the sine generator’s output. The first stage increases the amplitude to $3 V_{\text{pp}}$. This intermediate signal is also used by the MCU as a phase reference for demodulation. The second amplifier drives the transducer, and must therefore be able to supply sufficient power. We selected the OPA1641 operational amplifier and amplify the sine amplitude to $\pm 15 \text{ V}$. With the op-amp’s maximum output current of up to 30 mA, it can provide output powers

5 Increasing the Modulation Order

up to 350 mW RMS, driving transducers with input impedances down to 375 Ω . The high supply voltage of the second stage is provided by a charge pump.

To receive a backscattered signal from the tag, the reader measures the current through the transducer. We place a small 3.3 Ω shunt resistor in series with the transducer. A JFET transistor buffers the voltage drop over the shunt, which is proportional to the current through the transducer. The transistor is biased in its linear region and provides a voltage gain of approximately 17 dB. To increase the SNR, the measured signal is passed through a bandpass filter with passband between 180 – 250 kHz—the usable range of our transducer—and finally amplified with variable gain. The MCU can configure the gain in 255 steps in a range from 0 – 34 dB by setting a digital potentiometer. We used the 12-bit Analog to Digital Converter (ADC) integrated in the MCU to sample the current signal. At the same time, the MCU also samples the generated sine wave as a reference voltage.


The filters and amplifiers in the receiver signal path are supplied with 3 V. To avoid saturation at any stage, the current through the transducer must not exceed 65 mA. Consequently, the transducer's impedance magnitude must be larger than 230 Ω . However, the impedances of the selected transducers always exceeded this impedance in the measurements on the GW and SWP channels presented in Fig. 4.6.

Software-Demodulation

To determine the transducer's impedance, the MCU must extract the amplitude and phase offset of the current and voltage signals, i.e., demodulate them. While most conventional communication systems implement demodulation in the analog domain, the carrier frequencies and bandwidths in ultrasonic through-metal communication are small enough to perform the demodulation in software, reducing device complexity as less analog circuitry is needed. In this prototype, we employed the Short-Time Fourier Transform (STFT) as a digital demodulation technique to recover amplitude and phase of the carrier.

Our prototype is designed for carrier frequencies up to 250 kHz. According to the Nyquist frequency theorem, a sampling frequency higher than twice the maximum frequency component in the signal, i.e., more than 500 kHz is required. At this rate, real-time processing of the samples was not feasible with the selected microcontroller. Therefore, we employ undersampling, a technique to sample band-limited signals below the Nyquist rate, leveraging that aliasing projects the passband to a lower frequency band in the sampled signal. The effect of undersampling and STFT in the frequency spectrum is sketched in Fig. 5.9 (a).

The STFT suffers from frequency leakage when the sampling rate does not match a frequency bin perfectly. In that case, the phase and amplitude are distorted, and the demodulated signal has a carrier frequency offset. At the same time, the sampling rate and carrier frequency are all derived from the MCUs clock and are subject to limitations given by the MCU's design. As a consequence, only a limited set of carrier frequencies



■ **Figure 5.9:** Undersampling shifts the transmission band into an alias band below the Nyquist frequency (a). Only some carrier frequencies can be demodulated without error when using the STFT (b).

exists at which suitable sampling rates can be selected on the microcontroller. To derive the carrier frequencies that can be demodulated without frequency leakage, we formulate the set of constraints given by the hardware.

- The microcontroller's maximum peripheral clock frequency f_p is 90 MHz.
- The sampling rates f_s must be derived from the peripheral clock using an integer divider M

$$f_s = \frac{f_p}{M}, \quad M \in \{1, 2, \dots, 2^{16}\}. \quad (5.9)$$

The highest sampling rate that still allows in-time processing of the incoming signals is desirable, as higher sampling rates enable better noise suppression through digital filters and, consequently, a better SNR.

- The waveform generator can generate carrier frequencies

$$f_c = \frac{f_{\text{mclk}} \cdot N}{2^{28}}, \quad N \in \{2, 3, \dots, 2^{28}\}. \quad (5.10)$$

- The rate of the demodulated signal f_{sd} , i.e., the output of the STFT, must be larger than the bandwidth B of the transmission band. The output rate is the sampling rate divided by the window size W of the STFT:

$$f_{\text{sd}} = \frac{f_s}{W} \geq B, \quad W \in \mathbb{N}_+. \quad (5.11)$$

This set of constraints strongly limits the viable carrier frequencies the reader can generate and demodulate in software. Figure 5.9 (b) shows the set of carrier frequencies in the passband between 200 – 250 kHz when we assume a transmission bandwidth B of 2 kHz and a maximum sampling frequency of 200 kHz. Only 41 frequencies can be demodulated with the STFT, yielding an average frequency resolution of less than 1.2 kHz.

5 Increasing the Modulation Order

This limitation given by the STFT is not ideal, as we observed in Section 4.2.2, where peaks in the channel gain of GW channels occurred in bands smaller than 100 Hz. It should be noted here that it is possible to demodulate arbitrary carrier frequencies in software without distortion by mixing the sampled signal with a sinusoid of the same frequency followed by a lowpass filter. However, this is more computationally complex and was not feasible on the selected MCU.

The samples are processed by the STFT, yielding a demodulated signal. As we are only interested in the transmission band around the carrier, we only compute the STFT for the frequency bin corresponding to f'_c , the alias frequency of the carrier in the undersampled signal, which can be derived as

$$f'_c = \text{round} \left(\frac{f_c}{f_s} \right) \cdot f_s - f_c, \quad (5.12)$$

where `round` rounds the fraction to the next integer.

The reader demodulates both i_R and v_R , from which we calculate the complex impedance z_T .


5.3.2 Tag

In this prototype, we constructed a tag PCB that also extends the *ahoi* modem's mainboard. Although the powerful ARM Cortex-M4 microcontroller on the modem's mainboard is not a sensible choice for an ultra-low-power tag, it is suitable to acquire first insights into the challenges and opportunities of higher-order modulation. In this chapter, we favor the flexibility of a powerful MCU over minimal power consumption. Chapter 7 will investigate the additional challenges introduced with ultra-low-power hardware in more detail.

Hardware

The core of the tag's circuit is the modulator. It can choose between a set of impedances that it can apply to the tag's transducer to modulate the reflection coefficient. Both resistance and reactance can be modified to enable more flexibility and range in the choice of constellations. We use a digital potentiometer, allowing resistance modulation with a high resolution. Unfortunately, no comparable integrated circuits for reactance modification exist. Instead, we use switches that can connect a subset of different fixed reactance elements in parallel. The total reactance is then connected in series with the potentiometer (Fig. 5.10).

Suitable values for potentiometer and reactances must be adapted to the transducer. We aim to generate a wide range of different reflection coefficients at the tag, which in turn depends on the internal impedance of the transducer. While the exact transducer impedance is not known during hardware design, we can observe a rough region of impedance from Fig. 4.6. For example, in resonant frequencies, the observed transducer



■ **Figure 5.10:** The modulator circuit of the tag can choose resistance with a potentiometer and can modify the reactance by dynamically connecting different inductors or capacitors.

Device	R_{\max}	R_w	N_p	$\Delta\Gamma$	Γ_{\min}	Γ_{\max}
AD5175	10 k Ω	50 Ω	1024	0.032	-0.818	0.905
AD5293	20 k Ω	50 Ω	1024	0.062	-0.818	0.951
AD5259	5 k Ω	50 Ω	256	0.063	-0.818	0.951
AD5160	10 k Ω	50 Ω	256	0.121	-0.818	0.975
AD5280	50 k Ω	50 Ω	256	0.478	-0.818	0.995
AD5245	100 k Ω	50 Ω	256	0.757	-0.818	0.998

■ **Table 5.1:** Different available potentiometer ICs and the reflection coefficients they can generate in the average case of $z_0 = 500 \Omega$.

had an internal impedance magnitude of $z_0 \approx 500 \Omega$ with slightly capacitive reactance in the resonance region.


Commercially available digital potentiometer ICs have three important properties. The wiper resistance R_w describes the minimum resistance that the potentiometer can produce. A resistance R_n can then be chosen from a number of steps N_p between the wiper resistance and the maximum resistance R_{\max} as

$$R_n = R_w + n \cdot \frac{R_{\max}}{N_p}, \quad n \in \{0, 1, \dots, N_p - 1\}. \quad (5.13)$$

Table 5.1 compares the resulting maximum and minimum reflection coefficients (Γ_{\min} and Γ_{\max}), and the largest difference between neighboring reflection coefficients $\Delta\Gamma$ with the characteristic transducer impedance of $z_0 = 500 \Omega$ for different commercially available potentiometer ICs. Finally, we selected the AD5175 from Analog Devices because of its good tradeoff between resolution and range.

Similarly, the reactances are selected subject to typical transducer reactances observed in GW channels. Additionally, we also add a short-circuit path to the switch lines,

5 Increasing the Modulation Order



■ **Figure 5.11:** The tag designed in this chapter. The set of inductors is bulky and dominates tag size. The harvester is addressed in a later chapter.

which can be selected to produce a purely resistive load impedance. Finally, the tag can disconnect all switches to produce an open circuit where the load impedance approaches infinity. The resulting circuit board is shown in Fig. 5.11.

Firmware

The tag's firmware must select the right setting for the switch and potentiometer for every symbol in time. Therefore, the MCU configures a timer interrupt that is executed with interval T_m , and configures the switch and potentiometer via an I²C bus. The potentiometer and switch connect to the MCU over the same I²C bus. Consequently, the microcontroller executes a transition from one symbol to the next in two steps, e.g., first setting the new resistance value and then switching the reactance. Selecting a new state at the potentiometer requires in total 20 I²C-clock cycles, including address, start and stop conditions, and acknowledgments. Similarly, the switch requires 30 clock cycles. As most standard microcontrollers, the selected MCU allows I²C clock speeds up to 20 MHz, which finished a complete symbol transition within roughly 3 μ s.

At high clock speeds, the delay between resistance and reactance change is negligible, as it is orders of magnitude smaller than the symbol duration. However, such high I²C speeds also require the tag to use high clock rates, which increases the power consumption considerably. The relation between clock rates and power consumption is investigated in more detail in Chapter 7. When scaling down the clock rate to achieve ultra-low power consumption, the delays introduced by the serial transmission may become considerable, degrading the modulation accuracy of the backscatter communication.

5.4 Evaluation

In this section, we thoroughly evaluate the presented communication system. Therefore, we first present the experiment setup that we use for real-world experiments throughout this section. We then take a detailed look at the presented load selection algorithm and at the prototype's functionality. Finally, we explore the maximum achievable data rates over multiple channels and the influence of symbol duration and channel coding on the data rate.

5.4.1 Setup

In the following experiments, we used the tag and reader hardware described in Section 5.3.1 to generate the physical signals. However, part of the receiver's signal-processing pipeline was run on a desktop computer to be more flexible in parameter choice. The computer was connected to the tag and reader via serial interfaces. The PC configured the vector of load impedances per symbol (z_1), while the reader generated the carrier and sampled and demodulated i_R and v_R , and then transmitted the impedance measurements (z_T) to the PC. In real applications, the reader MCU would process the received samples locally. The setup is sketched in Fig. 5.12.

We used the experiment setup explained in Section 4.2.2 and kept the same convention to label the individual channels. In total, we used one SWP channel ($CH_{3 \rightarrow 5}$) and three GW channels with transducers 1 m ($CH_{3 \rightarrow 2}$), 2 m ($CH_{3 \rightarrow 1}$) and 3 m ($CH_{4 \rightarrow 1}$) apart on the 3.5 m long steel beam. To check if the method's performance is similar also in other structures, we experimented with a further channel on an additional specimen with a different cross-section. Therefore, the last channel ($CH_{6 \rightarrow 7}$) is a GW channel on a smaller steel beam with dimensions 1 m \times 2 cm \times 1 cm.

5.4.2 Dynamic Frequency Selection

Conventional backscatter systems usually employ a fixed carrier frequency. We want to quantify the gain of dynamically adapting the carrier frequency in GW channels. Figure 5.13 shows the SNR versus the frequency acquired for three channels on the same specimen. The measurements were repeated ten times in every channel and the average is shown. The standard deviation between the SNRs in the ten repetitions was roughly 3 dB in all channels. The experiment shows that the SNR varies by more than 40 dB within just 1 kHz carrier frequency change. Further, no frequency yields the optimal SNR in all three channels.

Among the three channels, 228 kHz is the frequency with the highest average SNR. This SNR would be acceptable for running a fixed carrier communication system for this specific set of channels. However, for an arbitrary set of channels, the best average frequency is not known before deployment, and there is no guarantee that it is also suitable in slightly different channels. Moreover, even in this small set of three channels,

5 Increasing the Modulation Order

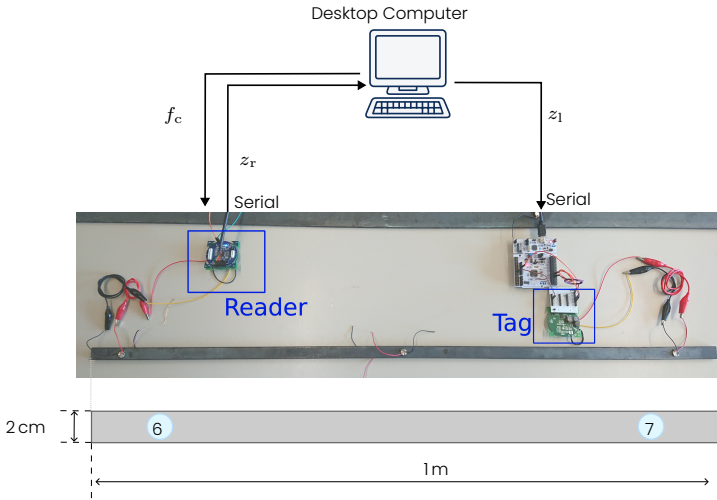


Figure 5.12: Setup with tag and reader. Their transducers are attached to a steel specimen with a length of 1 m. To conduct experiments more flexible, the desktop computer transmits the load impedances to the tag and processes the observed impedances from the reader.




Figure 5.13: Signal to Noise ratio versus carrier frequency in different channels. Picking an individual carrier frequency for each channel increases SNR.

every channel can achieve a better SNR when using a frequency dynamically adapted to the specific channel. Most prominently, the channel $CH_{3 \rightarrow 2}$ can gain additional 15 dB SNR when switching to its optimal frequency. Channels $CH_{3 \rightarrow 1}$ and $CH_{4 \rightarrow 1}$ improve only slightly by 2 dB, or 3 dB respectively. Moreover, an unlucky pick in the carrier frequency, e.g., 227 kHz would perform 70 dB worse than the optimum in $CH_{3 \rightarrow 2}$, and 48 dB worse in $CH_{3 \rightarrow 1}$. Note that the three channels were recorded on the same specimen with the same transducer type. The only difference between them is the transducer placement. Hence, we expect even more frequency dependency on varying structures. Therefore, dynamic frequency selection increases communication range and reliability for plug-and-play-like deployed readers and tags.

5.4.3 Dynamic Load Picking

In this section, we ask if dynamic load picking improves the range of possible constellation points compared to fixed load selections, as commonly used related work on higher-order backscatter modulation. Dynamic loads may require fewer components to generate high modulation orders, especially saving bulky inductors: With dynamic load picking, inductors can be reused for several constellation points or combined in parallel branches by activating multiple switches simultaneously. In contrast, with fixed loads, each load requires a dedicated fixed inductor. Furthermore, even when fixed loads are considered, optimal constellations may not be feasible in practical applications because they may require unrealistically dimensioned components, i.e., inductors in the range of tens to hundreds of millihenries.

We define the *modulation range* as the set of all z_r that can be generated with arbitrary loads at the tag. Based on the Moebius transform channel model, the modulation range can be represented as a circle, where $\text{Re}\{z_r\} = 0$. In our presented tag, the minimum load resistance is constrained by the potentiometer's wiper resistance—at least 50 Ω , reducing the modulation range to a smaller circle (See Fig. 5.14, left). To quantify the effect of the wiper resistance, we measured the achievable modulation range with our prototype (constrained). We then evaluated the unconstrained range with zero resistance by short-circuiting the potentiometer (unconstrained). The potentiometer reduced the modulation range (radius of the circle) of the impedances by 7% on average over all channels and carrier frequencies under test. The minimum radius reduction was 3.4% and the maximum 12%.

Section 5.2.2 discussed that the minimum distance between any two constellation points is important for the performance of the communication system. The presented tag is constrained in the choice of constellations it can generate. We now investigate how much this constraint reduces the distance between constellation points in real channels compared to ideal 2-PSK, 4-PSK, and 8-PSK constellations. Figure 5.14 (right) compares the minimal distances of ideal constellations versus the constrained minimal distances for all channels. In some channels, the distances are nearly halved, while in others, especially with just two symbols, the constraint has nearly no effect.

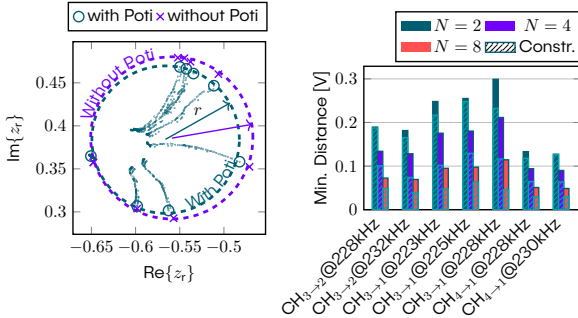



Figure 5.14: The maximum modulation range is constrained by the circle defined by $\text{Re}\{z_1\} = 0$. Every z_r within this circle can be achieved with a suitable load, exemplarily shown for channel $\text{CH}_{3 \rightarrow 2}$ (left). Using the constrained set of loads at the tag reduces the minimum distance between constellation points compared to the full, unconstrained range (right).

Unsurprisingly, ideal unconstrained loads perform better if optimally chosen for a specific channel. However, they also require manual adaption of the tag hardware for every channel. To decide if dynamic loads are favorable over fixed loads in real-world applications, we now investigate if a single set of fixed, arbitrary loads fitted to one channel outperforms the constrained loads in all considered channels so that the tag can be used in various channels without manual adaption. In Figure 5.15, we fitted a Moebius transformation to match the observed z_r and calculated the set of load impedances to yield optimal constellation points. We then compared the minimum distance with fixed loads (picked optimal for the first channel) and with the constrained, dynamic loads from our tag prototype. The plot shows the normalized reduction of the minimum distances relative to ideal loads for the particular channel. We observe that the constrained loads perform similarly to the fixed loads in most channels. Channel $\text{CH}_{3 \rightarrow 1}$ at 228 kHz is an exception, where the fixed loads perform particularly badly, and the dynamic loads from our tag perform significantly better, achieving a 35% increase in minimum distance when picking two symbols, and a 67% increase in distance at four symbols. The presented results suggest that the channels significantly influence the optimal load impedances, even if these channels are on the same structure, using the same type of transducer, and work within the same frequency range. The measurements do not clearly show that dynamic loads are superior to fixed loads in every channel. However, they show that dynamic load picking performs comparatively to fixed loads while reducing the count of bulky components.



■ **Figure 5.15:** A set of loads picked optimally for one channel ($\text{CH}_{3 \rightarrow 2}$ @ 228kHz) performs significantly worse on other channels.

Channel Model

The channel model presented in Section 5.2.1 predicts the transformation from load impedance to observed samples. In this evaluation, we ask how accurately the model represents the actual channel. Therefore, we first perform the load-picking algorithm presented in Section 5.2.2, which determines the transformation from three measured points. We then compare the predicted constellation points with actual measurements, as we have previously illustrated in Fig. 5.5 (right). On average, in all evaluated channels, the deviations between predicted and measured z_r is just 3.9 mV with little variance between the channels (maximum 6.1 mV, minimum 2.7 mV). This is merely 4.0% of the minimum distance between constellation points in a 4-QAM constellation. However, the prediction error becomes more significant when the modulation order increases. For example, when picking 16 symbols, the relative error rises to 34% of the minimum distance of constellation points on average.

We suspect that these deviations between prediction and measurement are caused by

1. non-ideal behavior of components, such as inductors and capacitors and ICs that have non-zero input capacities and stray inductances,
2. noisy measurement of the base points from which the Moebius transform is calculated, and
3. simplifications in the equivalent circuit used as a basis for the prediction.

5.4.4 Transient and Measurement Noise

The received signal is compromised by noise and distortion. Noise is usually assumed to be additive Gaussian noise, e.g., introduced by the analog circuits during the ADC conversion. Furthermore, ISI introduced by multipath propagation distorts the transmitted signal. We compare the magnitude of both in Fig. 5.16 for a GW channel and an SWP

5 Increasing the Modulation Order




Figure 5.16: Additive white gaussian noise, ISI, and SINR development depending on symbol duration in a GW channel ($CH_{3 \rightarrow 2}$, $f_c \approx 200$ kHz), and in an SWP channel ($CH_{5 \rightarrow 3}$, $f_c \approx 1$ MHz). The magnitude of ISI shrinks exponentially with symbol time.

channel. We measured the thermal noise during a steady state when the tag does not modulate its load impedance, and multipath effects vanish. The ISI curve was acquired during message transmissions. It shows the deviation of the received signal from the expected constellation points, averaged over 200 messages. We repeated the experiment with symbol durations from 1 – 8 ms. The ratio of signal power and the observed deviation is the Signal to Interference and Noise ratio (SINR), which is shown in the orange curve.

The measurements show that the ISI dominates the acquired signal over Gaussian noise. ISI intensity shrinks exponentially when longer symbol durations are chosen. However, as already observed in Section 4.2, the SWP and GW channels display very different behavior, as ISI in the GW channel drops below the thermal noise level only for symbol duration over 5 ms, while ISI in the SWP channel already decayed after 1 ms. The results highlight that the main limitation for higher data rates is ISI. Furthermore, we confirm that ISI is much more limiting in GW than in SWP channels.

5.4.5 Transmission Speed and Reliability

This section explores the achievable data rates of the proposed communication scheme. For this evaluation, we transmit pseudo-random messages from tag to reader and calculate the bit error rate (BER) and packet error rate (PER) with different parameters. For all channels and all modulation orders, we used the automated setup procedure described in Section 5.2.2 to adapt the carrier frequency and load impedances without manual tweaking. For every setup, 115200 payload bits were transmitted (200 messages with 576 bits per message).

We evaluate the effective gain of higher-order modulation by comparing the highest achievable data rate with the presented modulation scheme against the achievable data rate with conventional binary modulation. For the following evaluation, we define a




Figure 5.17: Comparison of the highest achievable data rates with higher-order modulation ($S > 2$) compared to simple binary modulation ($S = 2$) in different channels. Higher-order modulation only increased data rates in combination with channel coding.

data rate as achievable, if it supports a packet error rate less than or equal to one percent during our experiment. We chose this criterion because the comparatively low data rates in through-metal communication practically prohibit transmitting large numbers of bits, which is required to measure small BERs accurately. We observed, however, that the BER during our experiment was either zero or very small for some configurations, or very large ($> 10^{-2}$) for other configurations. Figure 5.17 shows the largest achievable data rates in all channels with different modulation orders.

Binary modulation ($S = 2$) is in none of the evaluated cases the best choice. However, a higher modulation order does not always improve data rates. The gain of using higher-order modulation strongly depends on the channel: The SWP channel ($\text{CH}_{5 \rightarrow 3}$) achieves a 211% increase in data rate. In contrast, the GW channels only achieve between 13 and 66% increase when using higher-order modulation. In SWP channels, the quick decay of multipath components (cp. Fig. 5.16) enables short symbol durations with little ISI, even at high modulation orders, where the distance between constellation points is small. In GW channels, in contrast, reduced symbol time strongly increases ISI, which inhibits the accurate discrimination of close constellation points. Figure 5.18 offers a different view on the presented results by showing at which symbol times these maximum data rates were achieved.

Furthermore, we have measured the highest achievable data rates without channel coding, as shown in the striped areas of the bars in Fig. 5.17. For all measurements, we used an LDPC code from the IEEE 802.16 (WiMAX) standard with a high code rate of $R = 0.83$, retrieved from [Hel+19]. In binary modulation, the uncoded communication achieves, in some cases, even higher data rates than coded communication because the lack of redundancy increases the number of payload bits when no bit errors occur without coding. Higher-order modulation, however, has only been advantageous in our tests when combined with error-correcting channel coding.

The effect of distance between transducers is inconclusive among the GW channels. In all channels under test, error-free communication was possible. However, the largest

5 Increasing the Modulation Order




Figure 5.18: Because of varying delay spreads, different symbol durations and modulation orders were achieved in the channels under test. The figure shows only the shortest feasible symbol time for every channel and every number of symbols.




Figure 5.19: BERs over time with and without the refined channel estimation. For this experiment, 32 symbols were used, and the symbol duration was 5 ms.

distance (3 m in $\text{CH}_{4 \rightarrow 1}$) achieves only half the data rate compared to $\text{CH}_{3 \rightarrow 2}$ and $\text{CH}_{3 \rightarrow 1}$, but the 2 m channel ($\text{CH}_{3 \rightarrow 1}$) achieves similar performance as the 1 m channel ($\text{CH}_{3 \rightarrow 2}$). For other acoustic channels, e.g., underwater, geometry plays a dominant role for shorter distances [RHS20], which is likely also the case here. Furthermore, we expect the effect of larger distances in GW channels to be much smaller than in conventional RFID or underwater channels because of the small attenuation that lamb waves experience in metals. In other words, the multipath characteristics of a channel have a stronger influence on the data rate than the distance.

5.4.6 Refined Channel Estimation

In Section 5.2.3, we presented an approach to correct the estimated channel constellation points. This experiment evaluates whether channel variations degrade communication reliability over time. Furthermore, we evaluate if the presented approach can successfully adapt to a changing channel. Therefore, we performed a channel estimation once to establish the mapping between z_1 and z_r . We then continuously sent one message per second for 1000 s and counted the number of bit errors per message. Figure 5.19 shows the BER over time. The BERs of every ten consecutive messages are averaged in the plot to smooth out fluctuations.

After a fresh channel estimation, the static and adaptive approaches both work well, and the messages are decoded without errors. After approximately 300 seconds, however, the reader without refined channel estimation starts to experience increasing bit error rates and becomes unusable. At the same time, the preamble correction approach keeps the same low BER during the whole experiment duration of roughly 17 minutes. Hence, without using the preamble for correcting channel variations, the full channel estimation would have to be repeated more often, reducing the effective data rate of the system. Similar results have been reproduced in all tested channels.

This effect is not as prohibitive when using a lower modulation order, e.g., 4 or 8 symbols, because the distances between constellation points are larger, and more channel deviation is tolerable. However, the increasing BERs show that the channel changes over time. We observed the channel variation in a static lab environment without significant temperature swings or changing forces on the structure. Channel deviations will be much more pronounced in practical scenarios, where environmental conditions are more dynamic. A deeper investigation of such scenarios will follow in Section 6.5.6.

5.5 Summary and Discussion

This chapter presented a method to implement higher-order modulation with acoustic backscatter communication by switching between dedicated load impedances per symbol. The challenge compared to conventional RF or underwater acoustic backscatter is that the transducer's impedance is unpredictable, making it impossible to select suitable loads before deployment. We presented a reader and tag design that enables the reader to dynamically select an SNR-enhancing carrier frequency and load impedances, using the bandwidth more efficiently. To establish efficient frequency and load selection, we implemented a setup protocol relying on an equivalent circuit model of the acoustic and electric channels. We also utilized the same model to adapt to time-variant channels without re-running the setup protocol, saving the reader and tag device's bandwidth and active time. Furthermore, we implemented the design in a hardware prototype with commercially available off-the-shelf components. We evaluated our prototype in real-world experiments with four *GW* channels and one *SWP* channel.

The evaluation demonstrated that dynamic frequency selection improves the SNR by up to 15 dB compared to a fixed frequency for all channels. Furthermore, dynamic load selection proved superior to fixed loads on average over seven inspected channels. Finally, when transmitting packets, higher-order modulation raised the data rates between 13 – 66% in *GW* channels and up to 211% in the *SWP* channel in combination with channel coding. However, the presented method has several downsides.

- **Hardware complexity:** Compared to conventional binary modulation, the required hardware is more complex. Tags require multiple bulky inductors and switches, while the reader has to demodulate the passband signal in software. In contrast,

5 Increasing the Modulation Order

binary modulation requires only an envelope detector at the reader and a transistor at the tag.

- **Setup required:** reader and tag must establish a common frequency and suitable load impedances. After setup, this state must be stored in both devices to enable successful communication. In binary modulation, no state is required, which reduces overhead and eliminates the requirement for persistent memory in the tag.
- **Channel Coding:** We could only achieve significant gains over binary modulation in combination with strong channel coding. In the scope of this work, we applied LDPC codes, which are known to have good error-correcting properties. However, LDPC codes also require considerable block lengths and are complex to decode, which might be unsuitable for some applications. Other types of channel codes may be a better fit for the envisioned applications. Therefore, future work should provide a detailed analysis of tradeoffs between coding gain and performance for different coding schemes.

Unfortunately, directly comparing the achieved data rates with prior work is not informative. The existing studies were conducted under very heterogeneous circumstances. In particular, no investigation of acoustic backscatter communication in *GW* channels exists in the literature. Moreover, most research groups used custom hardware, which we could not access. Piezo transducers and specimen material and geometry were also different. However, our results in the *SWP* channel are well in line with the 2.7 kbit s^{-1} reported in [Cun+16a], although the piezo transducers used in our experiments had only about one-fifth of the size.

Section 5.4.4 showed that the tag produces a strong signal over the short distances under test (up to 3 m), reaching *SINRs* up to 60 dB. However, ISI is the limiting factor for higher data rates. Equalization is therefore highly promising to further increase data rates in *GW* channels. However, the nonlinear backscatter channel renders conventional equalizers infeasible. We will analyze the problem and explore potential solutions in detail in Chapter 6.

The ultra-low power requirement for communication is the biggest advantage of backscatter communication. Our tag prototype deviates from common backscatter tags using dedicated switch- and potentiometer ICs instead of just a transistor. This raises the question of how these ICs affect the power consumption of the tag. Our prototype was developed as a proof of concept and did not satisfy low-power requirements. Hence, we can not present sensible measurements of the power consumption of the method compared to conventional backscatter. However, we can provide a rough estimate of the power consumption based on the component's datasheets. The selected potentiometer consumes about $1 \mu\text{A}$, and the switch consumes on average another $1 \mu\text{A}$. Therefore, the power consumption introduced by the ICs is likely negligible.

Battling Inter-Symbol Interference

6.1 Overview

Section 5.4 demonstrated that bitrates in GW channels do not surpass 2 kbit s^{-1} —although SNRs are relatively high with up to 25 dB. These bitrates may already suffice for many SHM applications that track long-term changes to a structure and can tolerate long intervals between individual measurements. Nevertheless, higher data rates reduce active times of sensors and enable a lower overall power consumption.

The major roadblock that prevents faster communication to date is ISI. In conventional wireless communication, equalizers are a common approach to mitigate ISI, which largely contributed to the increased data rates that wireless links achieved over the past decades. Unfortunately, the established equalizers can not be used for backscatter communication in GW channels because the backscatter channel is nonlinear. In contrast, conventional equalizers rely on the assumption of a linear channel.

After investigating the steady-state channel model for GW channels in Chapter 5, this chapter closely examines the channel's transient behavior. We characterize the source of the nonlinearity and propose a mechanism to linearize the GW channel, enabling the use of conventional equalizers to mitigate ISI. While the resulting method is developed and tested in through-metal GW channels, the fundamental linearization method also applies to other domains, such as acoustically coupled biomedical implants and underwater backscatter communication. In these environments, channel nonlinearity is likely not as pronounced as in GW channels but still plays a role.

6.1.1 Problem Analysis

In conventional communication systems, we assume that the channel is a Linear Time-Invariant (LTI) system, in which case the channel impulse response $h_c(t)$ fully describes




Figure 6.1: In conventional communication systems, the channel is represented by the channel impulse response $h_c(t)$ with additive noise $n(t)$ (a). In GW channels, the representation is much more complex, as it includes several impulse responses and the nonlinear mapping from load impedance to reflection coefficient (b).

the channel characteristics (see Fig. 6.1 (a)). The channel output $y(t)$ is then described as the convolution of the channel input $x(t)$ with $h_c(t)$ as

$$y(t) = x(t) * h_c(t) + n(t) = \int_{-\infty}^{\infty} x(\tau) \cdot h_c(t - \tau) d\tau + n(t). \quad (6.1)$$

As $h_c(t)$ is usually unknown to the receiver, it must be estimated. Therefore, the sender transmits a training signal known to the receiver. With knowledge of $x(t)$ and $y(t)$, the receiver estimates the channel's impulse response.

The backscatter channel, however, has two nonlinear components, which render the representation of the channel through an impulse response impossible. The first source of nonlinearity, appearing in all backscatter systems, including RFID, is reader self-interference. The reader generates an unmodulated carrier. Some of the carrier's energy reaches the tag over the downlink channel ($h_{d1}(t)$), is modulated, and reflects back to the reader over the uplink channel ($h_{u1}(t)$). At the same time, some—usually larger—fraction of the carrier's energy is reflected from the static environment with impulse response $h_{si}(t)$, which the reader also picks up. The received signal is, therefore, the sum of the reflected signal from the carrier and the constant reflection from the boundaries. With this static offset, the received signal no longer linearly depends on the tag's reflection coefficient. However, if the reader can estimate the self-interference, this nonlinearity component can be removed by subtracting the estimated self-interference from the received signal.

The second source of nonlinearity arises because the tag can not precisely control its reflection coefficient but merely controls the load impedance at the transducer. The mapping between load impedance and reflection coefficient, which we described in Section 5.2.1, is highly nonlinear. To complicate matters, the mapping is unpredictable because it is subject to varying transducer impedances. Consequently, this raises the question, what we consider as input for the backscatter channel. When considering the

load impedances as channel input, we obtain a nonlinear relation between input and output that can not be represented with a channel impulse response. Alternatively, if we consider the reflection coefficients as the channel's input, the receiver does not know the channel input, which is required for channel estimation.

6.1.2 General Approach

This section focuses on the transient transfer function and aims to enable equalization in nonlinear GW channels. To achieve this goal, we first thoroughly analyze the channel and derive the equations governing the channel transfer function, which includes the mapping from load impedance to reflection coefficient. We then prove that—under certain assumptions—the described nonlinear channel is equivalent to a symbol mapping followed by a linear channel. Moreover, the reader can estimate the mapping by acquiring the steady states of each symbol. Consequently, the reader can subsequently estimate the transient behavior of the linear channel as in a conventional communication system.

Leveraging this mapping, we perform a real-world analysis of two common equalizer types, DF and MMSE, in the GW channel and provide a detailed analysis of the potentials and limitations. Challenges that emerge in this context are a strong dependency on symbol timing and slowly varying channels. We investigate methods for symbol timing recovery and propose a variation of standard adaptive equalization techniques tailored to the nonlinear channel. We evaluate the magnitude and effect of these channel variations in GW channel with changing environmental conditions, showing that adaptive equalization can track such channel variations. Furthermore, we evaluate the gain that channel coding can achieve on the equalized signal.

To progress towards real-world implementations of an acoustic backscatter system on low-power devices, we also improve the hardware prototype presented in Section 5.3.1. The enhancements include the design of an analog demodulator at the reader, which reduces the computational burden of software demodulation and high sampling rates while largely expanding the range of selectable carrier frequencies compared to the previous prototype.

6.1.3 Structure

The remainder of this chapter is structured as follows. First, Section 6.2 analyzes the nonlinear backscatter channel in detail, describes the linearization method, and delivers a formal proof of its correctness. Next, we describe the implementation of the communication system, focussing on the implementation and training procedure of the selected equalizers in Section 6.3, and the proposed equalizer adaption method. Then, Section 6.4 describes the changes to the hardware. The evaluation in Section 6.5 first investigates tradeoffs and determines suitable values for equalizer parameters, i.e., the number of weights and length of the training sequence, then compares the performance of symbol timing recovery methods and the gain provided by additional channel coding.

Finally, we investigate the improvement in achievable data rates with equalization and the equalizer's capability to adapt to a changing channel. Section 6.6 summarizes the results.

6.2 Channel Analysis

This section aims to derive an equivalent channel that can be equalized with linear techniques. That goal is achieved when the equivalent channel

- is linear, i. e., it can be described with a channel impulse response, and
- the channel input for the training sequence is available to the reader, or the reader must be able to estimate the channel input from the received signal.

We describe the channels in baseband, i. e., we model the signals after demodulation. The received baseband signal is denoted as $y(t) = a_i(t) \cdot e^{j \cdot \phi(t)}$ with instantaneous amplitude a_i and phase ϕ .


6.2.1 Channel Modeling

The propagation paths in the metal structure are fundamental to our derivation of the channel model. We identify three distinct groups of propagation paths depicted in Fig. 6.2 (left). The first group consists of paths that originate from the reader's transducer and reflect from the structure's boundaries such that they reflect back to the reader. This group forms the impulse response of the self-interference $h_{\text{si}}(t)$. The second group contains paths that lead from the reader's transducer to the tag's transducer, defining the uplink $h_{\text{ul}}(t)$ and downlink $h_{\text{dl}}(t)$ impulse response. The last group are paths that originate from the tag's transducer and also end at the tag's transducer, leading to rescattering, i. e., the backscattered signal from the tag is reflected back to the tag and is again backscattered, characterized by $h_{\text{rs}}(t)$. These observations lead to the block diagram depicted in Fig. 6.2 (right).

The rescattering path complicates the channel model significantly. However, it has most likely only a minor influence on the channel output because the backscattered signal is at least an order of magnitude weaker than the incident carrier. Hence, in the remainder of this chapter, we neglect the rescattering paths in our model.

Formalizing the block diagram (without the rescattering path) in an equation yields the channel output $y(t)$

$$y(t) = c(t) * h_{\text{si}}(t) + [c(t) * h_{\text{dl}}(t)] \cdot \Gamma(t) * h_{\text{ul}}(t) + n(t), \quad (6.2)$$



■ **Figure 6.2:** Three different groups of propagation paths are relevant for the backscatter channel (left): Paths contributing to self-interference (orange), down and uplink (cyan), and rescattering (purple). Rescattering introduces a feedback loop (right).

where $c(t)$ is the carrier generated by the reader. The complex reflection coefficient $\Gamma(t)$ is now time-dependent because the tag modifies it during message transmission. The transducer's characteristic impedance is assumed to be constant. Hence,

$$\Gamma(t) = \frac{z_1(t) - z_0}{z_1(t) + z_0}. \quad (6.3)$$

Finally, noise occurs during the signal's generation, transmission, and reception, which we model as additive zero-mean white Gaussian noise $n(t)$.

We can simplify Eq. (6.2) for the case that the reader generates a single-frequency unmodulated carrier, i.e., $c(t)$ is a constant in baseband. When any channel impulse response is convolved with a constant α , the convolution is equivalent to the multiplication of the integral over the impulse response, i.e.,

$$\alpha * h(t) = \alpha \cdot \int_{-\infty}^{\infty} h(t - \tau) d\tau = \alpha \cdot g, \quad (6.4)$$


where we denote this integral over the channel impulse response as gain g . In the following equations, g_{si} , g_{d1} , and g_{u1} denote the integrals over $h_{si}(t)$, $h_{d1}(t)$, and $h_{u1}(t)$. Using this in Eq. (6.2), we can simplify the channel function to

$$y(t) = c_0 \cdot g_{si} + c_0 \cdot g_{d1} \cdot \Gamma(t) * h_{u1}(t) + n(t), \quad (6.5)$$

with $c(t) = c_0$ representing the constant carrier amplitude and phase.

6.2.2 Equivalent Linear Channel

Instead of learning the reflection coefficients, we propose to equalize an equivalent linear channel with transmit symbols y_{ss} and impulse response $\hat{h}_c(t)$. As the transmit symbols of the equivalent channel are generally not known beforehand, we now derive them



■ **Figure 6.3:** The nonlinear backscatter channel is equivalent to a learnable mapping followed by a linear channel.

analytically. We start by equating our desired linear channel with the original channel function (neglecting additive noise for brevity), i.e.,

$$y_{ss}(t) * \hat{h}_c(t) = c_0 \cdot g_{si} + c_0 \cdot g_{d1} \cdot \Gamma(t) * h_{u1}(t). \quad (6.6)$$

As the offset $c_0 \cdot g_{si}$ is constant, we can factor it into the convolution as

$$y_{ss}(t) * \hat{h}_c(t) = \left(\frac{c_0 \cdot g_{si}}{g_{u1}} + c_0 \cdot g_{d1} \cdot \Gamma(t) \right) * h_{u1}(t), \quad (6.7)$$

which already yields a linear channel function (with additive noise). We now further factor out the division by the gain of the uplink channel, yielding

$$y_{ss}(t) * \hat{h}_c(t) = \underbrace{(c_0 \cdot g_{si} + c_0 \cdot g_{d1} \cdot g_{u1} \cdot \Gamma(t))}_{\text{Transmit Symbols}} * \underbrace{\frac{h_{u1}(t)}{g_{u1}}}_{\text{Equiv. Channel}}. \quad (6.8)$$

When inspecting the equation for the transmit symbols, we can observe that they are equal to the received signal in steady-state when the tag does apply a constant load impedance (hence also a constant reflection coefficient) at its transducer. More formally,

$$y_{ss}(t) = c_0 \cdot g_{si} + c_0 \cdot g_{d1} \cdot g_{u1} \cdot \Gamma(t) = y(t)|_{\Gamma(t)=\text{const}}. \quad (6.9)$$

This is advantageous, as the reader can learn the transmit symbols easily if the training sequence includes for every possible transmit symbol a part where the received signal converges toward a steady-state. Therefore, it is not necessary to know the reflection coefficient anymore. In conclusion, the nonlinear channel is split into two distinct parts: a mapping from original load impedances to steady-state values, followed by an equivalent linear channel (see Fig. 6.3). Both can be estimated individually from a suitable training sequence.

6.3 Communication System

After determining an equivalent linear channel, this section describes the architecture of a communication system that leverages the equivalent channel. We start by describing the training protocol and message structure and then go into detail about the equalizer implementation.

6.3.1 Protocol and Message Structure

The communication system we describe here extends the system described in Section 5.2, which requires estimating the ideal load impedances at the tag and channel estimation. In this chapter, however, channel estimation must serve an additional purpose. After estimating the channel's steady-state symbols, the equalizer must learn the optimal weights to invert the equivalent linear channel. Figure 6.4 shows an example of a training sequence generated by the tag and the received signal at the reader. The message starts with the steady-state estimation, where the reader acquires knowledge about the mapping from symbols to steady-state. To obtain the steady state, each symbol is kept constant for a sufficient duration to converge, e.g., 5 ms in this example. The reader then samples the value at the end of each symbol period and stores it as the steady state for this symbol. For the sake of simplicity, the figure only shows a binary modulation scheme here. The steady-state mapping must be extended to include additional symbols when using a higher-order modulation scheme. We denote the acquired steady-state values as $\mathcal{Y}_{ss} \in \mathbb{C}^S$. A predefined sequence for equalizer training is transmitted directly after the steady-state estimation. The training sequence is a pseudo-random bit sequence known to the reader.

The training sequence adds up to a considerable length. If prepended at every message, the overhead quickly obliterates the potential increases in data rate gained through equalization. Hence, we investigate the required length of the training sequence in Section 6.5.1. However, the through-metal channel is primarily static, i.e., the transducers do not move. Furthermore, potential channel variations, e.g., induced by temperature, are expected to happen slowly, which we investigate in Section 6.5.6. Therefore, the reader can store the estimated equalizer parameters for messages between training intervals and update the channel estimation only rarely. To extend this update period, Section 6.3.3 investigates adaptive equalization that tracks changes to the optimal equalizer parameters over time.

After initializing the equalizer parameters, the reader is ready to receive payload messages. In our implementation, we prepend every message with a 13-symbol Barker-Sequence as a preamble to detect the start of a message. Following the preamble, the tag transmits the data symbols by switching to a unique load impedance for every symbol with symbol frequency $1/T_m$.

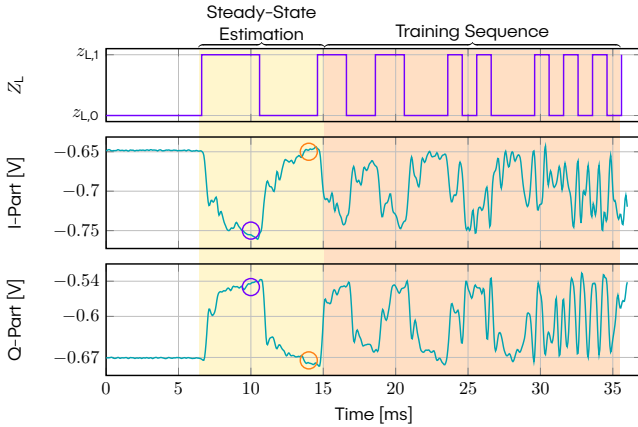


Figure 6.4: An exemplary training sequence for a binary modulation scheme. In higher-order modulation, the steady state estimation is extended to include every symbol.





Figure 6.5: Receiver structure showing the notation of signals at different stages in the analog and digital signal processing pipeline, their value range, and the sampling rates.

6.3.2 Receiver Structure

The receiver includes a hardware module that demodulates and filters the channel output in the analog domain, and a software-defined signal processing pipeline. Figure 6.5 provides a concise overview of the structure and the notation of signals at different stages in the pipeline. The demodulator removes the carrier and generates the base-band signal. Hence, the output is a complex signal where the real and imaginary parts represent the In-Phase and Quadrature components.

The receiver samples the signal at a multiple of the symbol rate N_m/T_m because timing recovery requires several samples per symbol. Consequently, the equalizer is fractionally spaced, i.e., it takes an input above the symbol rate. Finally, the equalizer outputs a single value for each symbol, which is the base for the final symbol decision.



■ **Figure 6.6:** Structure of the fractionally spaced MMSE equalizer, which consists of a forward filter only.

As in many modern receivers, the sampling clock is free-running, i.e., not synchronized with the transmitter's clock. Therefore, the samples may be acquired at any point within a symbol. This arbitrary symbol timing leads to a problem in the presented communication system because the equalizer trains once and then reuses the estimated weights in subsequent messages. If the symbol timing in a message deviates from the sample timing in the training message, the estimated equalizer taps are not optimal for this message, and equalizer performance degrades. We evaluate the effect of Symbol Timing Error (STE) on equalizer performance in Section 6.5.2.

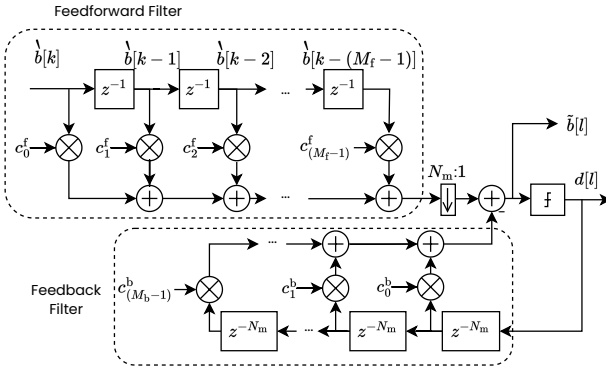
To mitigate the STE, we implement a square-timing-error detector that estimates the STE of the whole message in one step. The square-timing detector likely has two disadvantages in real-world systems. First, it requires the receiver to buffer the samples of the whole message before processing them, requiring significant memory and delaying reception. And second, the method is only applicable if the symbol timing offset does not change significantly throughout the message. However, for the evaluation in this chapter, these downsides are acceptable, as the receiving and transmitting devices are not ultra-low-power devices and have accurate clocks. Under these circumstances, the square-timing-error detector is more robust than dynamic tracking techniques, which we investigate for ultra-low-power communication devices in Chapter 7.

6.3.3 Equalizer

Next, the samples pass through an equalizer. In this work, we explore two common types of equalizers, MMSE and DF equalizers. Figure 6.6 shows the structure of the MMSE equalizer. It convolves the input signal with the learned coefficient vector $\mathbf{c} = [c_0, c_1, \dots, c_{M-1}]$. The required number of coefficients M depends on the maximum delay τ_{\max} of the impulse response, the symbol rate, and the number of samples per symbol. It can be estimated by

$$M \approx \left\lceil \frac{\tau_{\max} N_m}{T_m} \right\rceil. \quad (6.10)$$

Finally, the equalizer's output is downsampled to the symbol rate.



■ **Figure 6.7:** Structure of the DF equalizer, which consists of a feedforward and a feedback filter.

The DF equalizer, shown in Fig. 6.7, is more complex. It consists of two filters: a feedforward filter with M_f taps—similar to the MMSE equalizer—and a feedback filter with M_b taps. Symbol decisions feed into the feedback filter, which has two advantages. If correct, the symbol decisions are noise-free, so no noise enhancement occurs in the feedback filter. Furthermore, the feedback filter takes input at symbol rate. Hence, it requires fewer coefficients for large channel delays. We can calculate the number of coefficients as

$$M_f + M_b = M_f + \left\lceil \frac{\tau_{\max} - T_m}{T_m} \right\rceil, \quad (6.11)$$

where the feedforward typically includes only the samples of the current symbol, i.e., $M_f \approx N_m$. In the GW channels that we described in Section 4.2.2, we observed a maximum channel delay of roughly 8 ms. With four samples per symbol and a symbol rate of 1 ks^{-1} , we obtain 32 MMSE equalizer coefficients and only 11 DF equalizer coefficients.

Training

Deriving the optimal equalizer coefficients is a two-step process. First, the receiver must derive the input to the equivalent channel, denoted as $\mathbf{b} = [b_0, b_1, \dots, b_{(N_T-1)}]$, where N_T is the number of symbols in the training sequence. Derivation of the input signal is simply a mapping from the predefined training sequence symbol to the acquired steady-state values. In the second step, a linear equation system is solved to obtain the optimal coefficient vector \mathbf{c}_{opt} in the sense of minimal least square errors, i.e.,

$$\mathbf{c}_{\text{opt}} = \arg \min_{\mathbf{c}} \|\mathbf{A} \cdot \mathbf{c} - \mathbf{b}\|_2. \quad (6.12)$$

The matrix \mathbf{A} and the vector \mathbf{c} are compiled differently for the MMSE and DF equalizers. For MMSE, the i -th row of $\mathbf{A}^{\text{MMSE}} \in \mathbb{C}^{N_T, M}$ is given by

$$\mathbf{A}_i^{\text{MMSE}} = \begin{pmatrix} \dot{b}_{k+iN_m-(M-1)} & \dot{b}_{k+iN_m-(M-2)} & \cdots & \dot{b}_{k+iN_m} \end{pmatrix}, \quad (6.13)$$

and $\mathbf{c}^{\text{MMSE}} = [c_0, c_1, \dots, c_{M-1}]$.

For the DF equalizer, the \mathbf{A}^{DF} combines the acquired samples (forward filter) and past transmitted symbols (feedback filter) in one equation system, yielding

$$\mathbf{A}^{\text{DF}} = \left(\mathbf{A}^f | \mathbf{A}^b \right), \quad \mathbf{c}^{\text{DF}} = (\mathbf{c}^f | \mathbf{c}^b)^T, \quad (6.14)$$

with i -th rows

$$\begin{aligned} \mathbf{A}_i^f &= \begin{pmatrix} \dot{b}_{k-(M_f-1)+iN_m} & \dot{b}_{k-(M_f-2)+iN_m} & \cdots & \dot{b}_{k+iN_m} \end{pmatrix} \\ \mathbf{A}_i^b &= \begin{pmatrix} b_{i+M_b-1} & b_{i+M_b-2} & \cdots & b_i \end{pmatrix}. \end{aligned} \quad (6.15)$$

Adaption

We discussed that equalizer training is repeated only infrequently to reduce the overhead from long training sequences. Within the training intervals, the channel may change, and the estimated equalizer parameters become outdated. Adaptive equalizers continuously update the parameters during the reception of payload messages. The Least-Mean-Square (LMS) algorithm is an established technique for coefficient adaption, where the equalizer estimates its error e after every symbol, and slightly adapts the coefficient vector \mathbf{c} in the direction that reduces the estimated error most.

In contrast to conventional communication, however, a variation of the physical channel not only causes a change in the optimal impulse response, i.e., the coefficients, but also changes the steady-state symbols. Incorrect steady-state symbols increase equalizer errors because the equivalent channel becomes increasingly nonlinear. Therefore, we extend the LMS algorithm to modify the steady-state estimation after every received symbol. We refer to this technique as Constellation Point Adaption (CPA).

The core idea is to split the parameter update into two steps. In the first step, we assume \mathcal{Y}_{ss} to be correct and perform the equalizer coefficient update according to the conventional LMS algorithm. In the second step, we assume that the equalizer coefficients perfectly represent the channel impulse response and that the equalizer outputs are distributed around the steady-state symbols with zero means. Therefore, during the second step, we slightly adapt the estimated steady-state symbol of the currently detected symbol in the direction of the error.




Figure 6.8: Schematic of a decision feedback equalizer with error term.

For brevity, we illustrate the error term computation only for a DF equalizer in Fig. 6.8; the same principle also applies to the MMSE equalizer. For every $\tilde{b}[l]$, the decision device selects the i -th steady-state symbol y_{ss}^i with the closest distance, i.e.,

$$d[l] = \arg \min_{y_{ss} \in \mathcal{Y}_{ss}} \|y_{ss} - \tilde{b}[l]\|. \quad (6.16)$$

After every detected symbol, the detector updates steady-state estimation to reduce the error. At the same time, the equalizer coefficients \mathbf{c}_f and \mathbf{c}_b are updated according to the LMS algorithm.

$$\begin{aligned} e &\leftarrow d[l] - \tilde{b}[l] \\ \mathbf{c}_f &\leftarrow \mathbf{c}_f + 2\mu_c \cdot e \cdot \tilde{\mathbf{b}}^*[k, \dots, k - (M_f - 1)] \\ \mathbf{c}_b &\leftarrow \mathbf{c}_b - 2\mu_c \cdot e \cdot d^*[l - 1, \dots, l - M_b] \\ y_{ss} &\leftarrow y_{ss} - \mu_z e. \end{aligned} \quad (6.17)$$

Scaling factors μ_c and μ_z are required to avoid large parameter fluctuations caused by noise in the error signal. The choice of the scaling factor is important: If chosen too small, the adaption is slow and can not track changes fast enough. If chosen too large, however, the equalizers become unstable.

6.4 Prototype

The prototype we use to evaluate the presented communication system is largely identical to the one previously presented in Section 5.3.1. This section describes the changes we made to improve on the shortcomings of the previous prototype.

6.4.1 Analog Demodulator

The previous reader prototype demodulated the incoming signal in software by sampling transducer voltage and current in the passband and calculating the STFT, which required high sampling rates and considerable computational power. Moreover, the STFT relied on certain relations between carrier frequency, sampling frequency, and symbol rate,




Figure 6.9: Principle sketch of the implemented analog demodulator. The operations were implemented with operational amplifiers, while the mixing occurs with an analog switch.

strongly limiting the selectable carrier frequencies (see Section 5.3.1). Therefore, we implemented an analog demodulator in the prototype.

Analog demodulators usually mix the received signal with a local oscillator with the same frequency as the carrier. Typically, due to manufacturing tolerances, temperature differences, and similar effects, the carrier generated at the transmitter deviates slightly from the local oscillator, leading to CFO. Monostatic backscatter readers, in contrast, generate the carrier and demodulate the received signal in one device, which enables us to avoid carrier frequency offsets entirely. Unfortunately, we had to construct a custom demodulator from operational amplifiers because no suitable off-the-shelf analog demodulator ICs were commercially available. Available ICs were not specified for low-frequency ranges, as they are designed for RF applications, working in the range of hundreds of MHz to several GHz.

Figure 6.9 (left) shows how the demodulator is integrated into the reader, while the figure on the right sketches its working principle. The core idea is to mix the incoming signal $i(t)$ with a square wave (range between $+1$ and -1), where the square wave is generated from the same sinusoid that also generates the carrier $v_R(t)$. First, to reduce noise, the voltage $v_R(t)$ applied to the transducer, and $i_R(t)$, which is proportional to the current through the transducer, are filtered with fourth-order Butterworth bandpass filters with a passband from 200 – 250 kHz. For mixing, we generate a negated copy of $i_R(t)$ (equivalent to multiplication with -1) using an inverting amplifier circuit and then switch between the original $i_R(t)$ and its negation using an analog switch. The switch requires a digital input signal, which we derive from $v_R(t)$ using a comparator. Finally, as we are interested both in the in-phase and quadrature components ($\hat{b}_I[k]$ and $\hat{b}_Q[k]$) of the received signal, we add a second switch controlled by a 90° phase-shifted version of the square wave. The phase shift is achieved with an operational amplifier in a derivative configuration.

From Section 2.2.2, it is clear that the multiplication with a sinusoid at the same frequency as the carrier recovers the baseband signal. The same applies to multiplication

with a square wave, which consists of a sinusoid at the carrier frequency and infinitely many harmonics. These harmonics introduce copies of the baseband signal in higher frequency bands during mixing. However, these copies are effectively removed by the lowpass filters in the demodulator. After mixing, we pass the in-phase and quadrature signals through low-pass filters to remove the high-frequency components that also occur in the mixing process, yielding $\hat{b}_I(t)$ and $\hat{b}_Q(t)$, which are then sampled by the MCU. The I and Q parts are then combined to the signal $\hat{b}[k]$ and processed in software as described previously in Section 6.3.2.

6.5 Evaluation

This section evaluates the effectiveness and performance of the presented communication system and directly compares the two equalizer types, MMSE and DF, with each other. Our evaluation starts by identifying the required number of equalizer taps and training sequence length. We then explore the influence of the STE and its potential mitigation. Next, we compare the total achievable data rate with and without the applied equalization techniques, where we also account for potential gains through channel coding. We then demonstrate how environmental factors introduce time variance into the channel and quantify how well adaptive equalization tracks them.

The error vector magnitude (EVM) is an established metric for quantifying the accuracy of an equalizer and will be used extensively during the following evaluation. The EVM is defined as the average RMS equalizer error relative to the average distance between constellation points. In our case, this is

$$\text{EVM} = \sqrt{\frac{1/N_b \sum_{i=0}^{N_b-1} \|\tilde{b}[i] - d[i]\|^2}{1/S \sum_{i,j:i \neq j} \|\mathcal{Y}_{ss}^i - \mathcal{Y}_{ss}^j\|^2}}}. \quad (6.18)$$

To foster an intuitive understanding of the metric, we show a constellation diagram with randomly taken samples for three different EVMs in Fig. 6.10. The figure shows that for an EVM of 20%, the samples are still clearly distributed around the constellation points, and only a few errors may occur, if any. However, many errors occur within one message at an EVM of 35%.

6.5.1 Training Sequence and Parameters

We first evaluate the required number of filter taps in the equalizers, which directly translates to the delay spread in the channel that the equalizer can remove. We saw no improvement in extending the duration over 9 ms on the tested structure. Figure 6.11 shows the average EVM achieved with different equalizer durations once for a symbol duration of 0.25 ms and once for 0.5 ms. The equalizer errors have been averaged over 50 messages recorded with five samples per symbol. As a reference, we also show the




Figure 6.10: Illustration of sample distribution for different EVMs in a 4-QAM constellation diagram.





Figure 6.11: The chosen number of equalizer taps directly translates to the delay spread it can cover. Both equalizers were trained with a 340-symbol training sequence. The duration refers to the length of the equalizer's impulse response.

cumulative energy of the impulse response in the same channel. The 9 ms optimum equalizer length spans close to 100 % of the impulse response's energy.

The DF equalizer performed significantly better at shorter symbol duration, reaching a minimum EVM of roughly 12 %, while the MMSE equalizer achieves only around 15 % EVM. However, in the longer symbol duration, both equalizer types reach the same minimum EVM. The much larger number of filter taps in the MMSE equalizer likely causes its poorer performance in the first case: We calculate the number of coefficients in this scenario using Eq. (6.10) and Eq. (6.11) with a delay spread of $\tau_{\text{TMS}} = 9 \text{ ms}$ and $N_m = 5$. When $T_m = 0.5 \text{ ms}$, the MMSE equalizer requires 90 coefficients, and the DF equalizer requires only 22. At symbol duration 0.25 ms, the difference becomes more pronounced with 180 vs. 40 coefficients for MMSE and DF, respectively. Both equalizers were trained with the same 340-symbol training sequence. Hence, the smaller number of coefficients can be learned more accurately, reducing noise.



■ **Figure 6.12:** The DF equalizer requires significantly fewer taps, so a shorter pilot sequence suffices for equalizer training.

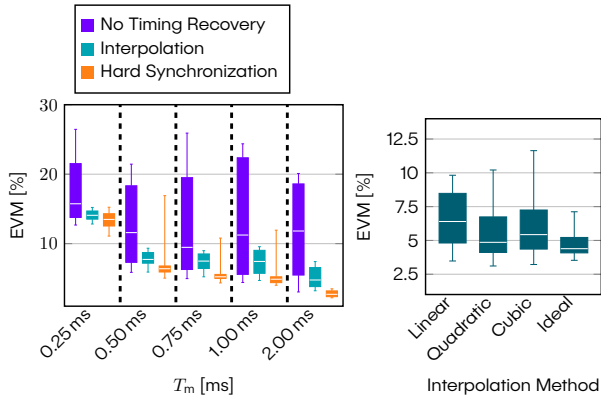
We used the same recorded messages to further investigate the relation between the training sequence length and equalizer error. In this experiment, we varied the training sequence length. After training, we again equalized all 50 messages and compared the average EVM. The results depicted in Fig. 6.12 show that MMSE and DF equalizers converge roughly towards the same EVM, but the MMSE equalizer requires a longer training sequence to achieve the same performance. With short symbol durations, this effect is very pronounced: For $T_m = 0.25$ ms, a training sequence of 170 symbols for the DF equalizer achieves an EVM of 18%, while the MMSE equalizer requires at least 400 symbols to achieve the same EVM. We can also observe the root cause of the MMSE equalizer's worse performance in Fig. 6.11: At a training sequence length of 340 symbols, it did not yet converge.

6.5.2 Symbol Timing Error

When the reader starts sampling the analog signal at a fixed rate, an arbitrary sampling delay τ occurs because of the lack of synchronization between reader and tag. When the training sequence is sampled with a specific delay τ_0 , the equalizer coefficients are optimized to remove the ISI at sampling instants with this specific delay. However, consecutive messages are likely sampled with a different $\tau \neq \tau_0$. In this case, the equalizer coefficients are not optimal for the channel, and the error at the equalizer's output increases.

We tested three cases to investigate the effect of the STE on our communication system and the developed synchronization approach.

- *No synchronization:* The samples are not corrected for symbol timing error. This case is the baseline with which we compare any mitigation strategy.
- *Timing recovery:* A square timing error detector was used to estimate the timing error in the sampled message. An interpolator then reconstructed the ideal sampling instants.



■ **Figure 6.13:** EVM in channel $CH_{3 \rightarrow 1}$ for 50 messages with different timing synchronization methods (left). The box shows the EVM range within which 90 % of the equalized messages fall, while the white line shows the median. Timing recovery reduces variations in output error between messages, although it can not reach the optimum (hard synchronization). When interpolating after timing error estimation (right), quadratic interpolation shows a good tradeoff between complexity and accuracy.

- **Hard Synchronization:** Reader and tag were connected with a cable, using a trigger signal to synchronize the sampling clock of the reader with the tag. As a cable defeats the purpose of wireless communication, we do not intend this method for real applications. However, it helps to determine the equalizer performance in the optimal case with perfectly synchronized signals.

Evaluation in a real channel shows that symbol timing errors significantly increase the equalizer's output error. Figure 6.13 (left) plots the variability of the EVM over 50 messages with different symbol times. Without symbol timing recovery, some messages achieve low EVMs (when $\tau \approx \tau_0$), while others show much higher errors, often causing packet errors. Timing recovery delivers more stable results, where the variability between messages is significantly reduced, and the average EVMs are very close to the best-detected EVMs in each run. In comparison, the hard synchronization produced even better and more consistent results, indicating that either the timing error detection or the interpolation diverged from the optimal results.

We further investigated four kinds of interpolators, i.e., linear, quadratic, cubic, and ideal interpolation. Ideal interpolation involves a fast Fourier transformation (FFT), multiplication in the frequency domain, and an inverse FFT. It is, therefore, very resource-consuming. We do not consider it a viable solution to run on the prototype, but it serves as a best-case reference. Differences in overall performance between the different interpolations are minor, as shown in Fig. 6.13 (right). As expected, the ideal interpolation performs best on




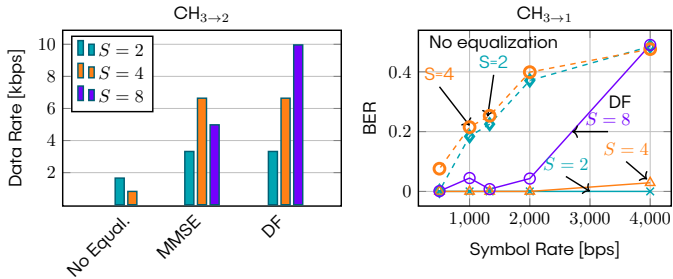
Figure 6.14: The equalized channel's packet error rate (PER) with and without decoding. Iterative belief propagation decoding was performed with a maximum of 10 iterations. Here, a symbol interval of $T_m = 0.25$ ms, and 4-QAM like modulation was used.

average. A quadratic interpolation scheme is a good tradeoff between computational complexity and EVM.

6.5.3 Noise and Channel Coding

To estimate the influence of equalization on the communication system, an isolated view at EVMs is insufficient. Instead, joint consideration of equalizers and channel coding may offer additional insights. For example, Goldsmith reports that channel coding can not improve upon the burst errors that appear through error propagation in DF equalizers [Gol05, p. 364f]. This limitation of DF equalizers raises the question of whether DF equalizers, in combination with channel coding, may have a worse overall outcome in terms of PERs despite yielding a slightly better EVM. To answer this question, we recorded 50 messages encoded with the LDPC code introduced in Section 5.4.5 (code rate $R = 0.83$). We then processed the messages once with a DF and an MMSE equalizer. We either performed the final symbol decision directly from the equalizer output (no coding) or after applying a belief propagation decoder with up to 10 iterations. To characterize the influence of SNR on the outcome, we repeated the experiment with simulated Additive White Gaussian Noise (AWGN) with increasing standard deviation.

The results in Fig. 6.14 indicate that the decision feedback equalizer still outperforms the MMSE equalizer in terms of noise sensitivity. Even without coding, the DF equalizer tolerated 9 dB lower SNR in the one-meter channel and about 3 dB in the two-meter long channel while producing the same PERs as the MMSE equalizer. Furthermore, channel coding increased the noise tolerance by further 3 – 4.5 dB in both channels. We also observe a significant coding gain of 9.5 dB and 5 dB with the MMSE equalizer, but the DF equalizer outperforms the MMSE equalizer in all cases.



■ **Figure 6.15:** Achievable error-free throughput with the different configurations. The data rates show the net throughput of information bits, not considering the redundant bits used for channel coding.

6.5.4 Achievable Data Rates

In the previous sections, we compared the performance of MMSE and DF equalizers and explored timing recovery. We now ask what benefit equalizers provide in terms of the achievable data rate of the complete communication system, i.e., data rates that support error-free communication in our experiments. In total, we recorded 100 messages for symbol durations $T_m \in \{0.25 \text{ ms}, 0.5 \text{ ms}, 0.75 \text{ ms}, 1 \text{ ms}, 2 \text{ ms}\}$, and for number of symbols $S = \{2, 4, 8\}$, i.e., binary, 4-QAM, and 8-QAM-like modulation.

The results show that without equalization, the highest achievable data rate is less than 2 kbit s^{-1} in the one-meter channel $\text{CH}_{3 \rightarrow 2}$, and higher-order modulation with eight constellation points can not be achieved at all even with a long symbol duration $T_m = 2 \text{ ms}$. The MMSE equalizer triples the achievable data rate, and the DF equalizer even reaches more than five times the data rate in this channel, establishing error-free communication with eight constellation points per symbol at 4000 symbols per second ($T_s = 0.25 \text{ ms}$). We achieved slightly lower data rates in the longer channels but could also observe significant speedups through equalization. In the two-meter channel $\text{CH}_{3 \rightarrow 1}$, the maximum achieved data rate was 6.6 kbit s^{-1} with a DF equalizer and only 830 bit s^{-1} without equalization, as shown in Figure 6.15. We corrected the STE in all configurations and used channel coding with an LDPC code.

6.5.5 Timing Imperfections

The previous section demonstrated that equalization using the equivalent channel enables higher data rates in the GW channel. However, experiments in a real-world scenario—especially with custom hardware—introduce numerous imperfections that may degrade the overall performance of the communication system. We test our receiver in an idealized simulated case to put the achieved rates into perspective.

When comparing the achieved data rates with the Shannon capacity limit, we find that a more than six-times speedup is theoretically possible in channel $\text{CH}_{3 \rightarrow 2}$. We computed the capacity limit using the acquired impulse response and the water-filling power allocation, yielding a maximum data rate of 66 kbit s^{-1} in channel $\text{CH}_{3 \rightarrow 2}$. Several causes come to mind to explain this shortcoming. The following effects may negatively impact the error rate of the real-world system.


- *Non-linear channel:* The presented channel model proved the equivalent channel's linearity, neglecting tag rescattering. However, rescattering might have tangible effects, in which case the channel linearity assumption is violated. Furthermore, the receiver hardware, i. e., filters and amplifiers, may also add nonlinearity to the channel.
- *Short Codes:* It is well known that approaching the channel capacity requires long channel codes. We only use relatively short codes, as the low data rates in the acoustic channel make very long messages impractical.
- *Clock Timing:* Sampling frequency offset between reader and tag distorts the signal nonlinearly, which the applied square timing error detector and interpolator can not correct.

To estimate the maximum possible performance without distortions from the hardware or the timing, we measured the channel impulse responses and simulated ideally linear received signals. The simulation convolves the transmitted signal with the acquired channel impulse responses. We then added noise to the received signals with a distribution similar to the recorded messages, i. e. gaussian noise passed through filters with a characteristic like the receiver hardware. This artificial generation guarantees that no non-linearities exist in the simulated signal.

Furthermore, we incorporated a random timing error into the simulation, i. e., the whole signal is resampled with a slightly different start time to mimic an unsynchronized reader and tag. The random offsets stem from a uniform distribution between zero and N_m . Lastly, we added an SFO, i. e., resampled the simulated signal with a slightly faster sampling rate.

Figure 6.16 shows the median EVM in the equalized signals under various conditions. Under ideal conditions, the EVM is only 58 % lower than in the real channel $\text{CH}_{3 \rightarrow 2}$, and 74 % lower in $\text{CH}_{3 \rightarrow 1}$. However, as soon as the timing error is randomized, the EVM rises significantly despite the timing error correction in the receiver. Adding SFO to the simulation further degrades the received signal. However, the effects of SFO are much more intense in $\text{CH}_{3 \rightarrow 1}$.

To put the simulated results in perspective, we also measured the clock difference between the MCUs on the reader and the tag, which was about $3.5 \cdot 10^{-4}$. With a similar SFO, the simulation yields very similar EVMs to the real transmissions. This indicates that the main contribution to the error is, in fact, the timing between the reader and the tags MCU.



■ **Figure 6.16:** Comparison of the EVM in the observed channels (real), a perfectly linear simulated channel (ideal) with and without sample frequency offsets (SFO), and a random timing error (STE).

6.5.6 Time-Variant Channels


This section characterizes the time-variability of the equivalent linear channel and explores the effectiveness of the adaptive equalization technique, including the CPA technique introduced in Section 6.3.3.

Static Evaluation

We first investigate adaption in a time-invariant channel, in which we assume that the channel parameters are not chosen optimally but do not change considerably throughout the experiment. In a real channel, we cannot change channel parameters in a controlled way. Instead, we deliberately add an error to the estimated steady-state values and the optimal equalizer coefficients and evaluate if the introduced adaption technique can remove the error successfully. We used a recording of 100 sequential messages in a time-invariant channel, in which we acquired the optimal steady-state values in a channel estimation step. We then distinguish three different scenarios:

1. *Ideal steady-state estimation:* The equalizer knows the correct steady-state values during training and message reception.
2. *Steady-state change after training:* The equalizer performed training using the correct steady-state values, but steady-state values changed after training, rendering the equalizer coefficients incorrect.
3. *Flawed channel estimation:* The steady-state values have been incorrectly estimated from the training sequence. Therefore, the equalizer has trained on a non-linear channel, yielding sub-optimal results.

In the latter two scenarios, we mimic a variation in the channel or an inaccurate channel estimation by adding an offset to the correct steady-state values, which the equalizer uses during training and reception.




■ **Figure 6.17:** Error with and without constellation point adaption (CPA) in different scenarios.

Figure 6.17 displays the EVM over time in the three scenarios. In the ideal scenario, CPA produces roughly the same average error as conventional LMS. However, the total EVM and the variation of the EVM between messages are slightly larger. This is expected, as the error used for adaption stems from noisy signals and also causes detrimental feedback. In the time-variant scenario, where the equalizer trains with the correct constellation points, the EVM is large initially because the trained equalizer coefficients no longer fit the actual channel. Both LMS and CPA can improve this error, but CPA strongly reduces the EVM from initially 40% to less than 13%, close to the ideal values at 10%. At the same time, without CPA, the LMS algorithm can adapt the channel parameters to reduce the error from 42% to only 21%. However, it can not fully adapt because the steady state constellation points are wrong, and the equalized channel is nonlinear.

Finally, in the last scenario, where the equalizer has already trained using an inaccurate estimate of the steady-state values, both LMS alone and LMS with CPA can slightly improve EVM. However, the CPA algorithm does so much quicker within the first few messages. However, none of the techniques achieves a similarly low error level as the ideal case, i. e., the adaption can not fully converge to the optimal coefficients and steady-state values. Hence, the results indicate that correct steady-state estimation from the training sequence is crucial for good equalization performance. However, CPA can be used to track time-variant constellation points. However, CPA will slightly degrade the EVM when there is little variation in the channel. This effect is already known from LMS adaption, which also enhances noise during equalization.

Real-World Evaluation

In a second experiment, we apply CPA in two real-world scenarios, showing how sensitive the metal channel is toward small environmental changes. In the first scenario, temperature variations in the metal cause the channel to change, while in the second scenario, we expose the specimen to external forces. We first performed channel estimation with




■ **Figure 6.18:** The setup for the dynamic channel experiments, with all used hardware.

equalizer training in both scenarios and then continuously transmitted messages over a 1-meter channel for several minutes. We then evaluated the same recorded samples without adaption, with LMS adaption only, and with LMS and CPA. Figure 6.18 depicts the setup with all equipment used for this experiment.

Figure 6.19 (left) shows the average EVM per message over time, while we heated the metal bar with a hot air gun from initially 24 °C up to 58 °C, measuring at a point in the middle between reader and tag. In outdoor environments, such temperature spans are feasible, e. g., when considering a bridge that heats up in the sun during the day and cools down overnight. However, in most realistic scenarios, we expect to observe more gradual temperature changes throughout the day.

Without adaption, the EVM quickly rises as we heat the metal. After about 35 s, the equalizer output is so distorted that subsequent correlation fails to detect a message. LMS improves the EVM and enables continuous reception during the entire experiment. Finally, with CPA, all messages were also successfully received, and the EVM was significantly lower than with only LMS during the whole experiment. We observed improvement, especially during the fast heat-up and early cool-down phases, where the channel changes much more rapidly than in the later cool-down phase.

We performed the same experiment in the second scenario, shown in Fig. 6.19 (right). However, instead of heating the metal bar, we applied mechanical forces. Forces included clamping the metal locally, putting a 5 kg weight on the metal bar, and stepping on it with body weight (≈ 80 kg) during message transmissions, causing it to bend and change geometry noticeably. In contrast to the first scenario, the channel changes much quicker. Without adaption, the EVM rises to over three times the initial EVM during transmission, causing high bit error rates. The LMS and CPA algorithms significantly reduced the EVM in response to the changed boundary conditions. Bit errors only occurred in a few messages directly after the change, before the equalizer could adapt. In this scenario, CPA did not add additional improvement over the conventional LMS algorithm, which indicates that the steady-state values do not change significantly in response to



■ **Figure 6.19:** With static equalizer coefficients, the EVM rises strongly upon temperature variations in the material. Adaptive equalization using LMS and CPA can reduce the EVM and mitigate communication failure due to channel variation. In the right plot, LMS and LMS+CPA achieve nearly identical outcome.

the mechanical forces.

6.6 Summary and Discussion

This chapter thoroughly analyzed the transient behavior of the GW backscatter channel and identified the combination of an unpredictable reflection coefficient and reader self-interference as the cause of nonlinearity. We derived an equivalent linear channel that only required knowledge of the steady-state values for each selected load impedance at the tag, which is easy to acquire in a training sequence. Based on the linearized channel model, we implemented two equalizer types, the linear MMSE and the nonlinear DF equalizer, and tested them on real channels. The results have shown that equalization with the linearized channel can significantly speed up communication rates to up to 10 kbit s^{-1} compared to only 2 kbit s^{-1} in the same channel without equalization. The DF equalizer has proven to be better suited between the two equalizer types, yielding lower error rates and requiring significantly shorter training sequences to achieve similar performance. However, for both equalizers, sample timing offset between messages played a significant role, and detection and interpolation were mandatory to achieve robust decoding. We also identified the **SFO** and **STE** due to unsynchronized devices to contribute most to errors at the equalizer output. Therefore, we will look into further clock synchronization methods in Chapter 7.

Requiring steady-state values complicated the implementation of adaptive equalization, as these steady state values must be tracked over time together with the equalizer coefficients. We proposed and implemented a suitable tracking technique and applied it to static errors of the channel estimation and dynamic cases with a time-variant channel. The evaluation demonstrated that environmental influences, such as temperature, stress, and strain, significantly influence the channel parameters. While conventional

LMS adaption successfully mitigated these effects, the proposed constellation point adaptation achieved lower equalizer output errors and, therefore, more robust adaptation. However, the experiments were under controlled laboratory conditions and might not fully replicate realistic conditions, which could also include structural vibrations. Follow-up research on real structures under different environmental influences must be conducted to clarify the practicability of the adaption technique.

This chapter also described refinements of the presented experimentation hardware that enabled the more fine-grained choice of the carrier frequency and a reduction of required sampling rates through the implementation of an analog demodulator. For this evaluation, most parts of the software-defined receiver were not yet implemented on the microcontroller but processed on a desktop computer. Among the receiver steps, the equalizer training is the most critical part, which requires significant computational power and memory. To verify the approach's feasibility, a master's student implemented the MMSE and DF equalizers on the reader MCU [Zor22]. The results from the implementation also highlight the advantages of the DF equalizer with its fewer weights that must be learned. Using the DF equalizer, the same training as in the presented offline receiver took about 25 ms on the MCU, and processing the samples during message reception on the MCU was fast enough to support a data rate of more than 200 thousand symbols per second with a 4-QAM like modulation and five samples per symbol. In comparison, under the same circumstances, the MMSE equalizer required over 220 ms for training and enabled only a maximum symbol rate of 130 thousand symbols per second. In both cases, the MCUs resources are more than enough to support equalization in a real-time system on an embedded device.

Wirelessly-Powered Sensor Nodes

7.1 Overview

The prototypes presented in the previous chapters relied on external power supply, e.g., with batteries or cable-bound. While backscatter communication alone is worthwhile to extend the lifetime of battery-powered devices, its full potential is only tapped when the nodes are batteryless. In applications in civil infrastructure with multi-decade lifetimes and sensors in hard-to-reach and potentially inaccessible places, regular battery replacement is uneconomical. Like RFID, reusing the carrier wave for wireless power transfer is a promising approach to release these tags from the shackles of limited energy storage. However, the literature on acoustic through-metal power transfer in GW channels is unsatisfying. Often, the researchers only inspect a single channel on a small specimen and tune their setup to work optimally in this specific configuration. Moreover, many leverage high-power and high-cost laboratory equipment and neglect the practical constraints of cheap and ultra-low-power hardware required for economically attractive monitoring solutions.

In general, while SWP channels are well predictable and enable highly efficient power transfer, GW channels are poorly understood and often only provide low power-transfer efficiencies—particularly over several-meter distances. Hence, the operator of such a system must expect to invest large amounts of power at the transmitter for every unit of consumed power at the tag. Therefore, ultra-low-power demanding tags are beneficial for supporting the efficient operation of such sensor networks.

7.1.1 Problem Analysis

In principle, harvesting and communication techniques for wireless power transfer and backscatter communication are well-known. However, using guided acoustic rather than electromagnetic waves poses an essential implication for impedance matching, which plays a central role in any energy harvesting system. As presented in Section 4.2.2, a transducer's output impedance in a GW channel is strongly coupled with the structure's

acoustic impedance and unpredictable before deployment. Hence, manually designing an impedance-matching circuit suited for the harvester and source—a requirement to achieve efficient energy transmission—is not viable in GW channels.

Moreover, using the same transducer for harvesting and backscattering introduces mutual dependencies between the two functions. For example, the modulator shorts the transducer's terminals during backscattering, prohibiting energy extraction from the transducer during this time. Similarly, a well-matched harvester prohibits the modulator from applying a large impedance at the transducer's terminals. While this problem also occurs in RFID systems, it is much more severe for GW channels, which have orders of magnitude longer symbol durations, during which harvesting and modulation impair each other.

Lastly, the requirement for ultra-low-power tag devices severely impacts the quality of backscatter signals. Inaccurate and unstable frequencies of low-power oscillators contribute to this signal degradation. We have shown in Section 6.5.5 that even small CFOs in more performant devices strongly inhibit correct decoding, which is much more severe with low-power oscillators. This chapter also investigates the limitations and potential remedies for these effects.

7.1.2 General Approach

This chapter makes a first step towards closing the gap between the vision and the practical applicability of wirelessly powered backscattering sensor nodes. Therefore, we analyze power transfer and energy-harvesting techniques in the GW channel. We validate the analysis by prototyping a fully batteryless sensor node. To our knowledge, this work is the first to construct and evaluate a wirelessly powered prototype leveraging simultaneous backscatter communication in GW channels.

In particular, we tackle the challenges by integrating a state-of-the-art energy harvester IC that performs MPPT to partially match the harvester's input impedance to the transducer. The interdependency between harvesting and modulation is analyzed, and a mitigation strategy, i.e., dynamically decoupling the harvester during transmission, is evaluated. Finally, we employ digital phase-locked loops from the software-defined radio field to track frequency and sample timing offsets introduced by the unstable and inaccurate oscillators on the tag device.

As this chapter focuses mainly on wireless power transfer, we implemented the batteryless tag not with the higher-order modulation presented in Chapter 5. This technique requires a complex setup protocol and shared state between reader and tag. Instead, we used the more common binary switch between open and closed circuits, which is well-known in RFID and underwater backscatter. Hence, integrating our promising new approach to higher-order modulation into a low-power tag is still an open topic for future research. Instead, we employ a modulation where the tag mixes the carrier with an intermediate frequency signal and uses PSK on this intermediate carrier.




Figure 7.1: The general architecture of the presented system. The backscatter modulator and the energy harvester work on the same piezoelectric transducer.


7.1.3 Structure

Section 7.2 first proposes a low-power tag architecture with an integrated energy harvester IC and modulator circuit and describes the duty cycle of the tag. We then focus on the backscatter communication scheme in Section 7.3, describing how we implement intermediate carrier generation and demodulation. Finally, Section 7.3.3 analyzes the timing errors introduced by the oscillators and investigates the software-defined techniques for time- and frequency synchronization.

Section 7.4 first explores the potential for energy transmission and harvesting in GW channels and the maximum power point for the piezoelectric discs. This is followed by a detailed evaluation of the harvester's efficiency and characteristics. We then quantify the instability and inaccuracy of the low-power oscillators in the selected microcontroller and investigate the power consumption of the MCU. Section 7.4.4 examines the implemented backscatter scheme and achievable data rates and the influence of the SFO mitigation techniques. Finally, Section 7.4.5 peeks into the potential of transmitting power through more realistic large structures.

7.2 Wireless Power Transfer

This section first describes the design of our hardware prototype and explains critical design decisions. Figure 7.1 sketches the system architecture. The addition of wireless power transfer to the prototype does not require modification of the reader compared to the previous chapters. The reader generates a single-frequency carrier signal for energy transmission. However, the tag also includes a harvesting unit that extracts the energy from the received signal at the transducer. In parallel, the modulator is attached to the same transducer, which controls the load impedance and, hence, the reflection coefficient. In the following sections, we first focus on the energy harvesting unit and then explain the backscatter modulator in more detail.



■ **Figure 7.2:** The harvesting circuit connected to the piezoelectric transducer.

7.2.1 Harvester Design


The energy harvesting circuit collects energy from the tag's transducer and provides a regulated supply voltage suitable to run the MCU. Our goal in this chapter is to test the limitations of WPT in metal channels, requiring a flexible harvester that functions in a broad range of scenarios. That poses three challenges, namely

1. the voltage on the receiving transducer may be smaller than the required supply voltage for the MCU and potential sensors, and
2. the maximum power that is available for extraction at the transducer may be insufficient to power the MCU continuously, and
3. the transducer's output impedance may vary strongly between frequencies and channels, demanding a harvester that works efficiently in a broad range of output impedances.

The first challenge can be overcome by raising the input voltage to a higher level with a voltage booster circuit, the second requires buffering energy in a storage capacitor and duty-cycle the MCU, and the third requires a load adaption mechanism, e.g., MPPT.

Generic PMUs are commercially available and pose a flexible solution integrating some or all of the required features, i.e., voltage boosting and regulation, duty-cycling support, and MPPT. As sketched in Fig. 7.2, we decided to use the ADP5091 from Analog Devices [Dev23], which supports all of the required features. Its inbuilt boost regulator, which works with input voltages down to 80 mV, distinguishes it from alternative ICs. However, in cold start, i.e., when the harvester is entirely out of energy, a minimum threshold voltage of 300 mV is required to start the boost regulator.

We first amplify and rectify the AC voltage from the transducer with a voltage quadrupler circuit, increasing the range of input voltages in which the cold start threshold is fulfilled. The rectified voltage charges a 10 μF input capacitor C_{in} . To support duty cycling, the PMU charges a storage capacitor C_{st} , and has a digital output pin (PGOOD), which switches on once the charge voltage surpasses a configurable threshold V_{th+}



■ **Figure 7.3:** Working principle of the power-harvesting system including MPPT and duty cycling.

and turns off once it falls below a lower threshold V_{th-} . In our prototype PGOOD enables a low-power LDO regulator that provides the regulated supply voltage for the MCU.

Figure 7.3 sketches the charging and duty-cycling process. If the source can provide more power than the MCU consumes, the MCU can run continuously while the carrier is active. Otherwise, the MCU runs in a duty-cycled mode. The choice of C_{st} and the threshold voltages depends on the energy requirements of the application. Ideally, the charged capacitor stores enough energy to power the MCU through a full task cycle, e. g., sensing and transmitting the result. However, storing more energy increases the tag's response time, i. e., the time it takes until the tag is charged and can respond with a sensor reading. The available energy in the capacitor E_{st} can be computed as

$$E_{st} = \frac{1}{2} C_{st} \cdot (V_{th+}^2 - V_{th-}^2), \quad (7.1)$$

while the consumed energy E_{MCU} is

$$E_{MCU} = V_{MCU} \cdot \bar{I}_{MCU} \cdot \Delta t, \quad (7.2)$$

with supply voltage V_{MCU} and average current \bar{I}_{MCU} over a period Δt . In the presented prototype, we set $C_{st} = 200 \mu\text{F}$, $V_{th+} = 5 \text{ V}$, and $V_{th-} = 2 \text{ V}$. Assuming no other losses, the buffered energy supports an output power of 2.1 mW for one second, which is sufficient for most sensing tasks.

The PMU implements an MPPT controller, realized by only transferring energy from the input capacitor to the storage capacitor when V_{in} exceeds the configurable MPPT threshold. The threshold is set relative to the open-circuit voltage, i. e., the input voltage that emerges if no current is extracted. The PMU implements this by stopping harvesting and sampling the open-circuit voltage every 16 s. The optimal MPPT threshold varies between

different harvesting sources; we investigate the optimal choice for piezoelectric transducers in Section 7.4.1. We want to stress, however, that the MPPT circuit is not equivalent to impedance matching, as the transducer has a complex-valued internal impedance. Maximum power transfer between the transducer and load is achieved when the load impedance is the complex conjugate of the transducer's output impedance. MPPT can only match the resistance of the load, which is suboptimal. Unfortunately, no viable option exists to dynamically adapt the reactive load of the harvester in the required range of impedances.

Low-Power MCU

In addition to maximizing the harvested power, minimizing the tag's power consumption reduces charging times and required capacitor sizes. Since backscatter modulation— the switching of the transistors—does not require significant power, the MCU is the dominant contributor to total power consumption. For CMOS technology, the dynamic power P_{dyn} that an MCU consumes scales quadratically with the supply voltage V_{DD} , and linearly with the MCU's clock frequency f_{mcu} [Ng+22], i.e.,

$$P_{\text{dyn}} \propto f_{\text{mcu}} \cdot V_{\text{DD}}^2. \quad (7.3)$$

Therefore, we select a dedicated low-power MCU (STM32L073RZ) and leverage all its power-saving features. This includes reducing the supply voltage to 1.8 V, saving nearly 70 % of power consumption compared to 3.3 V. Furthermore, choosing the internal low-power oscillator and scaling the clock frequency to the minimum requirement is advantageous. As a downside, the on-chip low-power oscillators are unstable and inaccurate, introducing a SFO that must be detected and compensated at the receiver, which we explain in detail in Section 7.3.3.

7.3 Backscatter Communication

This section describes the backscatter communication scheme, which adopts a different modulation scheme than the previous chapters, as motivated in Section 7.1.2. We first describe the general principle of the modulation scheme and then explain the modifications necessary in the receiver hardware. We then explain the details of the software-defined receiver used to mitigate the inaccurate oscillator at the tag.

7.3.1 Modulation Scheme

The modulation scheme relies on only two states: short-circuiting the transducer, generating a reflection coefficient of $\Gamma_0 = -1$, or the impedance at the transducer ap-

proaches infinity, i.e., $\Gamma_1 = 1$. The tag switches between both states with an intermediate frequency f_i , generating a square wave reflection coefficient

$$\Gamma(t) = \text{sq}(\omega_i t + \phi(t)) = \frac{4}{\pi} \sum_{k=1}^{\infty} \frac{\sin((2k-1)\omega_i t + \phi(t))}{2k-1}, \quad (7.4)$$

where $\omega_i = 2\pi f_i$, and $\phi(t)$ is the phase of the square wave. The last equation stems from the well-known spectrum of a square wave, which has infinitely many frequency components at odd multiples of its base frequency. The tag modulates $\phi(t)$ to implement phase shift keying.

Mixing with the intermediate frequency enables the reader to remove the self-interference at f_c from the received signal because f_c is not included in the passband spectrum. To illustrate that, we write the backscattered signal in the passband $s(t)$ as

$$\begin{aligned} s(t) &= a_0 \sin(\omega_c t + \phi_c) \cdot \Gamma(t) \\ &= \frac{2a_0}{\pi} \sum_{k=1}^{\infty} \left[\frac{\cos([\omega_c - (2k-1)\omega_i]t + \phi_c - \phi(t))}{2k-1} \right. \\ &\quad \left. - \frac{\cos([\omega_c + (2k-1)\omega_i]t + \phi_c + \phi(t))}{2k-1} \right], \end{aligned} \quad (7.5)$$

where a_0 is the amplitude of the carrier signal at the tag. From Eq. (7.5) and the visualization in Fig. 7.4, we observe that the backscattered signal is shifted away from the carrier by odd multiples of ω_i . The intermediate frequency must be chosen appropriately to support the correct reception of the transmitted signal. First, f_i must be larger than half the bandwidth of the phase signal $\phi(t)$, because otherwise, the spectra of the passband signal overlap with f_c , and adjacent copies of the transmission signal overlap in the passband. However, f_i can not be too large because $f_c \pm f_i$ must still be within the transducer's limited bandwidth. Furthermore, generating larger intermediate frequencies becomes more demanding for the resource-constrained tag. We chose $f_i = 8$ kHz as a reasonable tradeoff with transducer bandwidth between 200 – 250 kHz and anticipated symbol rates below 4000 symbols per second.

Modulation and Decoupling

The tag realizes binary load modulation with two NMOS transistors in a differential configuration, as shown in Fig. 7.5. The differential configuration is required because the transducer generates a sinusoidal voltage, and a single transistor can not isolate currents on positive and negative drain-source voltages. In the differential configuration, at least one transistor isolates the transducer terminals at any drain-source voltage when the gate is pulled to ground. We selected transistors with a low gate-source thresh-




Figure 7.4: At the tag, the incoming sine carrier wave is modulated with a PSK-modulated square wave at an intermediate frequency from the MCU.


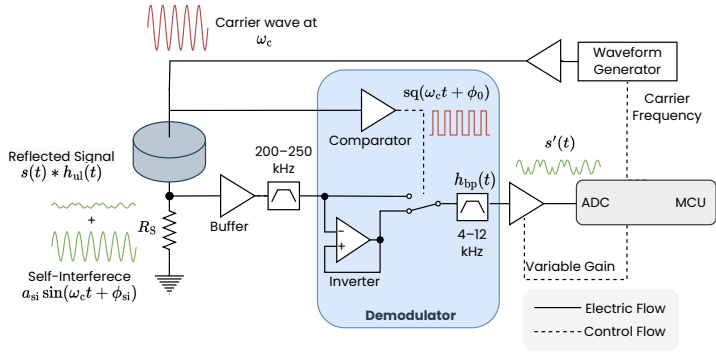


Figure 7.5: The modulator circuit at the tag, and the decoupling switch to detach the harvester during message transmission.

old voltage to ensure that the low-power microcontroller can switch the transistors completely on with its limited 1.8 V supply voltage.

A potential distortion of the backscattered signal arises because the harvester is connected in parallel with the modulator. Hence, when the transistors are switched off, the load at the transducer is dominated by the harvester’s input impedance. In the best case (for harvesting), this input impedance matches the transducer’s output impedance through MPPT, resulting in total signal energy absorption, i.e., $\Gamma = 0$. In reality, the effective harvester impedance varies over time, e.g., the MPPT threshold slowly changes and the harvester switches off when fully charged, which affects the backscattered signal. To evaluate if the coupling considerably distorts the backscattered signal, we added an analog switch to decouple the harvester from the transducer during message transmission.



■ **Figure 7.6:** The analog frontend removes the carrier from the received signal, with only the backscattered signal at the intermediate frequency remaining.

7.3.2 Analog Receiver

The analog frontend in this prototype is similar to the reader circuit presented in Chapter 6 (see Fig. 7.6). It buffers and filters the received current through the transducer with a fourth-order filter with passband from 200 – 250 kHz. Then, a demodulator mixes it with the carrier frequency f_c . As the passband in this prototype is not overlapping with the carrier frequency, the demodulator only has a single branch instead of individual I and Q branches. Furthermore, the filter with impulse response $h_{bp}(t)$ following the mixer is a fourth-order bandpass between 4 – 12 kHz instead of a lowpass filter. This filter completely removes the self-interference at f_c . The remaining signal is amplified and sampled by the microcontroller. Using Eq. (7.5), we can write the signal output signal of the analog frontend, $s'(t)$, as

$$\begin{aligned} s'(t) &= [s(t) * h_{ul}] \cdot \text{sq}(\omega_c t) * h_{bp}(t) \\ &\approx \frac{4a_0}{\pi} \sin(\phi_c) \cdot \sin(\omega_i t + \phi(t)) * h'_{ul}, \end{aligned} \quad (7.6)$$

where h'_{ul} is the uplink channel impulse response shifted to the intermediate band.

Demodulation with an intermediate frequency, also referred to as *heterodyne* demodulation, has two advantages over the direct conversion receiver used in Section 6.4.1. First, only one demodulator branch is required, as the relevant phase information is encoded in the intermediate frequency instead of the carrier. The I and Q components are extracted later in software, where digital filters make it much easier to achieve I and Q filter branches with identical impulse responses. Second, after removing the self-interference, the isolated backscattered signal can be amplified before sampling, avoiding possible amplifier saturation due to strong self-interference.

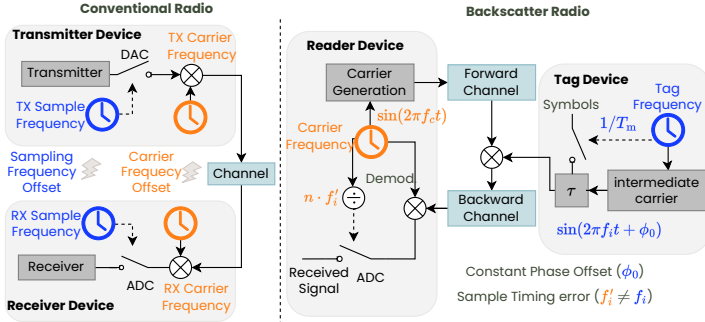


Figure 7.7: The intermediate frequency generated at the tag and the sampling clock are not synchronized, leading to a constant phase offset, and sampling frequency error. In contrast to conventional communication devices, all clock errors stem from the same mismatch between the reader and tag main oscillators.

7.3.3 Software-Defined Receiver

After sampling, the backscattered signal is again bandpass-filtered digitally to reduce the sampling noise. We implement a second-order Butterworth bandpass filter and adapt the passband dynamically to the symbol rate, i.e., the filter's passband is between $f_i - 1/(2T_m)$ and $f_i + 1/(2T_m)$. The transformation into baseband can be done very efficiently when the sampling rate is four times f_i as

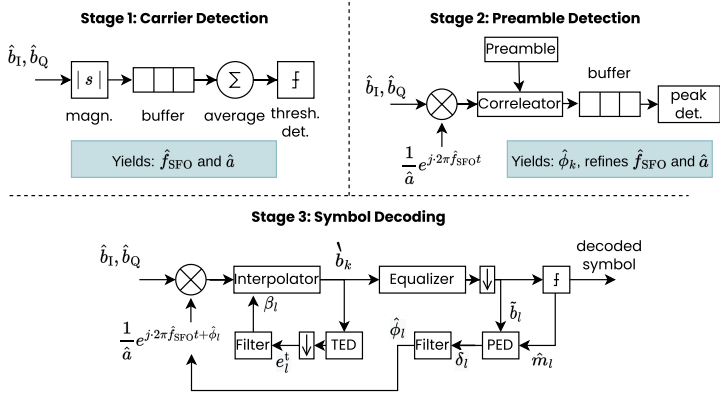
$$\hat{b}_I[i] = (-1)^{i+1} s' (2i \cdot T_m) = (-1)^{i+1} s' \left(\frac{i\pi}{\omega_i} \right) \quad (7.7)$$

$$\hat{b}_Q[i] = (-1)^i s' ((2i + 1) \cdot T_m) = (-1)^i s' \left(\frac{(i + 0.5)\pi}{\omega_i} \right) . \quad (7.8)$$

Finally, we further downsample the signal to a desired multiple of the symbol rate N_m , yielding the final baseband signal.

Because signal generation and reception are based on different physical clocks, visualized in Fig. 7.7, two types of timing errors arise in the baseband signal. A constant phase offset (PO) occurs because the sampling clock and the intermediate carrier run out of phase. Furthermore, both run at slightly different frequencies, causing an SFO. The SFO causes the received signal to rotate over time, which looks similar to a carrier frequency offset. Furthermore, the symbol timing changes throughout the signal due to the SFO. Lastly, this SFO may vary over time because the oscillators used in ultra-low-power devices are unstable.

This clock setup differs from traditional wireless communication, where the carrier is typically in a much higher frequency range and is generated by a separate oscillator.



■ **Figure 7.8:** The software-defined receiver pipeline during message reception.

Therefore, the carrier frequency and sample timing offsets are typically independent and must be estimated separately. In our setup, all timing errors are caused by the same clock mismatch, i.e., the mismatch between the main oscillators at the reader and tag. We can, hence, estimate the clock mismatch jointly.

We apply standard techniques from software-defined radio receivers, using DPLLs to track and correct the SFO. The reception process consists of the three stages shown in Fig. 7.8. First, in the *signal detection stage*, the receiver detects the presence of a signal with a signal power threshold. Once the signal power exceeds the threshold, the receiver performs a rough SFO and signal amplitude estimation. The SFO estimate \hat{f}_{CFO} is essential because subsequent frame synchronization fails when the SFO is large. At the same time, the signal is scaled to unity power for each message, as the optimal gains and parameters for DPLLs in subsequent stages depend on signal amplitude \hat{a} . We oversample the signal with a multiple N_m of the symbol rate to support timing error detection and compensation.

In the second stage, the receiver corrects the incoming samples by phase-shifting and scaling them with the estimated phase offset and amplitude. It then searches for the start of a frame by cross-correlating it with the preamble. Once a peak in the correlation exceeds a fixed threshold, the start of the message is detected. In the following steps, the DPLLs track the phase and timing errors throughout the message. The detected preamble is also used to refine the estimate of the initial SFO and STE α for DPLL initialization. Second, we also use the preamble to train the equalizer. As the preamble is used for parameter estimation and equalizer training, its length is crucial. A longer preamble increases the accuracy of the estimated parameters and equalizer weights but also adds overhead and keeps the tag active without transmitting any payload.

In the last stage, the payload symbols are decoded. At the input of the stage, the signal is scaled and phase-shifted to compensate for *SFO* and Phase Offset (*PO*). Then an interpolator corrects the *STE*. The samples are subsequently equalized with a decision feedback equalizer to remove ISI. Finally, a symbol decision is performed. After each symbol, the timing and phase error *DPLLs* are updated.

The *DPLLs* consist of an error detector and a loop filter. The exact phase error of the k -th symbol δ_l is

$$\delta_l = \angle(\tilde{b}_l) - \angle(\hat{m}_l), \quad (7.9)$$

with equalized symbol \tilde{b}_l and symbol decision \hat{m}_l . To avoid computing the phase angle explicitly—which is computationally expensive—we approximate the phase error by

$$\delta_l \approx \text{Im}\{\tilde{b}_l \cdot \hat{m}_l^*\}. \quad (7.10)$$

With the PSK constellation points on the unit circle ($|\tilde{b}_l| \approx 1$) and assuming small phase errors, the approximation is valid as

$$\text{Im}\{\tilde{b}_l \cdot \hat{m}_l^*\} = \left| \tilde{b}_l \right| \sin(\angle(\tilde{b}_l) - \angle(\hat{m}_l)) \approx \angle(\tilde{b}_l) - \angle(\hat{m}_l). \quad (7.11)$$

The loop filter for the phase *DPLL* is an integrator with a filter gain g_p . After every symbol, the estimate of the total phase error $\hat{\phi}_l$ is updated as

$$\hat{\phi}_l = \hat{\phi}_{l-1} + g_p \cdot \delta_l. \quad (7.12)$$

Analogously, the error for the sample timing error *DPLL* is computed using a derivative timing error detector. The timing error for the l -th symbol, e_l^t , is approximated as

$$e_l^t = \text{Re} \left\{ \frac{d\hat{b}_k}{dt} \cdot \hat{b}_k^* \right\} \propto \text{Re}\{(\hat{b}_{k+1} - \hat{b}_{k-1}) \cdot \hat{b}_k^*\}. \quad (7.13)$$

The loop filter is again an integrator with filter gain g_t , such that the update of the interpolator input β_l is computed as

$$\beta_l = \beta_{l-1} + g_t \cdot e_l^t. \quad (7.14)$$

7.4 Evaluation

In the following evaluation, we will first focus on wireless energy transfer. Then, we characterize the properties, e.g., the clock stability and power consumption in our prototypes. A detailed investigation of the backscatter communication performance and the gains of *SFO* tracking will also be given. Lastly, we explore longer distances and conditions found in real-world scenarios.

7.4.1 Wireless Energy Transfer

This section first analyzes the potential for acoustic energy transfer through GW channels over several meters. Afterward, we measure the efficiency that our harvester achieves.

Harvesting Potential

The maximum power point for piezoelectric transducers is essential to configure the MPPT threshold to maximize extracted power from the source. Piezoelectric sources are typically modeled as voltage sources with a nonzero internal series impedance. In that case, the maximum power point is established when the load matches the source's internal resistance. We validated this assumption by measuring voltages and currents over the piezoelectric source with different load resistances and repeated the experiment with various carrier frequencies. The results and experiment setup are shown in Fig. 7.9 (a), yielding two conclusions. First, the relationship between current and voltage is approximately linear, indicating that maximum power point occurs when the MPPT threshold is half the open-circuit voltage. This result confirms observations from [TMÖ22] and [STM20]. Second, the absolute value of the maximum power point varies enormously with the chosen carrier frequency, which emphasizes the requirement for a dynamic MPPT mechanism compared to a fixed threshold voltage.

Next, we ask how much AC power we can expect to extract at the sensor node at a certain distance. The selected carrier frequency strongly influences this power level, i.e., frequencies that resonate in the structure enable more efficient power transfer. In Fig. 7.9 (c), we show the maximum harvestable powers with optimal resistive matching in a frequency band from 200 – 250 kHz in channel $CH_{4 \rightarrow 1}$. The measurements show that resonant frequencies produce narrow spikes with more than 8.6 times the average harvestable power over all frequencies. At 3 m, the best frequencies provided an output power of 4 mW (6 dBm). At that frequency, the reader has generated an input power of 141 mW (21 dBm), leading to a transfer efficiency of 2.8% (-15 dB).

Repeating the measurement for multiple channels with varying distances shows no clear trend. The maximum power over a 1 m channel is about 10 mW (10 dBm) with roughly 8.3% transfer efficiency (-11 dB). The 2 m channel shows almost identical power levels to the 3 m channel. However, the optimal frequencies varied strongly between channels. These results indicate that within distances of a few meters, the distance is not the dominant factor determining transfer efficiency. Strong frequency dependency suggests that positive interference patterns are essential. These are influenced by structure geometry and the exact placement of the transducers, which are generally unknown during system design.

Harvester Performance

The previous section investigated the maximum extractable AC power at the transducer. However, the harvester converts AC into DC power to supply microcontrollers and sensors.

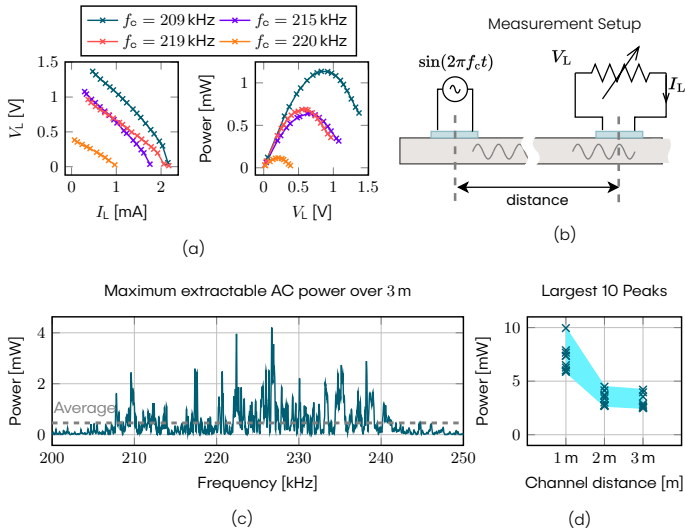


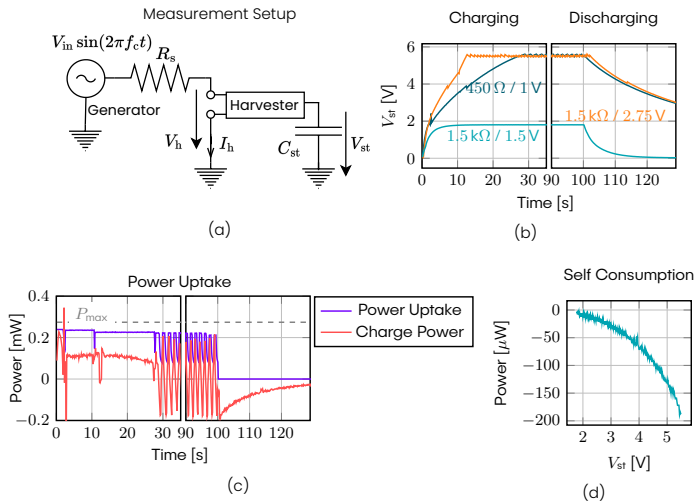
Figure 7.9: The maximum power point (a) is achieved at half the open circuit voltage, measured with (b) the experiment setup. Wirelessly transmitted power (c) is highly frequency dependent, and (d) the distance between transmitter and receiver is not dominant.

In this section, we inspect the harvester’s efficiency. Multiple factors may degrade the harvester’s performance. First, the MPPT tracker of the PMU may mismatch the input impedance of the harvester, not extracting the maximum available power from the source. Second, the harvester’s self-consumption reduces the energy available to the application. Self-consumption includes power consumed to boost the input voltage in the PMU and leakage currents from the storage and input capacitors.

We define the harvester efficiency as the power charging the storage capacitor divided by the power that is theoretically available at the source. We evaluate the input power by applying a sine generator with variable series resistance R_s on the harvester inputs (see Fig. 7.10 (a)). Measuring the current I_h through, and voltage drop V_h over the harvester, we determine the power uptake. Furthermore, the maximum available power to the harvester $P_{h,max}$ is, according to the maximum power transfer theorem,

$$P_h^{\max} = \frac{V_{in,eff}^2}{4 \cdot R_s} = \frac{V_{in}^2}{8 \cdot R_s}, \quad (7.15)$$

where the sinusoidal source generates the voltage V_{in} with effective (RMS) voltage $V_{in,eff}$. Finally, we approximate the harvested power P_h from the storage capacitor



■ **Figure 7.10:** (a) The measurement setup. The charge curves (b) show the voltage at the storage capacitor over time. The harvester's power uptake is suboptimal (c), and the PMU consumes part of the power itself. Self-consumption (d) is dependent on the voltage at the storage capacitor.

voltage

$$P_h = \frac{\Delta E}{\Delta t} = \frac{C_{st} (V_{st}^2(t + \Delta t) - V_{st}^2(t))}{2 \cdot \Delta t}, \quad (7.16)$$

where ΔE is the change of stored energy over a short period Δt . During this measurement, we disconnected the MCU from the harvester so that all consumed power is harvester self-consumption.

Figure 7.10 (b) shows the storage capacitor voltage over time for different source resistances and input voltages, with a sinusoidal carrier active for 100 s. The three displayed curves were recorded with different source power levels, realized by varying R_s and V_{in} . The time required to charge the storage capacitor up to the maximum voltage varies with source power, and if the power is insufficient, the storage voltage never reaches its terminal voltage but plateaus at a level where the harvested power and self-consumption are in equilibrium. Figure 7.10 (c) provides a more detailed analysis of the harvested power. In the beginning, the harvester takes up about 239 μW , 87% of the maximum available power from the source. That indicates a suboptimal impedance match between the source and harvester, i. e., MPPT does not match the source's resistance perfectly. Furthermore, the self-consumption is considerable, as only around 115 μW (48% of the uptake) end up in the storage capacitor. Once the source stops, the harvester loses up

to 200 μW due to leakage and quiescent current. However, the self-consumption shown in Fig. 7.10 (d) depends on the voltage over the storage capacitor, indicating that leakage is the dominant factor. In the performed experiments, the harvester efficiency is around 40 – 50%. Exact efficiency is highly variable, as it depends on the source power and on the voltage over the storage capacitor. We can confirm, however, that these observed efficiencies are consistent with previously reported harvesters with comparable source powers, e.g., as reported in [IAS22]. When available power reduces to tens of microwatts, the efficiency shrinks to only several percent, as the harvester consumes most of the available power itself.

In light of the presented measurements, the 300 mV nominal minimum turn-on voltage of the PMU is not a significant limitation. Typical piezoelectric sources have considerable internal series resistances—the utilized transducers have between 500 Ω and 1 k Ω . Hence, voltages significantly higher than the turn-on voltage are required to provide sufficient power to overcome self-consumption. However, further optimization of the PMU can reduce leakage and harvest more efficient from low-power sources: When the harvester must work with sources providing very little power, the charging threshold at the storage capacitor should be reduced, and a larger capacitor should be chosen instead. However, a disadvantage of this strategy is that a larger capacitor increases the charge-up time required until the capacitor voltage exceeds the minimum required threshold to start the MCU.

7.4.2 Low-Power Oscillator

This section compares the internal medium-speed 4.194 MHz oscillator integrated into the MCU with a 4 MHz external crystal oscillator. The internal oscillator is energy efficient, does not require additional space, and enables the MCU to run in its lowest power consumption mode. Its drawback is that it is less accurate and more unstable than an external crystal, i.e., the clock frequency varies over time. In this section, we first investigate the absolute deviation from their nominal frequency and the variation over time of the internal and external clocks and measure their influence on tag power consumption. We then evaluate how clock instability affects communication performance and if the presented mitigation strategy is effective.

To evaluate the clock stability, we record the square wave produced by the tag's MCU with an oscilloscope and compare it with an ideal square wave with the nominal carrier frequency. Figure 7.11 shows the measured frequency offset over time. The MCU's supply voltage significantly impacts the frequency of the internal oscillator, which is roughly 70 Hz larger—more than 800% increase—at $V_{DD} = 1.8\text{ V}$ compared to 3.3 V. Additionally, even with constant supply voltage, the instantaneous frequency varies throughout a single message by more than 5 Hz. To put that in perspective, consider a 4-PSK modulation scheme in which a phase offset of up to 45° can be tolerated without decoding errors. This would allow an instantaneous SFO of 5 Hz for up to 20 ms. However, since we have to use a nonlinear equalizer to mitigate the strong ISI, even smaller phase




Figure 7.11: Sample Frequency Offset measured throughout multiple messages (different colors). The supply voltage level strongly influences the generated carrier frequency when using the internal oscillator, but it also fluctuates significantly during the message. The external oscillator is independent of the supply voltage and much more stable.

offsets may lead to large errors in the equalizer output. Therefore, the offset of 5 Hz is a serious problem for correct message reception.

For comparison, we also measure the frequency offset of the external crystal. The instantaneous frequency is much more stable and only varies less than a tenth of a Hertz during a message transmission. However, it also has a constant frequency offset compared to the nominal frequency of about 6 Hz. Therefore, a receiver based on the external oscillator may rely on a single *SFO* estimate at the beginning of a message. However, depending on the internal oscillator, a receiver must dynamically track and correct phase offsets during message reception. Moreover, as the *SFO* also affects the symbol timing, the receiver must also track and correct symbol timing dynamically. Considering these severe disadvantages of the low-power oscillator, it is preferable to use an external oscillator. However, the following section will compare the differences in power consumption between both clocks, showing that the external crystal requires significantly more power.

7.4.3 Tag Power Consumption

The tag's power consumption determines the minimum required energy buffer size, and lower total energy consumption enables better duty cycles. In our characterization, we distinguish consumed energy in the *PMU self-consumption* (evaluated in Section 7.4.1), *static consumption* from the *MCU*, e.g., for supplying the clock and the required peripherals, and *dynamic consumption* for modulating the backscatter signal. In this section, we investigate the contributions of the different factors to the total power consumption.




Figure 7.12: Tag power consumption with different oscillators. External crystal oscillators at equivalent frequencies consume significantly more power than low-power internal oscillators. The backscatter modulation itself only consumes around $20\ \mu\text{W}$.

The MCU provides various power-saving features, i. e., several low-power modes and clock frequency scaling. To reduce static consumption as much as possible, we supply the MCU with 1.8 V, scale the clock frequency down to 4 MHz, and use the STOP mode, in which the clock and most peripherals are halted, whenever possible. However, hardware timers and a pulse-waveform generator peripheral must run to generate the intermediate carrier, preventing STOP mode usage during transmission. Figure 7.12 shows the tag's power consumption during the transmission of one message. We observe the consumption of only $6\ \mu\text{W}$ in STOP mode, but total consumption rises strongly during message transmission. In direct comparison at 4 MHz the external crystal consumes additional $500\ \mu\text{W}$ compared to the internal oscillator—an increase of 80%. Additionally, the external oscillator requires a startup time of at least 2 ms after wakeup, and power demand spikes during oscillator startup. Given the high power requirement, we conclude that using the internal oscillator is highly desirable.

Furthermore, dynamic frequency scaling is another advantage of the internal oscillator. It can be configured dynamically to different rates depending on the requirement of the current task. For example, we found that a clock frequency of 1 MHz is sufficient for backscattering data at $2\ \text{kbit s}^{-1}$ ($T_m = 1\ \text{ms}$ with 4-PSK). At that rate, the MCU power consumption drops to only $300\ \mu\text{W}$ ($-5\ \text{dBm}$). With the external oscillator, the MCU may use prescalers and PLLs to change the main clock in the MCU and reduce the power consumption within the MCU. However, the oscillator will always run at its resonance frequency, and driving it requires the same amount of power.

We measured dynamic consumption by measuring power consumption with and without connected modulator circuits. The difference between both reveals the dynamic power required for backscattering, i. e., switching the transistor's gate voltage. The backscattering modulation only consumes about $20\ \mu\text{W}$. Hence, the—mostly idling—MCU consumes most of the power.

The presented results show that it is feasible to perform communication with ultra-low power consumption suitable for batteryless systems using general-purpose MCUs and off-the-shelf components. Static power consumption is the dominant factor, indicating

that more integrated solutions, e.g., based on FPGAs or ASICs, may have a large potential for reduction of tag power consumption, similar to the ASIC-based RFID tags that can operate on tens of microwatts (cf. Table 3.1). This section also shows that using the internal oscillator is beneficial because of the reduced overall power consumption and dynamic frequency scaling options.

7.4.4 Backscatter Communication

This section focuses on the reliability of backscatter communication. We first characterize the examined channels. The setup for the following evaluations is the 3.5 m long beam with several attached piezo transducers presented in Section 4.2.2. To explore the effectivity of the SFO adaption, we compare the reliability of the stable and accurate external crystal oscillator with the low-power internal oscillator. Afterward, we look at BERs of the backscatter communication in six different narrowband channels.

We first analyze the characteristics of the backscatter channel in the test structure. The top graphs in Fig. 7.13 show the SNR for different carrier frequencies. As with the higher-order modulation scheme in Chapter 5, we observe strong fluctuation by more than 20 dB within a few hundred hertz. Remarkably, the peak SNRs in the 3 m long channel $CH_{3 \rightarrow 1}$ were even higher than in the 2 m long channel $CH_{4 \rightarrow 1}$, again indicating that interference patterns dominate attenuation over short distances. To test our approaches thoroughly in various channels, we pick three distinct carrier frequencies for each distance for further evaluation. The bottom graphs in Fig. 7.13 show the impulse responses for the selected channels. All channels exhibit extended reverberation for about 10 ms. However, channel impulse responses vary significantly with the carrier frequency. Hence, SNR alone is likely unsuitable for estimating a carrier frequency's communication reliability.

SDR Pipeline

Next, we transmitted 250 messages through each channel for each selected frequency and decoded them with the SDR pipeline described in Section 7.3.3. We compare transmissions in which the tag generated the signal with the internal oscillator versus the external oscillator. The modulation scheme was 4-PSK at a symbol rate of 1000 symbols per second, yielding a total data rate of 2 kbit s^{-1} . Each packet contained 200 payload bits.

First, we validate that phase and timing error tracking is advantageous. Therefore, we decode all messages without tracking, i.e., only estimating the SFO from the preamble and assuming it is constant throughout the message. We then use the DPLL to track the SFO dynamically. The packet error rates (PER) displayed in Fig. 7.14 reveal that nearly all packets are correctly decoded when using the stable external oscillator, even without SFO tracking. However, about 25% of the packets contain errors when using the internal oscillator. In both cases, the DPLL enables error-free reception of all packets.

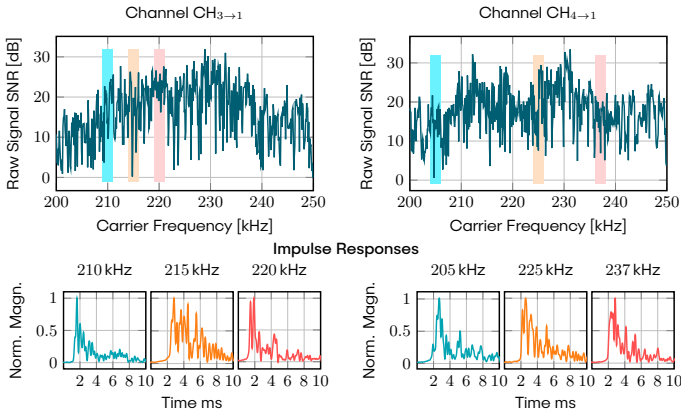


Figure 7.13: SNR of different carrier frequencies in two selected physical channels and the magnitude impulse responses of six selected narrowband channels.






Figure 7.14: Tracking phase and timing errors with a PLL improves packet error rate drastically when using the unstable internal oscillator.

We then compute the Error Vector Magnitude (EVM) of the equalizer output. Figure 7.15 displays the distribution of EVMs in all messages. The PLLs successfully mitigate the instability of the internal oscillator. The equalizer output’s SNR with the internal oscillator is comparable to the EVM with the external oscillator, and, in some channels, the internal oscillator even outperforms the external one. While this seems paradoxical at first, a possible explanation might be that the underlying channels are not equal—even with the same nominal carrier frequency: the low-power oscillator had an SFO of up to 70 Hz, which slightly shifts the transmission band. Surprisingly, the results for the different channels are very similar, even though the SNR of the received signal varies by more than 10 dB between the channels.

Furthermore, we also evaluated if decoupling the harvester from the modulator circuit improves the signal quality. Surprisingly, we did not observe a significant difference



■ **Figure 7.15:** When using a PLL, internal and external oscillators perform similarly, and all messages are received and decoded correctly in the channels under test.



■ **Figure 7.16:** The effect of the decoupling switch on the decoding performance is marginal.

between the cases where the harvester is connected—and influences the tag’s load impedance—and when it is decoupled with the analog switch. A direct comparison of 250 messages with and without decoupling in Fig. 7.16 shows even a 1% lower EVMs without the decoupling switch. Therefore, we conclude that the decoupling switch described in Section 7.3.1 is not required for concurrent backscatter transmission and harvesting. The distortion introduced by varying harvester impedances seems to be negligible.

Simulation Study

In real-world experiments, we have limited control over all variables, i.e., the channel impulse response, the extent to which the tag oscillator frequencies vary, and the amount of noise in the signal. To characterize the influence of various parameters on the receiver, we also implemented a channel simulation incorporating multipath propagation, time-variant *SFO*, and noise levels. In Figure 7.17 (left), we show the influence of different channel impulse responses on the output EVM. We compare three channel impulse responses: An ideal channel without ISI and two previously recorded channels. Gaussian noise with varying standard deviation was added after convolving the transmitted signal with the

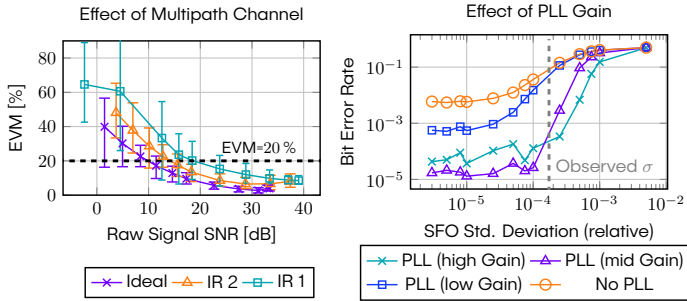


Figure 7.17: The influence of the channel impulse response on the receiver’s EVM is characterized (right), where IR1 was recorded from CH₄→₁ at $f_c = 225$ kHz, and IR2 stems from CH₃→₁ at $f_c = 210$ kHz. The choice of PLL gain (right) depends on SFO variability. During this evaluation, simulation signals had an SNR of 15 dB.

impulse response. We observe that the impulse response influences the receiver’s output error significantly: In the ideal channel, the EVM is more than 30 % lower at comparable SNR than with impulse response IR1, and channel IR2 produces a significantly lower EVM at the equalizer output than IR1. In other words, some channels are more resilient against noise than others. To reach an acceptable EVM of 20 %, the ideal channel requires an SNR of 10 dB, while IR1 requires a 9 dB higher SNR.

In Figure 7.17 (right), we simulated a random time-variant SFO similar with a distribution similar to the observed variations from the low-power oscillator shown in Fig. 7.11. We varied the simulated SFO standard deviation and measured the BER over 5000 simulated messages (1 million bits). The results indicate that even with an SFO close to zero, the decoder without PLL has a high BER ($\approx 5 \cdot 10^{-3}$). This counterintuitive result can be attributed to errors in estimating the SFO during preamble detection: Even if the SFO is completely constant during the whole message, its estimation is not exact due to noise in the signal. Without phase and frequency tracking, the small estimation error can accumulate into larger phase and timing errors throughout the message. The phase tracking PLL mitigates the error accumulation. In direct comparison, a decoder with PLL significantly reduces the bit error rate by more than two orders of magnitude to roughly $1 \cdot 10^{-5}$. However, the gain of the PLL plays an important role. With an SFO deviation in the range of $2 \cdot 10^{-4}$ (200 ppm), which we observed on the low-power oscillator, the decoders with medium to high gain PLL achieved BERs of about $3 \cdot 10^{-4}$.

7.4.5 Larger Structures

In prior work—including the previous chapters—WPT was only evaluated on small and isolated structures, e. g., small plates or pipes. This section peeks into the potential of WPT




Figure 7.18: Experiment on a large structure with transducers placed in distances 1, 2, 3, 4, and 5 meters. The graph shows the amplitudes of the ten best frequencies over different distances.

on more realistic real-world structures. We, therefore, attached several transducers to a steel beam supporting a crane shown in Fig. 7.18, with transducer distances of 1, 2, 3, 4, and 5 meters. We used the reader prototype to generate a carrier wave and measured the signal amplitudes at the transducers.

Compared to the smaller specimen we used in the previous evaluations, the amplitudes in the large beam are much smaller. The dominant cause is likely the larger thickness (2 cm compared to 5 mm in the previously tested specimen), where more energy is required to excite a comparable wave amplitude in the structure. When choosing the resonant frequencies, amplitudes in the large beam reach up to 400 mV over 1 and 2m distances, and still about 300 mV at 5 m. While these amplitudes are only a fraction of what we observed in the smaller specimen, the observed amplitudes also suggest that the acoustic waves experience little attenuation when traveling through the metal beam, i.e., a fivefold distance increase from 1m to 5 m only corresponded to a 25% decrease in amplitude. Hence, large structures may require more excitation power, but large distances are feasible once a wave is excited.

Unfortunately, the wirelessly transmitted power over this structure was insufficient to fully charge our prototype tag. However, several improvements may still enable batteryless tags on such large structures. First, increased *sender power* can provide sufficient energy: The sender power in our tag is in the range of 10 dBm, while typical RFID readers transmit up to 30 dBm, and in underwater backscatter, transmission powers of up to 59 dBm were used. Second, the chosen transducers and their frequency band do not fit well to the large beam, with a thickness of more than 2 cm. *Larger transducers* combined with *lower carrier frequencies* likely improve the efficiency of the wave excitation. Third, the harvester can also be optimized to work with less available power in a longer duty cycle, as discussed in Section 7.4.1.

7.5 Summary and Discussion

This chapter motivated that backscatter communication can be of significant interest in reducing power consumption during communication between sensor nodes embedded in metals. While prior studies employed only active communication in guided wave channels together with WPT in metal structures, this chapter demonstrated the feasibility of combined power transfer and backscatter communication.

Using the presented prototype, we investigated the potential of wireless power transfer over multi-meter distances, where a maximum of more than 10 mW was transferred over 1 m with an efficiency of 8%. Over 3 m distance, efficiency is reduced drastically to only 2.8%, but still, 4 mW were transferred. These quantities support a wide range of sensing applications: Tang et al. [TMÖ22] report SHM measurements using lamb waves at sub-milliwatt power. Many other sensors require even less power, e.g., low-resolution image recording [Afz+22] can be achieved at several hundred microwatts, acceleration measurements at tens of microwatts, strain gauge, temperature, and humidity measurements consume even less. Furthermore, a duty cycle can realize higher-power applications requiring extended computations. The presented harvester requires the source to provide at least 200 μ W to charge the tag fully, but we discussed potential optimizations to reduce this requirement.

The overall harvester characteristics relate well to harvesters presented in RF harvesting, e.g., the harvester presented in [IAS22] showed dramatic drops in efficiency when the energy source only provides several tens of μ W. The WPT efficiency of 8% over 1 m exceeds the observations from [Kiz+17] (0.1% at 1 m) and is orders of magnitude larger than efficiencies known from RFID. However, transfer efficiencies are strongly related to the structure, the transducer's geometry, and the bounding method. Therefore, a direct comparison of the achieved efficiencies is not informative.

In theory, the signal power of guided lamb waves is attenuated only linearly with distance. We performed practical measurements on a large structure, supporting the low attenuation over large distances, i.e., the maximum observed voltage at the receiver after 5 m was still about 75% of the observed voltage over 1 m.

Tailoring the energy transfer method to a specific structure could increase the overall transfer efficiency we achieved with a general WPT system—less than 3% over 3 m distance. Different optimization techniques have been proposed, such as determining the resonance frequencies of the structure with simulation or measurement, and then choosing transducer's with the exact same resonance. A different approach is to use transducers that produce directed waves aligned between transducer and tag. However, such optimizations are hard to achieve as real structures often have complex geometry and unknown boundary conditions, prohibiting previous determination of optimal locations and frequencies. We, therefore, aimed to investigate a *plug-and-play* solution that does not rely on prior knowledge about the channel. However, dynamic beamforming using a piezo array may be a promising future optimization.

We tackled two challenges for real-world deployments of acoustic backscatter devices, i.e., a) reader self-interference limits the viable communication distances, and b) low-power devices use inaccurate and unstable oscillators that significantly impede message reception. The presented prototype showed that using digital phase-locked loops, we can successfully track the errors induced by the unstable oscillators and decode messages at a data rate of 2 kbit s^{-1} over 3 m, which was not feasible without tracking mechanisms in place. These data rates support low-cost and low-complexity implementations of battery-free embedded sensor nodes in closed metal containers for structural and industrial monitoring, which can be constructed from off-the-shelf discrete components.

The presented backscatter tag in this study required less than $300 \mu\text{W}$ during transmission, of which only $20 \mu\text{W}$ were required for modulation, and the **MCU** required the rest. In comparison, Shaik [STM20] reports power consumption of $100 - 1000 \mu\text{W}$ during active uplink communication, dependent on the data rate. While an ASIC implementation is potentially much more energy efficient, we showed that a working system is achievable with commercially available off-the-shelf components. For research prototypes and small series production, that is desirable as ASIC production is very costly and time-consuming. The backscatter communication decoding with **DPLLs** proved robust against the **SFO** induced by low-power oscillators. The **WPT** will likely be the limiting factor for longer distances in large structures. However, backscatter communication can potentially be used over much longer distances in battery-powered sensor nodes to reduce the power requirement of communication and increase battery life.

Wideband Downlink Communication

8.1 Overview

The previous chapters focused on uplink communication from the tag to the reader. However, a downlink is required

- for many applications, which may need central orchestration of multiple sensor nodes to perform measurements simultaneously,
- for a data-link layer, which includes multiple-access schemes that are often based on interrogation from the reader, and
- for parameter identification and negotiation in the physical layer itself, as demonstrated in the backscatter scheme presented in Chapter 5 demonstrated, where the reader estimates the optimal load impedances and transmits these to each tag.

Similarly to the uplink communication, high data rates are also beneficial in the downlink. Even if the application does not transmit large amounts of data, sampling and decoding require significant computation at the tag, quickly depleting its scarce energy resources. Therefore, keeping the message duration as short as possible reduces its overall energy demand. Furthermore, in applications with many tags, control messages, e.g., for tag interrogation, may take up most of the available transmission time, limiting the number of sensor nodes a single reader can manage effectively.

The downlink can be regarded as active communication because the reader can embed information in the actively generated carrier it uses to power the tags. Unfortunately, strong ISI also exists in the downlink channel, inhibiting achievable data rates. Accordingly, many prior studies on active acoustic communication through GW channels report data rates in the range of only 20 – 200 bit s⁻¹. ISI mitigation strategies, such as OFDM or equalization, are unsuitable for resource-constrained tag devices that lack the computational power and energy budget to demodulate and decode such complex signals.

Simultaneous communication and energy harvesting on the same transducer is another challenge that complicates downlink communication in two ways: On the one hand, carrier modulation influences the signal power available at the tag. On the other hand, extracting current from the tag's transducer changes the observed voltage over the transducer, distorting the signal the tag uses to detect and decode incoming messages. Finally, the *SFO* between the reader and tag also affects the correct reception of the received messages; lightweight mechanisms must be found to mitigate it with few resources.

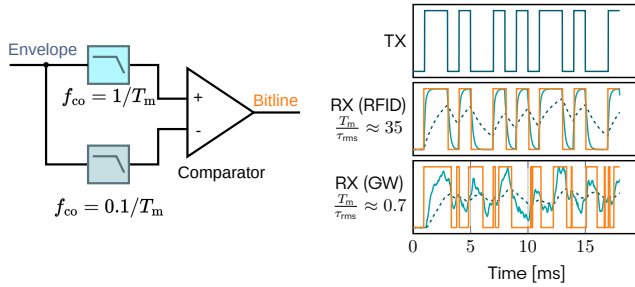
8.1.1 Problem Analysis

Compared to the uplink, the downlink communication flips the resource imbalance between the transmitter and receiver. The tag is now the receiver. Its scarce resources prohibit complex active filters, high sampling rates, or computationally intense operations. In practically all backscatter applications, the tags employ a simple diode-based envelope detector to demodulate the received signal. Consequently, modulation schemes that modulate the phase or frequency of the carrier are incompatible with such a receiver. In contrast, the reader, which generates the transmitted signal, can leverage much more resources, favoring ISI mitigation on the reader side.


Conventional backscatter systems used Continuous Wave (*CW*) modulation schemes, e.g., OOK, PDM, or PIM. CW modulation generates a narrowband signal, which is particularly disadvantageous in the strong frequency-selective GW channel. Consequently, the chosen carrier frequency strongly impacts ISI intensity. Finding optimal carrier frequencies dynamically during runtime is not straightforward because the resource-constrained tag can not measure channel impulse responses accurately.

Moreover, conventional backscatter systems typically decode the envelope using a comparator to compare the current envelope with an average of the envelope over multiple symbol periods, as demonstrated by the WISP platform (cf. Fig. 3.5). This approach allows the MCU to remain in sleep mode and save energy most of the time. Analog demodulation works well in RFID applications, which report RMS delay spreads below 80 ns at symbol durations of several tens of μ s. The measured GW channels, however, had delay spreads of several milliseconds, where more advanced decoding schemes are required. To demonstrate this, we show the result of such analog demodulation in Fig. 8.1, once with a channel characteristic typical for RFID channels and once for GW channels.

Instead, we propose wideband modulation as a promising alternative to eliminate the frequency dependency in downlink communication and increase reliability in multipath channels. Wideband pulses distribute their energy over a large bandwidth. Hence, they are less susceptible to the exact distribution of resonances and spectral nulls in the channel, promising a more robust performance over various channels. While tight regulation on the RF spectrum makes wideband modulation difficult in conventional RFID systems, no limitations exist for acoustic signals in metallic structures.



■ **Figure 8.1:** Many backscatter systems, e.g., WISP, decode the received signal analogously by comparing the current envelope with the average over multiple symbol periods. This approach is only reliable when the symbol duration is long compared to the channel's RMS delay spread τ_{RMS} .



■ **Figure 8.2:** The PMU actively controls the current i_{PMU} to keep V_{in} at $\|V_s\|/2$ (MPPT), thereby inhibiting the produced signal V_e .

Sensing the received signal is impeded when harvesting energy from the same transducer simultaneously. Because of the transducer's large internal impedance z_0 , any current i_{PMU} drawn by the PMU diminishes the measured amplitude (cf. Fig. 8.2). When using active MPPT tracking, the PMU extracts current from C_{in} when V_{in} exceeds the MPPT threshold, prohibiting the signal on the envelope detector from exceeding this threshold. Furthermore, commercially available PMUs require a large input capacitance C_{in} to work effectively—at least $10 \mu\text{F}$ [Dev23]. Together with the observed transducer's output impedances, the input capacitance forms a lowpass filter with a cutoff frequency below 40 Hz, prohibiting short symbol durations.

8.1.2 General Approach

This chapter investigates continuous wave modulation in GW channels using several encodings typically employed in RFID standards and existing acoustic backscatter systems. To assess achievable data rates of these schemes in the GW channel, we record

the channel impulse responses of several physical channels and perform a simulation study with the selected encodings. We then evaluate the proposed modulation with wideband pulses in the same channels, using the same simple envelope detector for demodulation and decoding as typically used for the narrowband signals.

In contrast to prior existing RFID and acoustic backscatter prototypes, we use a software-defined receiver that oversamples the envelope with several times the symbol frequency. We show that this is necessary to achieve low BERs in the presence of ISI. Furthermore, we devise and test an SFO mitigation scheme that uses acknowledgments from the tag to estimate the frequency offset and adapts the subsequent messages accordingly, thereby shifting the burden of estimation and compensation to the reader.

We prototype a transmitter and receiver pair and test the most promising wideband modulation scheme from the simulation in a realistic scenario with an ultra-low-power MCU as receiver. Based on that prototype, we analyze the interference between receiving and harvesting and explore a potential remedy: Deactivating the harvester during the reception of downlink messages. Furthermore, we explore the potential of using wideband carriers for WET and the interference of concurrent harvesting and signal reception.


8.1.3 Structure

In the remainder of the chapter, Section 8.2 first describes the design of the wide- and narrowband modulation schemes. Then, Section 8.3 describes the implementation and setup of our simulation and evaluates the performance of the schemes in the simulation, exploring various parameters. Using the simulation study results, we design a prototype transmitter and receiver that implements the most promising modulation type in Section 8.4. We then evaluate the performance of this prototype in a real-world test in Section 8.5.

8.2 Communication System

This section explains the design choices and assumptions made for the downlink communication. We then characterize the modulation and encoding schemes—selected because of their popularity in RFID and piezoacoustic backscatter systems—and the proposed wideband modulation.

We chose a common message structure and protocol, shown in Fig. 8.3, for all modulation schemes in this chapter. Following a reader-talks-first approach, which is used in the EPCGlobal RFID standards, the reader starts by generating a continuous-wave carrier to charge the tag. To make it easy for the tag to detect the start of a message without continuously correlating with a preamble sequence, the reader initiates a message with a short period T_{off} , where no carrier is transmitted. The tag can easily detect such an initiation pulse running on a low sampling rate to conserve power. The initiation is



■ **Figure 8.3:** The general message structure is valid for all modulation schemes in this chapter. In this example, the message and preamble bits are encoded in OOK.


followed by a 7-symbol Barker-Sequence that serves as preamble to synchronize with the received signal. Then, N_b symbols are transmitted with symbol rate $1/T_m$. The choice of N_b depends on the application and is investigated in Section 8.3. Finally, the reader starts the sinusoidal carrier again, during which the tag can backscatter an acknowledgment.

8.2.1 Modulation

The tag architecture is based on an envelope detector, which has several implications for the modulation schemes. First, only the signal envelope can be detected, prohibiting any phase or frequency modulation. Second, no active gain control is available. Hence, weak signals below the noise power can not be detected, and strong signals may cause saturation of the ADC. Weak signals are not a limitation for passive tags because any signal sufficient to power on the tag will be detectable. However, saturation is a realistic problem that can distort a signal, e.g., if a carrier frequency resonates strongly in the structure. In the following paragraphs, we introduce the selected modulation schemes and discuss their advantages and disadvantages regarding these limitations. We only investigate binary modulation to achieve maximum robustness and keep the receiver simple. The investigated modulation schemes are illustrated in Fig. 8.4.

AM modulates the amplitude of the CW. Several backscatter systems realize OOK using rectangular pulses, where the amplitude of the carrier switches between $A_{c,w}$, and a reduced amplitude A_{low} . The lower level can be adjusted to trade signal SNR for a higher average transmitted power: While the average available signal power is higher with a larger A_{low} , enabling the tag to harvest more energy during message reception, AM is susceptible to saturation if the lower amplitude level exceeds the maximum ADC value at the receiver, the signal is entirely eliminated. Therefore, we set A_{low} to zero to prevent saturation.

Using **AM**, we can choose different encodings. In **OOK**, the amplitude is either low or high throughout the symbol period, depending on the encoded data. The average received signal power in **OOK** is, hence, data-dependent, i.e., when a long streak of zeros is transmitted, no power can be harvested during this period. In Pulse Position Modulation (**PPM**), a low-level pulse is placed in the first or the second half of a symbol



1: Assuming that the low-frequency is entirely attenuated by the channel.

Figure 8.4: Overview of the investigated continuous wave modulation schemes and encodings, their magnitude spectrum, and their average signal power relative to the continuous wave power P_{cw} .

interval, depending on the encoded data. To achieve maximum robustness, we set the duration of the low-level pulse to $T_m/2$. Then, the average signal power is half the power of the continuous wave signal and is independent of the transmitted data. Furthermore, in Pulse-Duration Modulation (PDM), a high-level pulse concludes every symbol period, but the duration of the preceding low-level pulse changes depending on the value of the bit. Average signal power is data-dependent, though it is at least $0.25 P_{cw}$ and at most $0.5 P_{cw}$.

Lastly, some acoustic backscatter systems used FSK in the downlink. At first glance, FSK seems unsuitable in this scenario because the envelope detector can not estimate the frequency of the received signal. However, due to the strong frequency-selectivity of the channel, the different symbols experience a different gain, leading to amplitude variation in the received signal. The authors of [Gon+22] chose frequency f_0 outside the transducer's resonance band. The tag uses an OOK receiver to decode the FSK-modulated message. While this seems counterintuitive, it may perform differently from OOK because the power of an FSK-modulated signal is distributed differently in the frequency domain. The authors report that FSK reduced ISI in their experiments compared to OOK.

In contrast to these narrowband modulation schemes, we propose wideband pulses instead. Various wideband schemes exist, e.g., chirp spread-spectrum modulation or OFDM. However, a short-duration pulse seems most suitable for receiving the pulses with the resource-constrained tag because it can be detected based purely on the signal envelope. The pulse energy should be concentrated within the resonance band of the transducer because the channel gain is diminishingly small outside of this band. Hence,




Figure 8.5: Overview of the investigated wideband modulation schemes and encodings, their magnitude spectrum, and their average signal power.

we define a transmit pulse $p_{wb}(t)$ in this regard as

$$p_{wb}(t) = A_{wb} \cos(2\pi f_c t) \text{sinc}(Bt), \quad (8.1)$$

where A_{wb} is the peak amplitude of the pulse and B is the pulse's bandwidth. This pulse has a frequency spectrum of

$$\mathcal{F}\{p_{wb}\}(f) = \frac{A_{wb}}{2B} \text{rect}\left(\frac{|f| - f_c}{B}\right), \quad (8.2)$$

where the signal energy is distributed equally over the specified bandwidth. The window function $\text{rect}(f)$ equals one for all f between -0.5 and 0.5 and is zero otherwise. For a fair comparison with the continuous wave modulation, we choose A_{wb} so that the wideband pulse has the same total energy as the rectangular continuous wave pulses with amplitude A_{cw} and duration T_m , i.e.,

$$A_{wb} = A_{cw} \sqrt{T_m \cdot B}. \quad (8.3)$$

With the described wideband pulse, we consider two encodings illustrated in Fig. 8.5. Wideband On-/Off-keying (WB-OOK) encodes a one-bit by the presence of a pulse during a symbol period, while a zero-bit is represented by the absence of a pulse. Similarly, in Wideband Pulse Position Modulation (WB-PPM), the pulse is shifted in time to peak at $T_m/4$ or $3T_m/4$ to encode a one-bit and zero-bit, respectively. By choosing A_{wb} according to Eq. (8.3), WB-OOK and WB-PPM have the same average signal power available to the harvester as OOK and PPM.




Figure 8.6: The receiver employs a passive lowpass filter before sampling. The cutoff frequency is inversely proportional to each modulation scheme's shortest pulse duration.

8.2.2 Receiver


In the receiver, the envelope signal V_e passes through a passive first-order lowpass filter with a cutoff frequency f_{co} that is inversely proportional to the shortest pulse duration in the signal (cf Fig. 8.6). We then sample the lowpass-filtered signal at rate N_m/T_m and continue processing the signal in the MCU. The software-defined receiver must be simple to run on the resource-constrained tag device without using high sampling rates or complex filters such as equalizers. We correlate the received digital signal $y[k]$ with the expected preamble for frame synchronization, yielding the start index of the first message sample k_0 . The symbol decision $d[l]$ for the l -th symbol is then made by correlating the symbol's samples with the ideal symbol, which depends on the selected modulation and encoding.

The OOK, WB-OOK, and FSK receiver compares signal mean throughout a symbol period with a threshold value \hat{b}_{th} , which is the average magnitude of the envelope and can be estimated by computing the mean of the received preamble samples. Because of the long delay spread in the channel, the optimal threshold value fluctuates depending on the transmitted data, e.g., it falls after a long streak of zero bits or rises after several consecutive one bits, requiring tracking during the message.

$$d_1[l] = \begin{cases} 0 & \text{if } c > \hat{b}_{th} \\ 1 & \text{else} \end{cases} \quad \text{with } c = \frac{1}{N_m} \sum_{i=0}^{N_m-1} \hat{b}[k_0 + lN_m + i]. \quad (8.4)$$

The PPM and WB-PPM schemes do not require any threshold to make a bit decision because they compare the mean of the first half of the symbol with the mean of the second half, making them independent of the average signal amplitude.

$$d_2[l] = \begin{cases} 0 & \text{if } c_1 - c_2 > 0 \\ 1 & \text{else} \end{cases}, \quad (8.5)$$



■ **Figure 8.7:** Simulation signal flow. The input baseband signal is generated differently depending on the selected modulation and encoding scheme.

with

$$c_1 = \sum_{i=0}^{N_m/2-1} \hat{b}[k_0 + lN_m + i] \quad \text{and} \quad c_2 = \sum_{i=N_m/2}^{N_m-1} \hat{b}[k_0 + lN_m + i].$$

Finally, the PDM scheme considers only the samples that differ between the zero and one-symbol, which we denote here as the set \mathcal{I} . It then also uses a threshold to make the decision.

$$d_3[l] = \begin{cases} 0 & \text{if } c < \hat{b}_{\text{th}} \\ 1 & \text{else} \end{cases} \quad \text{with} \quad c = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \hat{b}[k_0 + lN_m + i]. \quad (8.6)$$

8.3 Simulation Study and Evaluation

We implement a simulation to explore the limitations and influences of various parameters for the downlink. Figure 8.7 depicts the signal flow of the implemented channel simulation starting with the transmitted baseband signal $b(t)$, which is generated from pseudo-random bit sequences depending on the selected modulation and encoding scheme. The first part of the simulation works at the sampling rate of the recorded frequency response, i.e., 4 MHz in our case. We first shift the baseband signal to pass-band by mixing it with a cosine at the carrier frequency. Then, the channel impulse response $h_{\text{pb}}(t)$ is convolved with the passband signal. To model the unsynchronized clocks between reader and tag, the simulation draws a delay from a uniform random distribution in the range from $[-T_m/2, T_m/2]$. The envelope detector is modeled by taking the absolute value of the signal and passing it through a first-order lowpass filter. Finally, the signal is downsampled to the receiver rate, and we add AWGN noise to the simulated envelope.

Channel Sounding

The simulation requires knowledge of the channel's impulse responses $h_{\text{pb}}(t)$, which we record from the test structure by transmitting chirp signals through the channel and recording the response using a combined signal generator and oscilloscope (TiePie Handyscope HS-5). The linear chirps are given by

$$x_{\text{chirp}}(t) = A_0 \sin \left(2\pi \int_0^t \frac{f_{\text{max}} \cdot \tau}{T} d\tau \right), \quad (8.7)$$

where f_{max} is the bandwidth of the chirp, A_0 is the amplitude and T is its duration. The chirp is well suited and commonly used for acoustic channel soundings [SR21] because its frequency response has a near-constant magnitude over the chirp's bandwidth. Furthermore, in contrast to other wideband excitations, with similar magnitude responses, e.g., short-duration pulses, chirps also have a constant amplitude, making them more practical to excite by distributing signal power more evenly over the chirp's period. We used $f_{\text{max}} = 2 \text{ MHz}$ and $T = 100 \text{ ms}$ for the channel sounding. The oscilloscope records the generated chirp and the response $y_{\text{chirp}}(t)$ at the receiving transducer, and the impulse response is given by

$$h_{\text{pb}}(t) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{y_{\text{chirp}}(t)\}}{\mathcal{F}\{x_{\text{chirp}}(t)\}} \right\}. \quad (8.8)$$

In the remainder of this section, we first determine a suitable bandwidth for the wideband pulse and then compare the reliability in terms of the BERs achieved with different modulation schemes and encodings. We then investigate the influence of the number of samples per symbol, the robustness against sampling frequency offsets, and tolerance against noise.

8.3.1 Wideband Pulse Characteristics

We choose the bandwidth of the wideband pulse according to the channel impulse responses. The goal is to reach a tradeoff between wide bandwidth and efficient use of transmitted energy. For the investigated specimen, we discussed and presented the magnitude frequency response of one of the GW channels already in Section 4.2.2, where we observe the largest channel gains between 200 – 250 kHz. To validate this observation and quantify this for all investigated channels, we show the cumulated spectral energy density of the frequency response in Fig. 8.8 for four GW channels on the test structure (cf. Section 4.2.2). Despite the differences between the channels, this frequency band contains over 95 % of the total spectral energy in all channels, which is why we choose a wideband pulse with $B = 50 \text{ kHz}$ and $f_c = 225 \text{ kHz}$.




Figure 8.8: Cumulative spectral energy of the investigated specimen for four channels.

8.3.2 Achievable Data Rates

We used the channel simulation to estimate bit error rates for the investigated modulation schemes at different data rates. For the continuous-wave schemes, we repeated the simulation with 11 different carrier frequencies between 200 – 250 kHz and repeated the procedure on four impulse responses (CH₃→₂, CH₃→₁, CH₄→₁, and CH₅→₁). To determine the maximum achievable data rates, we started simulation for each modulation scheme at a data rate of 3 kbit s⁻¹ and then sequentially reduced data rates in steps of 50 bit s⁻¹ until no bit errors occurred through ten thousand messages with 100 bits each. We then selected the highest data rate that achieved a BER of less than 10⁻⁵ as the maximum achievable data rate for this channel and modulation. The wideband modulation schemes always occupy the same band; therefore, we acquire only a single maximum data rate for each channel impulse response. The results are shown in Fig. 8.9

Examination of the highest symbol rates (with BER below 10⁻⁵) for the narrowband modulation schemes confirms 1) that the throughput varies between different carrier frequencies with more than ten-fold increases between the best and worst carrier frequencies. Furthermore, 2) which carrier frequency permits high rates varies significantly between the different channels, and 3) at the same carrier frequency, some modulation schemes work significantly better than others. However, there is not one superior modulation scheme. The distribution of the achieved narrowband data rates shown in Fig. 8.10 suggests that PPM yields, on average, the most consistent performance, while PDM is, in most cases, the least favorable option. FSK performs very similarly to OOK, with only minor deviations.

The wideband modulation schemes achieve a more reliable data rate in all four channels. In channel CH₃→₂, they exceed the maximum data rate of all narrowband schemes considerably with up to 2500 bit s⁻¹ compared to 1950 bit s⁻¹ with PPM at 245 kHz. In channel CH₃→₁, WB-PPM is nearly identical with the fastest narrowband option at 1100 bit s⁻¹. In channel CH₄→₁, WB-OOK is about 30 % below the maximum narrowband channel, yet still well above the average throughout all carrier frequencies. Finally, in

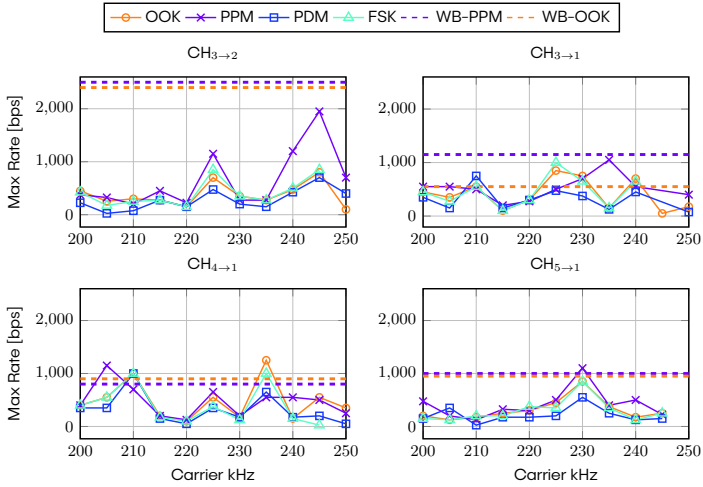


Figure 8.9: Maximum achievable data rates with a BER smaller than 10^{-5} in simulation with additive white Gaussian noise ($\sigma = 10$ mV, roughly corresponding to the observed ADC noise in the MCU). All simulations were performed with $N_{\text{in}} = 10$.





Figure 8.10: The distribution of the achieved data rates with the narrowband modulation schemes at various carrier frequencies and the wideband schemes in multiple channels.



■ **Figure 8.11:** Higher numbers of samples per symbol (SPS) increase the decoding reliability for WB-PPM. However, increasing N_m beyond 10 does not increase data rates significantly.

channel $CH_{5 \rightarrow 1}$, the WB-PPM is 10% below the fastest narrowband scheme. In all cases, even in the 3m channel, which supports the lowest throughput, the observed data rates of WB-PPM are still above 800 bit s^{-1} , with a median of 1075 bit s^{-1} —a 132% increase over the median of the fastest narrowband scheme. Hence, without prior knowledge about the channel, wideband modulation is favorable, and WB-PPM is likely more reliable than WB-OOK. Therefore, the following evaluations will focus on understanding the influence of different parameters for WB-PPM.

8.3.3 Samples Per Symbol

The number of samples per symbol N_m has to be chosen carefully. A low value is desirable in terms of computational complexity, as higher N_m require higher sampling rates and consequently more resources on the tag. However, if chosen too low, the decoding results may be severely impacted, leading to higher BERs.

To quantify the tradeoff between both effects for wideband modulation schemes, we performed the channel simulation with different N_m . We find that $N_m = 10$ is a good tradeoff between reliability and computational complexity. Figure 8.11 shows bit error rates for six, ten, and twenty samples per symbol. When using only six samples instead of ten, the bit error rates are up to two orders of magnitude higher, and the highest achievable data rates are nearly halved. At the same time, even doubling the number of samples does not increase the usable data rates and reduces bit error rates only at data rates where we already achieve BERs below 10^{-5} with ten samples per symbol. For brevity, we only show the graphs for two channels. However, the results for the other channels are very similar.

8.3.4 Sampling Frequency Offset

We already discussed the inaccurate clocks of ultra-low-power microcontrollers in Section 7.4.2, where we explained that the low-power clock has a constant offset from

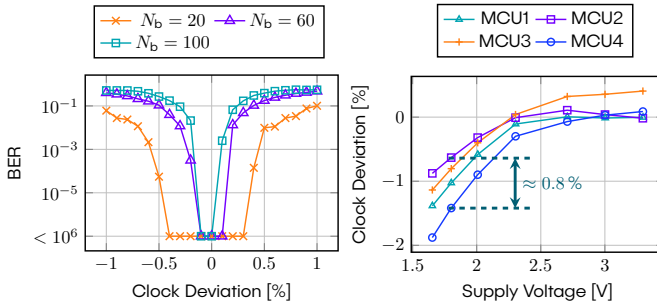


Figure 8.12: SFO causes the BER to rise (left), even in short messages with only 20 symbols. The actual MCU clock frequency compared to its nominal frequency (right) depends on the supply voltage but also deviates significantly between different devices from the same MCU model.

the nominal clock rate, depending on the supply voltage of the MCU, and a time-variant component. Tracking algorithms for SFO correction during message reception are complex, and executing them at the tag should be avoided. We used the simulation study to quantify how much SFO the receiver tolerates without requiring additional measures. Therefore, we resampled the simulated samples with an increasing sample clock deviation before feeding them into the decoder. We repeated the experiment with 20, 60, and 100 bits per message, which influence the susceptibility to SFO: throughout a long message, even a slight frequency offset can add up to a significant delay.

Figure 8.12 (left) reveals that the decoding algorithm is tolerant to a few per mill of SFO, and the BER rises considerably when exceeding this threshold. Short messages of just 20 bits are more resilient but tolerate no more than roughly 0.4% SFO. To estimate if ultra-low-power MCUs can achieve this clock accuracy, we measured the voltage-dependent clock deviation for four MCUs of the same model (STM32L073RZ). Every measurement was repeated ten times to validate consistent behavior. Figure 8.12 (right) clearly shows that the four devices follow a consistent trend, where the actual clock frequency is up to 2% below the nominal frequency when the supply voltage is reduced to the specification minimum of 1.65 V.

The reader could compensate the clock frequency offset if it is systemic, i.e., consistent between all MCUs of the same model. However, the absolute offset varies between the four MCUs. At 1.8 V, the span between the devices is within 0.8%. This clock variation exceeds the tolerances of the decoding algorithm. Hence, dynamic compensation is necessary to achieve robust communication.




Figure 8.13: (a) The receiver is tolerant to signal SNRs below 15 dB in all channels. (b) With the observed noise power ($\sigma_n = 7.3$ mV), the required SNR allows a reduction of the maximum sender amplitude to only 8 V.


8.3.5 Noise Tolerance

This section evaluates the receiver's noise tolerance. We measure BERs while increasing the standard deviation of the additive white Gaussian noise in the simulation. We also compute the SNR for each measurement, where signal power is estimated from the envelope before adding noise. Figure 8.13 (a) shows that the BERs exceed 10^{-5} when the SNR falls below 10 – 15 dB.

To put this result into perspective, we measured the noise level experienced by the tag upon receiving an unmodulated envelope. The observed ADC samples were distributed with a standard deviation of roughly 7.3 mV. To exceed an SNR of 15 dB with this noise distribution, the signal must have a standard deviation of at least 41.1 mV. The SNR requirement is not a limitation for practical wirelessly powered devices, as incoming signals must exceed the required SNR significantly to deliver sufficient power to the harvester. Semi-passive or active devices, however, may facilitate communication over longer distances with small channel gains, where this limitation becomes relevant.

From another perspective, we can utilize the derived SNR limits and the observed noise strength to estimate how much the sender can reduce its TX power without compromising reliability. Figure 8.13 (b) investigates the relationship between the wideband pulse amplitude and the SNR of the signal. A peak amplitude of only 8 V would be sufficient for reliable reception in the investigated channels.

Finally, we have to note that we measured the noise observed by the MCU under ideal conditions on a resting structure. Furthermore, we only considered white Gaussian noise in this evaluation. Various noise sources with non-gaussian distribution may exist in real applications, e.g., structural vibrations through traffic or other loads. Therefore, noise may still cause problems in practical deployments.



■ **Figure 8.14:** The reader utilizes an acknowledgment to estimate the tag's sampling clock offset and adapts its symbol interval accordingly in the next message.

8.4 Prototype Implementation

We extended our MCU-based tag implementation presented in Chapter 7 to validate the simulation results and assess their practical applicability. The results from the simulation study guide the design decisions made during prototype development. WB-PPM was the most promising modulation scheme in simulation, so we tailored the software-defined receiver to this scheme. Furthermore, we tested mitigation techniques for practical problems, such as interdependence between harvesting and receiving and SFO mitigation.

8.4.1 Sample Frequency Offset Mitigation

Section 8.3.4 revealed that the typical SFOs with the selected MCU require some adaption mechanism. However, running complex PLLs on the tag is impractical because of the resource limitation and focus on low power consumption. Therefore, we shift the burden of SFO tracking to the reader, which has much more computational power. We design a communication protocol where the tag transmits a short Acknowledgement (ACK) following every successfully received packet or a non-acknowledgment (NACK) after every incorrect packet. The reader then utilizes this response to detect the tag's clock offset and adapts its symbol interval accordingly when transmitting the next packet.

The symbol interval the tag's sampling is based on, T'_m and the intermediate frequency used for backscattering messages are derived from the same main clock. Hence, the reader can use the intermediate frequency to estimate the tag's SFO. We denote the nominal intermediate frequency f_i and the actual frequency generated at the tag as f'_i . Moreover, we denote the estimate obtained at the reader as \tilde{f}_i . Similarly, the symbol interval estimation denoted \tilde{T}_m . Figure 8.14 sketches the estimation sequence during transmission.

In our prototype, we designed the ACK/NACK packets to enable simple frequency estimation. Therefore, the tag backscatters a pulse with continuous frequency f'_i , and the reader estimates \tilde{f}_i by first sampling the backscattered signal with a sampling rate of $4f_i$ and retrieving the complex baseband samples as described in Section 7.3.3. It

then uses static frequency offset detection (cf. Eq. (2.19)) to recover the average phase shift between consecutive samples $\Delta\phi_i$ as

$$\Delta\phi_i = \sin^{-1} \left(\frac{1}{N_a - 1} \sum_{k=0}^{N_a - 2} \frac{y[k+1] \cdot y^*[k]}{\|y[k+1] \cdot y[k]\|} \right), \quad (8.9)$$

where N_a is the number of samples in the acknowledgment message. With $\Delta\phi_i$, we can derive the estimate of the tag's actual intermediate frequency as

$$\tilde{f}_i = f_i + \frac{\Delta\phi_i}{2\pi} \cdot f_s = f_i + \frac{\Delta\phi_i}{2\pi} \cdot \frac{N_m}{T_m}, \quad (8.10)$$

and a matching symbol interval as

$$\tilde{T}_m = T_m \cdot \frac{f_i}{\tilde{f}_i}. \quad (8.11)$$

We separate between ACK and NACK through the pulse duration, i.e., an ACK is 20 ms long, and a NACK is 30 ms long.

8.4.2 Reader

The wideband pulses are more complex to generate than the continuous sine carrier used in the previous chapters. For our prototype, we generate the transmission signals in software with a desktop computer and then use an off-the-shelf signal generator to convert it into the analog domain. We then used the analog frontend of our prototype reader from Section 5.3.1 to amplify and buffer the signal to drive the transducer. Similarly, we used an oscilloscope to record the demodulated backscatter signal from the reader for SFO estimation. Figure 8.15 depicts the entire reader setup. The oscilloscope and signal generator are coupled in a single device, the Handscope HS-5 from TiePie.

The short pulses require a higher peak amplitude than the sinusoidal carrier to generate a similar average signal power. For example, using Eq. (8.3), for a symbol duration of 1 ms and a bandwidth of 50 kHz, the peak amplitude should be more than seven times higher than the amplitude of a comparable sinusoidal carrier. At the same time, Section 8.3.5 demonstrated that peak amplitudes exceeding 8 V may already facilitate robust communication. We chose to boost the signal voltages with a 4:1 transformer (Coilcraft WA8351-AL [Co23]), yielding a maximum amplitude of 60 V without additional changes to the priorly presented reader hardware. To avoid signal distortion through these adaptations, we verified experimentally that the selected transformer had a flat frequency response in the bandwidth between 200 – 250 kHz.

for 20 ms, or 30 ms respectively, and then switches back into the charging phase by reducing the clock frequency and sampling rate.

Harvester Coupling

Section 8.1.1 discussed that coupling between the harvester and envelope detector is challenging for downlink communication. Due to the transducer's large output impedance, the current the harvester draws reduces the detected envelope. The PMU's power cycle and MPPT mechanism are designed to optimize harvesting, dynamically changing the harvester's input impedance, which leads to the following problems:


1. The large required input capacitance acts as a lowpass filter on the received envelope, leading to considerable distortion. Unfortunately, smaller input capacitors are incompatible with the harvester's voltage-boost circuit.
2. When the MPPT threshold is too low, the harvester extracts charge from the input capacitor at a voltage level below the wideband peak amplitude, clipping the wideband pulses altogether.
3. If the storage capacitor is charged up to the maximum voltage, the harvester switches off to avoid overcharging. It then extracts no current from the source. This sudden switch-off leads to a rapid increase in the envelope, which may saturate the ADC or even damage sensitive electronics. Even without saturation, the rapid change likely induces bit errors if it occurs during a message.

To mitigate these undesired effects, we modified the harvesting circuit as highlighted in Fig. 8.16. The first problem is addressed by placing a diode between the rectifier and the harvester. The diode prevents current from flowing back from the input capacitor to the envelope detector, enabling the detection of a rapid drop of the signal envelope despite the large input capacitor. Additionally, the diode's forward voltage drop also enables the detection of a spike despite C_{in} . However, power dissipation over the diode also reduces the power available for harvesting, compromising the harvester's efficiency.

The second problem, the saturation induced by a low MPPT threshold, is addressed by deactivating the dynamic MPPT sensing mode and operating the MPU with a fixed MPPT threshold. While this reduces harvesting efficiency during the charging phase, it improves the reliability of the downlink communication. Unfortunately, selecting an MPPT threshold well above the average envelope of the downlink signal means that the harvester can not extract significant power during reception and requires a sufficiently large energy buffer to power the MCU throughout the whole downlink message.

Finally, we addressed the third problem with a Schottky diode with 5 V breakdown voltage in parallel to the storage capacitor, which prevents fully charging the storage capacitor and, consequently, prevents the harvester's switch-off during message reception.

The presented measures trade harvesting efficiency against more robust downlink communication. In the long term, we believe that an integrated solution provides a



■ **Figure 8.16:** The harvester setup with a series diode to prevent the lowpass filtering effect of the large input capacitance C_{in} , and a Zener diode to prevent harvester switch off due to full storage.

much better tradeoff, where the MCU has fine-grained control over the PMU states, e.g., ensures that MPPT optimizes harvesting performance during the charging phase but disables during message transmission. However, commercially available PMUs are not optimized for the simultaneous use of the transducer as a harvesting source and sensing element and lack such fine-grained control. Furthermore, existing RFID and piezoacoustic backscatter prototypes do not present solutions to these problems either: As they assume a fixed output impedance of the harvesting source, they do not employ dynamic MPPT matching techniques and implement a harvester with fixed input impedance.

8.5 Real-World Evaluation

This section assesses the potential of wideband modulation for the downlink and investigates practical limitations. We evaluate our prototype in several GW channels on the test structure presented in Section 4.2.2 and compare the maximum achievable data rates in the experiments with simulation results. Furthermore, we characterize harvesting efficiencies and power consumption in different stages of the tag's duty cycle and analyze the average throughput over multiple messages when considering available power.

8.5.1 Achievable Data Rates

We test the maximum data rates our prototype achieves in the GW channels. This evaluation focuses on the robustness of the modulation. Therefore, we eliminate the tag's duty cycle as a potential source of missed packets by connecting the tag to a constant voltage source. This way, the tag is constantly active and ready to receive a packet. Nonetheless, we still run the tag with a strictly limited clock rate to mimic the conditions in the wirelessly powered scenario.




Figure 8.17: Data rates achieved with the receiver implementation on MCU compared to achieved data rates from the simulation. The rates of narrowband schemes were only acquired by simulation.

Figure 8.17 compares the practically achieved data rates with the rates from the simulation. For every measurement point, we transmitted 1000 messages with 96 bits per message. We transmitted fewer bits in total than in the simulation experiment because messages are transmitted in real time in the experiments. As in the simulation, we assessed a data rate achievable if the BER was below 10^{-5} . During all experiments, we applied the SFO estimation technique at the reader and adapted the symbol interval dynamically.

We observed lower achievable data rates in real-world tests. Some of the differences are related to practical limitations: In CH₃→₂, we did not observe any bit errors at a data rate of 1600 bit s⁻¹, but could not increase the data rate because the MCU was not able to process sampling rates higher than 16,000 samples per second. In the other channels, MCU speed was not a limitation, but bit errors occurred already at slightly lower data rates than in the simulation.

A likely explanation for the remaining difference between simulation and real-world tests is the SFO between the reader and the tag: Although we adapt the symbol interval dynamically, the estimates are always noisy, and the actual SFO may vary throughout the message. Furthermore, the diode-based envelope detector is not ideal; therefore, its output is slightly distorted compared to the simulation output. During the simulation, we additionally adapted the cutoff frequency of the lowpass filter according to the symbol rate. In the real-world experiment, we used a fixed cutoff frequency for all data rates to avoid manual resoldering of components for different symbol rates. Hence, reduced performance in the real-world experiment must be expected.

Despite the reduced throughput, the wideband pulse-based modulation scheme still achieves significantly higher data rates than the 80 – 200 bit s⁻¹ that previous studies reported. Comparison with the narrowband modulation schemes is ambiguous. The best carrier frequency facilitates higher rates than WB-PPM in real-world tests in all channels. However, the best frequency is usually unknown, and the median narrowband




Figure 8.18: The accuracy of the SFO detection scheme based on the ACK/NACK messages. The first five milliseconds of the ACK are unsuitable for detection (left). With a low SNR of the backscattered signal, the individual estimates deviate more (right).

data rate of all frequencies is, in all cases, below the wideband rates. Therefore, the wideband scheme provides robust and faster communication without adapting the carrier frequency to each channel. Furthermore, we acquired the maximum data rates of the narrowband schemes only in simulation. We expect them to drop in a real-world implementation similarly to the wideband schemes, supporting the usefulness of WB-PPM.

Sample Frequency Offset

Section 8.3.4 showed that correct decoding of the downlink messages requires an SFO mitigation strategy. The observed SFO of our prototype is in the range of 0.8%—too large for decoding even short, e.g., 20-bit messages. The mitigation technique presented in Section 8.4.1 relies on an estimate of \hat{f}_i derived from the backscattered ACK/NACK packet. This section evaluates the estimation accuracy and robustness to noise and the effect of multipath propagation on the estimate. Therefore, we generated nine thousand ACK packets at 30 carrier frequencies and three different clock offsets. For each acknowledgment, we executed the presented estimation algorithm.

Figure 8.18 (left) shows the instantaneous SFO estimates computed from the phase difference of two consecutive samples. At the start of the acknowledgment, the estimation error is large due to the multipath propagation. However, after about 5 ms, a steady state emerges, and the average estimation error over all recorded acknowledgments converges toward zero. However, the standard deviation of the instantaneous estimations is roughly 200 Hz, corresponding to a clock mismatch of 2.5% at an intermediate frequency of 8 kHz—much more than the receiver can tolerate. Hence, an instantaneous measurement alone is insufficient for SFO estimation.

To reduce the estimation error for a full acknowledgment, we omit the first 5 ms of each packet and average the remaining instantaneous estimates. Figure 8.18 (right) shows the standard deviations of estimates derived from all observed acknowledgments. We observe that estimates become less accurate when the SNR decreases. However, with SNRs down to 7 dB, most of the estimates have a standard deviation below 2 Hz, corresponding to a clock mismatch of 0.025 %, which is negligible for the receiver. For clarity, the plot does not show one outlier at a carrier frequency where the SNR was only 3 dB. In this case, the standard deviation of the estimates rose to roughly 155 Hz. In such cases, additional measures will be necessary to enhance the accuracy, e.g., averaging the estimates of multiple ACK packets. However, such low SNRs are rarely observed, especially when fine-tuning the carrier frequency for a particular channel. Hence, in most cases, the frequency estimation is accurate enough to mitigate the sampling frequency offset at the reader.

8.5.2 Wideband Power Transfer

To evaluate how efficiently the PMU harvests energy from wideband pulses, we compared the transfer efficiencies with CW and wideband signals. The transmitted power was derived from the measured current and voltage of the generated signal. We measured the harvested power by placing a 3.3 k Ω resistance at the 1.8 V regulated DC output of the harvester—yielding an output power of roughly 1 mW while the harvester output is activated. We then obtain the average harvested power by measuring the duty cycle. Hence, the obtained efficiency incorporates all losses from the channel, rectifier, PMU, and leakages. Due to the maximum power transfer theorem, the harvester can, at most, harvest half of the transmitted power (see Section 2.4.1).

In Fig. 8.19 (left), we first measure the transfer efficiency in an ideally flat channel. Therefore, we placed an 1 k Ω resistor between reader and tag to emulate a source with fixed output impedance. The power transfer achieved an efficiency of roughly -7 dB. At the same time, the efficiency using wideband pulses was only about -12.7 dB—only a quarter of the CW efficiency. Hence, as the idealized channel is frequency-agnostic, the selected harvester is better suited for CW than WB signals.

When repeating the same experiment on a real GW channel, the carrier frequency is crucial for efficient power transfer. We first chose two frequencies that had a high voltage gain on the specimen, 219.54 kHz and 222.74 kHz, and compared them with wideband pulses. In the real channel, all efficiencies are lower than in the flat channel, with roughly -15 dB for the best CW frequency. The difference between the wideband pulses and the best CW frequencies increased to about -7.5 dB. However, the WB pulses perform better than CW signals with medium-gain frequencies, e.g., 220.56 kHz. Ultimately, when choosing a frequency with a particularly low voltage gain, the power was insufficient to charge the harvester at all.

The takeaway from these results is that we should choose a suitable frequency CW signal for efficient power transfer. Furthermore, we can expect to harvest about 5 to 6

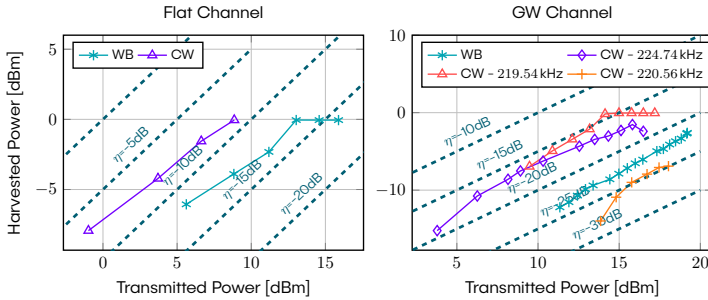



Figure 8.19: Power Transfer Efficiency with wide- and narrowband pulses in an ideal flat channel (left) and in a real GW channel, where the carrier frequency is important (right).

times less power during downlink communication with wideband-modulated messages. However, wideband power transfer may still be useful during tag discovery, when the reader does not yet know suitable frequencies for each tag.

8.5.3 Power Consumption and Availability

In the previous sections, we considered the maximum achievable data rates at which we can still decode the messages correctly. However, the data rates in productive batteryless systems are also limited by available power: The tag must have enough stored energy to receive a whole message, and between messages, it requires time to charge up before the next message starts. However, to estimate these maximum average data rates, we must first acquire the MCU's power consumption.

We conducted a low-side current measurement by placing a $10\ \Omega$ shunt resistor between the MCU's ground pin and the global ground and supplied the MCU with 1.8 V. Figure 8.20 shows the power consumption in four states. Message reception is the most demanding stage with up to 1.2 mW power consumption, followed by acknowledgment generation. Between messages, during the charging phases, the MCU is actively trying to detect the start of the following message. Here, reducing the clock frequency during the packet detection phase saves about 57% power compared to the active reception mode. While this is a considerable reduction, an analog message detection mechanism without active sampling, similar to WISP, would be beneficial. Then, the MCU could remain in the STOP mode with only $5\ \mu\text{W}$ power consumption while waiting for an external pin interrupt.



■ **Figure 8.20:** Power Consumption of the MCU during four states while receiving a packet with $T_m = 1$ ms and $N_b = 96$. Each boxplot shows the distribution of power measurements during the state.

Maximum Rates

To explore the effectiveness of optimizing the detection stage, we estimate the maximum achievable data rates when taking the charging and ACK phases into account. To estimate the maximum, we assume that the reader transmits a message as soon as the tag has charged up sufficiently to receive it. These rates serve as an upper bound and are not achievable in real applications because the reader does not know the tag's state of charge and will—most of the time—either transmit too early, risking packet loss, or too late, not utilizing the maximum potential rate.

In our model, we denote the durations the tag spends in the states *charging and detection*, *reception*, and *acknowledgment* as t_c , t_r , and t_a . Similarly, P_c^c , P_r^c , and P_a^c are the consumed powers in each state, and P_c^h , P_r^h , and P_a^h are the harvested powers per state. During continuous maximum-rate communication, the tag must harvest at least as much energy as it consumes throughout message reception. Hence,


$$t_c(P_c^h - P_c^c) + t_r(P_r^h - P_r^c) + t_a(P_a^h - P_a^c) \geq 0. \quad (8.12)$$

This equation yields the minimum charge time between messages

$$t_c \geq \frac{-t_r(P_r^h - P_r^c) - t_a(P_a^h - P_a^c)}{P_a^h - P_a^c}. \quad (8.13)$$

The consumed powers have been estimated in Fig. 8.20. Furthermore, we assume that $P_a^h = 0.5P_c^h$ because the transducer is short-circuited half the time while transmitting the acknowledgment. From Section 8.5.2, we assume $P_r^h \approx 0.2P_c^h$. Then, we derive the maximum rate as

$$R = \frac{N_b}{t_c + t_r + t_a} \quad \text{with} \quad t_r = T_{\text{off}} + (7 + N_b)T_m, \quad (8.14)$$



■ **Figure 8.21:** The maximum data rate of a batteryless tag when using frame synchronization with full MCU speed, reduced MCU speed, and analogously with the MCU in stop-mode. The messages were $N_b = 96$ bits long.

with N_b as the number of payload bits in the message, and a preamble length of 7 symbols. The acknowledgements are 20 ms long.

Figure 8.21 shows the maximum rates for the different frame synchronization implementations. Two methods rely on the MCU to sample and detect an initialization pulse, once with full MCU speed and once with reduced speed, as implemented in our prototype. The third method assumes an analog detector, so that the MCU can remain in the power-conserving stop mode during the charging phase. The model shows that reducing the power consumption in this stage is worthwhile as a more efficient detection reduces the minimum harvested power required to charge the tag and enables higher data rates when the harvested power is below roughly 4 mW. With higher harvested power, the tag harvests more than it consumes in all states, hence, nothing is gained by reducing power consumption further.

8.6 Summary and Discussion

This chapter analyzed the downlink communication in GW channels. The conventional techniques of implementing a receiver on the resource-constrained tag are infeasible in GW channels because of the large delay spread of the channel. We investigated wideband pulses to increase the robustness against multipath propagation as an alternative to commonly used continuous wave modulation schemes. Therefore, we implemented a channel simulation including impulse responses of four real downlink channels, incorporating real-world effects such as noise and randomized symbol timing. Using this simulation, we compared the maximum achievable data rates of four different encodings using CW modulation with two encodings based on wideband pulses.

The simulation revealed that the achievable data rates with CW schemes depend strongly on the chosen carrier frequency. We identified Pulse Position Modulation as the most robust scheme both with CW and WB pulses. In two out of four channels,

communication with WB pulses enabled higher data rates than any CW scheme at any frequency. In the remaining two channels, some frequencies existed where CW was better at at least one frequency. However, in all cases, WB modulation was faster than CW on average over all carrier frequencies.

Inaccurate clocks also pose a challenge for the downlink. The simulation study revealed that the receiver only tolerated a clock derivation of less than 0.4% for short messages and considerably less when messages become longer. However, measurements on several MCUs yielded an expected clock offset of more than 0.8%. To reduce tag complexity, we developed a scheme where the reader estimates the tag's frequency offset from backscattered acknowledgments and subsequently adapts the symbol interval according to the estimated clock offset. This scheme successfully enabled error-free transmission with high clock deviations without manual tuning for backscattered signal SNRs down to 7 dB.

We demonstrated that active MPPT during energy harvesting interferes with message reception, as the extracted current directly influences the transducer voltages. Therefore, we proposed a modification of the harvester circuit that still enables message reception at the expense of harvesting efficiency. In the future, we propose investigating more integrated solutions where the MCU has fine-grained control over the PMU states, i.e., optimizing for efficiency during a charging phase and for signal integrity during the reception.

Furthermore, we experimentally investigated if wideband pulses are a viable alternative for wireless power transfer. We found that, even with an idealized channel, the transfer efficiency of wideband pulses is about 5 dB lower than that of CW signals. In real channels, the efficiency of CW signals depends strongly on the carrier frequency. WB pulses were about 7.5 dB worse for power transfer than the best-frequency CW signals. However, WB pulses are beneficial when the channel is yet unknown, i.e., during tag discovery.

Implementing better message detection schemes potentially improves tag availability and increases data rates when harvested power is scarce. We show that an analog comparator-based detector can reduce the power consumption during the charging phase by roughly 93% compared to the software-defined detection, which allows the MCU to reduce charging times significantly.

Conclusion and Outlook

Structural health monitoring becomes increasingly important to efficiently maintain aging and new infrastructure, such as bridges, pipelines, wind turbines, and similar structures. Unfortunately, the vast number of sensors required for continuous and real-time monitoring poses a challenge, as deployment and maintenance of the devices become very costly. Although batteryless sensor nodes using backscatter communication, e. g., RFID, may reduce the costs of such sensor networks immensely, some environments are unsuitable for radio-frequency wave-based communication. Acoustic backscatter communication and power transfer are promising opportunities to realize economically scalable wireless sensor networks embedded in hard-to-reach and electromagnetically shielded places.

This dissertation first analyzed the state-of-the-art in acoustic through-metal communication. Backscatter communication has never been tested in guided-wave-based channels. Furthermore, research in related active communication and power transfer in such channels is scarce and very heterogeneous. In particular, existing studies tailored systems to a single sensor position in a specific structure, which is not scalable to large numbers. Lastly, most of these studies neglected challenges arising in practical implementations, e. g., limited computational resources, clock mismatches between different devices, and power consumption.

Our first goal was the feasibility analysis of backscatter communication in GW channels and its limitations. Therefore, we started by analyzing the techniques used by RFID and existing acoustic backscatter systems in underwater or in-body applications. The investigation of multiple guided-wave metal channels showed that backscatter communication is possible but comes with several challenges. First, the multipath propagation in guided-wave channels prohibitively limits the achievable data rates. Second, the input impedance of the piezoelectric transducers is strongly coupled with the structure's acoustic impedance, making it impossible to perform static impedance matching, and third, efficient wave propagation in these channels is limited to low frequency ranges, reducing the available bandwidth. To increase the low achievable data rates despite the limited bandwidth, we investigated higher-order modulation, improving bandwidth

efficiency. The strong coupling of the transducer's input impedance with the structure makes selecting ideal load impedances for the backscattering tag before deployment infeasible. We, therefore, proposed and implemented a method that dynamically selects fitting load impedances for the nonlinear channel and a setup procedure based on an equivalent circuit channel model. A thorough evaluation showed that the channel model agrees well with the actual channel.

As inter-symbol interference from the multipath propagation is a major limitation for further data-rate increases, we investigated equalizers for guided-wave channels. The nonlinear nature prohibits the direct application of standard equalization techniques. Therefore, we derived a linearized channel model that the receiver can learn from a training sequence. Using this model, we compared two common equalizer types, showing that decision feedback equalizers are better suited due to their robustness to noise and reduced number of taps. Equalizers were very beneficial in enhancing the data rates, yielding rates up to 10 kbit s^{-1} .

Our second goal was to analyze and improve the downlink, which uses active communication but experiences the same reverberation as the uplink. However, downlink communication requires a different approach because the resource-constrained receiver can not execute complex equalizers. We investigated short-time wideband pulses as an alternative to continuous-wave modulation, which improved data rates and robustness without increasing receiver complexity. Evaluation in simulation and with a practical prototype demonstrated that wideband modulation is more robust against inter-symbol interference than narrowband modulation for most carrier frequencies. Although the narrowband modulation may be better at optimally chosen carrier frequencies, wideband modulation is beneficial without prior knowledge of the channel. However, the long message duration in through-metal communication is challenging for batteryless nodes, as little energy can be harvested during reception, requiring large energy buffers and limiting the availability of a tag device.

Our third goal was to characterize the potential of wireless power transfer and the tradeoffs arising from simultaneous energy harvesting and communication at the tag. We found that unpredictable transducer impedances also interfere with efficient wireless energy transfer, which relies on impedance matching. Instead of tailoring our harvester to the specific channel, we evaluated dynamic maximum power point tracking in the harvester. The guided-wave channel provides opportunities to transfer energy with 2.8% efficiency over three meters, which is acceptable to power ultra-low-power sensor nodes. Experiments on more realistic structures suggested that distances of tens of meters are feasible because of the low attenuation of guided waves, although wave excitation might require a higher TX power initially in thick structures. While simultaneous harvesting did not significantly interfere with backscattering at the tag, the downlink messages proved to be highly distorted, requiring deactivation of the MPPT mechanism during message reception.

Throughout all communication methods in this dissertation, we could show that inaccurate oscillators arising in practical implementations significantly impact robustness,

which most existing studies neglected in their lab and simulation setups. To fulfill our fourth goal, we proposed static and dynamic adaption techniques to alleviate the effects of sampling clock mismatches. However, simulations have shown that, although the implemented methods considerably mitigated the effects of frequency offsets, a clock mismatch still strongly impacted receiver robustness. Furthermore, we demonstrated that the metal channel is time-variant and is sensitive to slight temperature changes and environmental conditions. Hence, we implemented adaption techniques that include the channel parameters of the proposed channel model. To ensure the generality of the proposed methods, we tested them on multiple physical channels in each evaluation, avoiding optimizing our prototype to only one specific channel.

Finally, our last goal was to construct and share a prototype from commercially available, cheap, and low-power components. All the proposed methods were first tested in simulation or on high-performance lab devices. We then designed custom prototypes to validate that the techniques are feasible with the resource-constrained devices that must be expected in realistic applications. To support our goal of fostering additional research in the field, we published the circuit diagrams, PCB layouts, firmware, and software of our prototype in [OR23] under an open-source license and plan to publish later prototypes similarly with this dissertation.

9.1 Outlook

As through-metal communication received little attention in the past—particularly in guided-wave channels—promising research paths emerge from this dissertation.

The encouraging results of wideband modulation for the downlink, i.e., improved robustness to ISI and less dependence on a suitable carrier frequency, motivate the application of wideband modulation schemes in the uplink as well. While many wideband schemes exist, the challenge is to modulate the signal with a microcontroller-based and backscattering tag. Frequency hopping spread spectrum (FHSS) is a likely candidate because it is easy to implement with a tag prototype we presented in this dissertation. However, the backscatter modulation prevents the tag from using a band-limited pulse shape, making it challenging to isolate the sub-channels used in FHSS from each other.

The energy demand of the batteryless tag is a limiting factor for availability and increased distances. Therefore, more energy-efficient tag implementations have a vast potential to improve distances and reduce overall energy consumption. As demonstrated with WISP and SOC-WISP, dedicated implementations on ASICs can potentially reduce energy consumption more than twenty times. Considering the linear decay of guided-wave power with distance, that translates to a twenty-times increase in charging distance at the same transmitter power. Furthermore, large-scale monitoring of civil infrastructure requires large amounts of devices, and a SoC implementation can cut the device cost and size considerably.

In this dissertation, simultaneous harvesting and communication were shown to impede each other because the employed power-management units were not designed for concurrent use of the power source as a sensing element. An integrated approach where the MCU has fine-grained control over the PMU, e.g., deactivating MPPT during message reception, can help to strike a balance between efficient harvesting and robust communication. Furthermore, a tag that is aware of its energy reserves can further reduce its energy consumption, e.g., by completely shutting down when the stored energy is insufficient to perform its task.

Finally, this dissertation focused on the physical layer of the backscatter communication, assuming that only one tag is available for a reader. In practical scenarios, multiple tags may be within the read range, requiring a data-link layer protocol. The diverse RFID standards already propose several approaches to tag discovery and addressing. However, they are based on application scenarios different from the envisioned monitoring tasks: In through-metal scenarios, the readers and tags are fixed to one location, data rates are much lower, and the uplink is considerably faster than the downlink because of the multipath propagation. The design of the existing protocols may, hence, be unfit for this purpose, and new or adapted protocols may prove beneficial.

Bibliography

- [23a] *WISP Related Publications*. 2023. url: <https://sites.google.com/uw.edu/wisp-wiki/community/publications> (visited on 11/08/2023).
- [23b] *Zustand der Brücken an Bundesfernstraßen in Deutschland im Jahr 2023*. 2023. url: <https://de.statista.com/statistik/daten/studie/166122/umfrage/zustand-der-strassenbruecken-in-deutschland/> (visited on 11/08/2023).
- [ACG21] Jonathan Akafua, Ryan Chapman, and Hongzhi Guo. "A Design of Wireless Communication and Wireless Energy Transfer System for In-Pipe Robots". In: *2021 IEEE International Conference on Wireless for Space and Extreme Environments (WISEE)*. 2021. doi: [10.1109/WiSEE50203.2021.9613833](https://doi.org/10.1109/WiSEE50203.2021.9613833).
- [Afz+20] Sayed Saad Afzal et al. "Enabling Higher-Order Modulation for Underwater Backscatter Communication". In: *Global Oceans 2020: Singapore – U.S. Gulf Coast*. Oct. 2020. doi: [10.1109/IEECONF38699.2020.9389061](https://doi.org/10.1109/IEECONF38699.2020.9389061).
- [Afz+22] Sayed Saad Afzal et al. "Battery-free wireless imaging of underwater environments". In: *Nature Communications* 13.1 (2022). doi: [10.1038/s41467-022-33223-x](https://doi.org/10.1038/s41467-022-33223-x).
- [Alg+19] J. M. Algueta-Migue et al. "Low-Power Ultrasonic Front-End for Cargo Container Monitoring". In: *IEEE Transactions on Instrumentation and Measurement* (2019). doi: [10.1109/TIM.2019.2939704](https://doi.org/10.1109/TIM.2019.2939704).
- [All+23] Ahmed Allam et al. "System-Level DC-to-DC Analysis and Experiments of Ultrasonic Power Transfer Through Metallic Barriers". In: *IEEE/ASME Transactions on Mechatronics* 28.1 (2023). doi: [10.1109/TMECH.2022.3195973](https://doi.org/10.1109/TMECH.2022.3195973).
- [All21] Ahmed Allam. "Acoustic Power Transfer Leveraging Piezoelectricity and Metamaterials". Ph.D. Thesis, Publisher: Georgia Institute of Technology. PhD thesis. July 29, 2021. url: <http://hdl.handle.net/1853/65095> (visited on 05/18/2022).
- [AMW12] Daniel Arnitz, Ulrich Muehlmann, and Klaus Witrisal. "Characterization and Modeling of UHF RFID Channels for Ranging and Localization". In: *IEEE Transactions on Antennas and Propagation* 60.5 (May 2012). doi: [10.1109/TAP.2012.2189705](https://doi.org/10.1109/TAP.2012.2189705).
- [Ash+13] Jonathan D Ashdown et al. "A full-duplex ultrasonic through-wall communication and power delivery system". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 60.3 (2013).
- [ASS16] M. N. Abdallah, T. K. Sarkar, and M. Salazar-Palma. "Maximum power transfer versus efficiency". In: *2016 IEEE International Symposium on Antennas and Propagation (APSURS)*. IEEE, 2016. doi: [10.1109/APS.2016.7695800](https://doi.org/10.1109/APS.2016.7695800).

Bibliography

- [Bah+22] Rudy Bahouth et al. "Lamb Wave Wireless Communication Through Healthy and Damaged Channels With Symmetrical and Asymmetrical Steps and Notches". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 69.7 (2022). doi: [10.1109/TUFFC.2022.3171729](https://doi.org/10.1109/TUFFC.2022.3171729).
- [Bie+21] Franklin Bien et al. "Metal Surface Guided-Wireless Power Transfer System for Portable Applications With Multiple Receivers". In: *IEEE Access* 9 (2021). doi: [10.1109/ACCESS.2021.3076507](https://doi.org/10.1109/ACCESS.2021.3076507).
- [Bru06] Michel Bruneau. "Equations of Motion in Non-Dissipative Fluid". In: *Fundamentals of Acoustics*. John Wiley & Sons, Ltd, 2006. Chap. 1. doi: [10.1002/9780470612439.ch1](https://doi.org/10.1002/9780470612439.ch1).
- [Cha+15] Sourmya Chakraborty et al. "Low-power, low-rate ultrasonic communications system transmitting axially along a cylindrical pipe using transverse waves". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 62.10 (2015).
- [Coi23] Coilcraft. *Ultrasonic Sensing Transformer WA8351-AL*. 2023. url: <https://www.coilcraft.com/getmedia/3a8a29e3-84c9-497d-9acd-96ca699b92c2/wa8351-AL.pdf> (visited on 09/14/2023).
- [CSR23] Filippo Campagnaro, Fabian Steinmetz, and Bernd-Christian Renner. "Survey on low-cost underwater sensor networks: from niche applications to everyday use". In: *Journal of Marine Science and Engineering* 11.1 (2023).
- [Cun+16a] Michael T Cunningham et al. "Low-rate ultrasonic communications and power delivery for sensor applications". In: *MILCOM 2016-2016 IEEE Military Communications Conference*. IEEE, 2016. doi: [10.1109/MILCOM.2016.7795307](https://doi.org/10.1109/MILCOM.2016.7795307).
- [Cun+16b] Michael T. Cunningham et al. "Low-rate ultrasonic communications and power delivery for sensor applications". In: *MILCOM 2016 - 2016 IEEE Military Communications Conference*. 2016. doi: [10.1109/MILCOM.2016.7795307](https://doi.org/10.1109/MILCOM.2016.7795307).
- [Cur+07] Jari-Pascal Curty et al. "RFID Tag design". In: *Design and Optimization of Passive UHF RFID Systems*. Springer, 2007. doi: [10.1007/978-0-387-44710-0_7](https://doi.org/10.1007/978-0-387-44710-0_7).
- [Dev23] Analog Devices. *ADP5091/ADP5092 Datasheet*. 2023. url: <https://www.analog.com/media/en/technical-documentation/data-sheets/ADP5091-5092.pdf> (visited on 09/04/2023).
- [Dig22] Bundesministerium für Digitales und Verkehr. *Verkehrsinvestitionsbericht für das Berichtsjahr 2020*. 2022. url: <https://bmdv.bund.de/goto?id=513572> (visited on 11/03/2023).
- [Din+20] Lianghui Ding et al. "Modeling and Evaluation of Piezoelectric Transducer (PZT)-Based Through-Metal Energy and Data Transfer". In: *MDPI Sensors* 20.11 (2020). doi: [10.3390/s20113304](https://doi.org/10.3390/s20113304).
- [Dob08] Daniel M. Dobkin. "Chapter 8 - UHF RFID Protocols". In: *The RF in RFID*. Newnes, 2008. doi: [10.1016/B978-075068209-1.50008-5](https://doi.org/10.1016/B978-075068209-1.50008-5).
- [Ens16] Joshua Ensworth. "Ultra-low-power Bluetooth Low Energy (BLE) compatible backscatter communication and energy harvesting for battery-free wearable devices". Ph.D. Thesis. 2016. url: <https://digital.lib.washington.edu/443/researchworks/handle/1773/38120> (visited on 09/28/2023).

- [EPC18] EPCglobal. *EPC Radio-Frequency Identity Protocols Generation-2 UHF RFID Standard*. 2018. url: <https://www.gs1.org/standards/rfid/uhf-air-interface-protocol> (visited on 08/05/2023).
- [Faw+14] Abdelfattah Fawky et al. "UWB chipless RFID system performance based on real world 3D-deterministic channel model and ZF equalization". In: *The 8th European Conference on Antennas and Propagation (EuCAP 2014)*. 2014. doi: [10.1109/EuCAP.2014.6902135](https://doi.org/10.1109/EuCAP.2014.6902135).
- [Gha+20] Reza Ghaffarivardavagh et al. "Ultra-Wideband Underwater Backscatter via Piezoelectric Metamaterials". In: *SIGCOMM '20: Proceedings of the ACM Special Interest Group on Data Communication*. ACM, 2020. doi: [10.1145/3387514.3405898](https://doi.org/10.1145/3387514.3405898).
- [GM20] Mohammad Meraj Ghanbari and Rikky Muller. "Optimizing Volumetric Efficiency and Backscatter Communication in Biosensing Ultrasonic Implants". In: *IEEE Transactions on Biomedical Circuits and Systems* 14.6 (2020). doi: [10.1109/TBCAS.2020.3033488](https://doi.org/10.1109/TBCAS.2020.3033488).
- [GNS11] D. J. Graham, J. A. Neasham, and B. S. Sharif. "Investigation of Methods for Data Communication and Power Delivery Through Metals". In: *IEEE Transactions on Industrial Electronics* 58.10 (2011). doi: [10.1109/TIE.2010.2103535](https://doi.org/10.1109/TIE.2010.2103535).
- [Gol05] Andrea Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.
- [Gon+22] Zheng Gong et al. "Empowering smart buildings with self-sensing concrete for structural health monitoring". In: *SIGCOMM '22: Proceedings of the ACM SIGCOMM 2022 Conference*. ACM, 2022. doi: [10.1145/3544216.3544270](https://doi.org/10.1145/3544216.3544270).
- [Gui+19] Raffaele Guida et al. "U-Verse: a miniaturized platform for end-to-end closed-loop implantable internet of medical things systems". In: *SenSys '19: Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. ACM, 2019. doi: [10.1145/3356250.3360026](https://doi.org/10.1145/3356250.3360026).
- [Hei+19] Alexander Heifetz et al. *Final Report for Transmission of Information by Acoustic Communication along Metal Pathways in Nuclear Facilities*. Sept. 2019. doi: [10.2172/1573242](https://doi.org/10.2172/1573242). url: <https://www.osti.gov/biblio/1573242> (visited on 11/08/2023).
- [Hei+21] A. Heifetz et al. "Transmission of Images on High-Temperature Nuclear-Grade Metallic Pipe with Ultrasonic Elastic Waves". In: *Nuclear Technology* 207.4 (2021). doi: [10.1080/00295450.2020.1782626](https://doi.org/10.1080/00295450.2020.1782626).
- [Hel+19] Michael Helmling et al. *Database of Channel Codes and ML Simulation Results*. 2019. url: www.uni-kl.de/channel-codes (visited on 10/07/2023).
- [Her+19] Cristian Herrojo et al. "Chipless-RFID: A Review and Recent Developments". In: *MDPI Sensors* 19.15 (2019). doi: [10.3390/s19153385](https://doi.org/10.3390/s19153385).
- [Hos+10] Thomas Hosman et al. "Multi-tone FSK for ultrasonic communication". In: *2010 IEEE Instrumentation Measurement Technology Conference Proceedings*. 2010. doi: [10.1109/IMTC.2010.5488066](https://doi.org/10.1109/IMTC.2010.5488066).
- [HS22] Xin Huang and Jafar Saniie. "Ultrasonic Video Transmission through Solid Metallic Channel". In: *2022 IEEE International Ultrasonics Symposium (IUS)*. 2022. doi: [10.1109/IUS54386.2022.9957154](https://doi.org/10.1109/IUS54386.2022.9957154).
- [Hua+18] X. Huang et al. "Applying EMAT for Ultrasonic Communication Through Steel Plates and Pipes". In: *2018 IEEE International Conference on Electro/Information Technology (EIT)*. 2018. doi: [10.1109/EIT.2018.8500148](https://doi.org/10.1109/EIT.2018.8500148).

Bibliography

- [IAS22] Gerald K Ijearu, Kenneth Li-Minn Ang, and Jasmine KP Seng. "Wireless power transfer and energy harvesting in distributed sensor networks: Survey, opportunities, and challenges". In: *International Journal of Distributed Sensor Networks* 18.3 (2022). doi: [10.1177/15501477211067740](https://doi.org/10.1177/15501477211067740).
- [JA19] Junsu Jang and Fadel Adib. "Underwater Backscatter Networking". In: *SIGCOMM 19: Proceedings of the ACM Special Interest Group on Data Communication*. ACM, 2019. doi: [10.1145/3341302.3342091](https://doi.org/10.1145/3341302.3342091).
- [Jia+20] Pengcheng Jiao et al. "Piezoelectric Sensing Techniques in Structural Health Monitoring: A State-of-the-Art Review". In: *MDPI Sensors* 20.13 (2020). doi: [10.3390/s20133730](https://doi.org/10.3390/s20133730).
- [Joh03] D.H. Johnson. "Origins of the equivalent circuit concept: the voltage-source equivalent". In: *Proceedings of the IEEE* 91.4 (2003). doi: [10.1109/JPR0C.2003.811716](https://doi.org/10.1109/JPR0C.2003.811716).
- [Joh09] Sarah J. Johnson. "Low-density parity-check codes". In: *Iterative Error Correction: Turbo, Low-Density Parity-Check and Repeat-Accumulate Codes*. Cambridge University Press, 2009. doi: [10.1017/CB09780511809354.003](https://doi.org/10.1017/CB09780511809354.003).
- [JYZ13] Yuanwei Jin, Yujie Ying, and Deshuang Zhao. "Data Communications Using Guided Elastic Waves by Time Reversal Pulse Position Modulation: Experimental Study". In: *MDPI Sensors* 13.7 (2013). doi: [10.3390/s130708352](https://doi.org/10.3390/s130708352).
- [JZY11] Yuanwei Jin, Deshuang Zhao, and Yujie Ying. "Time reversal data communications on pipes using guided elastic waves: Part I. Basic principles". In: *Health Monitoring of Structural and Biological Systems 2011*. Ed. by Tribikram Kundu. Vol. 7984. SPIE, 2011. doi: [10.1117/12.880271](https://doi.org/10.1117/12.880271).
- [Kiz+17] M. E. Kiziroglou et al. "Acoustic power delivery to pipeline monitoring wireless sensors". In: *Ultrasonics* 77 (2017). doi: [10.1016/j.ultras.2017.01.017](https://doi.org/10.1016/j.ultras.2017.01.017).
- [KM21] Nabil Khalid, Rashid Mirzavand, and Ashwin K. Iyer. "A Survey on Battery-Less RFID-Based Wireless Sensors". In: *Micromachines* 12.7 (July 2021). doi: [10.3390/mi12070819](https://doi.org/10.3390/mi12070819).
- [Kur14] Aleksander Kural. "Wireless Energy Supply to Aircraft Structural Health Monitoring Nodes Using Ultrasonic Lamb Waves". In: *Fatigue of Aircraft Structures* 2014.6 (2014). doi: [10.1515/fas-2014-0002](https://doi.org/10.1515/fas-2014-0002).
- [Lam+17] Rolf Lammering et al. *Lamb-wave based structural health monitoring in polymer composites*. Springer, 2017. doi: [10.1007/978-3-319-49715-0](https://doi.org/10.1007/978-3-319-49715-0).
- [LH10] Michael Lunglmayr and Mario Huemer. "Least Squares Equalization for RFID". In: *Second International Workshop on Near Field Communication*. Apr. 2010. doi: [10.1109/NFC.2010.19](https://doi.org/10.1109/NFC.2010.19).
- [LP17] Fuyun Ling and John Proakis. "Carrier Synchronization". In: *Synchronization in Digital Communication Systems*. Cambridge University Press, 2017. doi: [10.1017/9781316335444.007](https://doi.org/10.1017/9781316335444.007).
- [Mäl+19] Moritz Mälzer et al. "Combined Inspection and Data Communication Network for Lamb-Wave Structural Health Monitoring". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 66.10 (2019). doi: [10.1109/TUFFC.2019.2925864](https://doi.org/10.1109/TUFFC.2019.2925864).
- [MG13] Rito Mijarez and Patrick Gaydecki. "Automatic guided wave PPM communication system for potential SHM of flooding members in sub-sea oilrigs". In: *Smart Materials and Structures* 22.5 (2013). doi: [10.1088/0964-1726/22/5/055031](https://doi.org/10.1088/0964-1726/22/5/055031).

- [Mur+19] Florian Muralter et al. "A theoretical and experimental study of passive computational RFID tags". In: *4th International Conference on Smart and Sustainable Technologies (SpliTech)*. 2019. doi: [10.23919/SpliTech.2019.8783177](https://doi.org/10.23919/SpliTech.2019.8783177).
- [Ng+22] Len Luet Ng et al. "Power Consumption in CMOS Circuits". In: *Electromagnetic Field in Advancing Science and Technology*. Ed. by Hai-Zhi Song, Kim Ho Yeap, and Magdalene Wan Ching Goh. IntechOpen, 2022. Chap. 5. doi: [10.5772/intechopen.105717](https://doi.org/10.5772/intechopen.105717).
- [OM88] M. Oerder and H. Meyr. "Digital filter and square timing recovery". In: *IEEE Transactions on Communications* 36.5 (1988). doi: [10.1109/26.1476](https://doi.org/10.1109/26.1476).
- [OR23] Peter Oppermann and Bernd-Christian Renner. "Acoustic Backscatter Communication and Power Transfer for Batteryless Wireless Sensors". In: *MDPI Sensors* 23.7 (2023). doi: [10.3390/s23073617](https://doi.org/10.3390/s23073617).
- [PF14] Enrico Paolini and Mark Flanagan. "Low-Density Parity-Check Code Constructions". In: *Channel Coding: Theory, Algorithms, and Applications*. Ed. by David Declercq, Marc Fossorier, and Ezio Biglieri. Academic Press, 2014. doi: <https://doi.org/10.1016/B978-0-12-396499-1.00003-0>.
- [PI 23] PI Ceramics. *Piezoelektrische Scheiben*. 2023. url: https://www.piceramic.de/fileadmin/user_upload/physik_instrumente/files/datasheets/PRYY-Datenblatt.pdf (visited on 07/23/2023).
- [Pri+07] Richard Primerano et al. "Echo-Cancellation for Ultrasonic Data Transmission through a Metal Channel". In: *2007 41st Annual Conference on Information Sciences and Systems*. 2007. doi: [10.1109/CISS.2007.4298426](https://doi.org/10.1109/CISS.2007.4298426).
- [RHS20] Bernd-Christian Renner, Jan Heitmann, and Fabian Steinmetz. "ahoi: Inexpensive, Low-power Communication and Localization for Underwater Sensor Networks and μ AUVs". In: *ACM Transactions on Sensor Networks (TOSN)* 16.2 (2020). doi: [10.1145/3376921](https://doi.org/10.1145/3376921).
- [Ric+22] Fabrizio Ricci et al. "Guided waves for structural health monitoring in composites: A review and implementation strategies". In: *Progress in Aerospace Sciences* 129 (2022). doi: [10.1016/j.paerosci.2021.100790](https://doi.org/10.1016/j.paerosci.2021.100790).
- [RMK09] Richard Primerano, Moshe Kam, and Kapil Dandekar. "High bit rate ultrasonic communication through metal channels". In: *2009 43rd Annual Conference on Information Sciences and Systems*. 2009. doi: [10.1109/CISS.2009.5054845](https://doi.org/10.1109/CISS.2009.5054845).
- [Roel15] M. G. L. Roes. "Exploring the potential of acoustic energy transfer". PhD thesis. TU Eindhoven, 2015. url: <https://research.tue.nl/en/publications/exploring-the-potential-of-acoustic-energy-transfer> (visited on 01/12/2021).
- [Ros14] Joseph L. Rose. *Ultrasonic Guided Waves in Solid Media*. Cambridge University Press, 2014. doi: [10.1017/CB09781107273610](https://doi.org/10.1017/CB09781107273610).
- [Rup19] Stefan Johann Rupitsch. *Piezoelectric Sensors and Actuators: Fundamentals and Applications*. Topics in Mining, Metallurgy and Materials Engineering. Springer, 2019. doi: [10.1007/978-3-662-57534-5](https://doi.org/10.1007/978-3-662-57534-5).
- [Sch+22] Thomas Schaechtle et al. "Acoustically Coupled Passive Wireless Sensor System With Mechanical Resonant Sensor". In: *2022 IEEE International Conference on Wireless for Space and Extreme Environments (WISEE)*. IEEE, 2022. doi: [10.1109/WiSEE49342.2022.9926923](https://doi.org/10.1109/WiSEE49342.2022.9926923).

Bibliography

- [Sho+07] D. A. Shoudy et al. "P3F-5 An Ultrasonic Through-Wall Communication System with Power Harvesting". In: *2007 IEEE Ultrasonics Symposium Proceedings*. 2007. doi: [10.1109/ULTSYM.2007.465](https://doi.org/10.1109/ULTSYM.2007.465).
- [SN22] Richard R. A. Syms and Aiman Noorwali. "Polyphase codes for multiplexed acoustic signalling and sensing on pipes". In: *Smart Materials and Structures* 31.9 (Aug. 2022). doi: [10.1088/1361-665X/ac8385](https://doi.org/10.1088/1361-665X/ac8385).
- [SR21] Fabian Steinmetz and Bernd-Christian Renner. "From the Long-Range Channel in the Ocean to the Short-Range and Very Shallow-Water Acoustic Channel in Ports and Harbors". In: *2021 Fifth Underwater Communications and Networking Conference (UComms)*. IEEE. 2021. doi: [10.1109/UComms50339.2021.9598094](https://doi.org/10.1109/UComms50339.2021.9598094).
- [SS13] Alanson P. Sample and Joshua R. Smith. "The Wireless Identification and Sensing Platform". In: *Wirelessly Powered Sensor Networks and Computational RFID*. Springer, 2013. doi: [10.1007/978-1-4419-6166-2_3](https://doi.org/10.1007/978-1-4419-6166-2_3).
- [STM20] Sameer M. Shaik, Xinyao Tang, and Soumyajit Mandal. "Self-Optimizing Wireless Networks on Structures". In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 67.5 (2020). doi: [10.1109/TCSII.2020.2983931](https://doi.org/10.1109/TCSII.2020.2983931).
- [Sto48] H. Stockman. "Communication by Means of Reflected Power". In: *Proceedings of the IRE* 36.10 (1948). doi: [10.1109/JRPROC.1948.226245](https://doi.org/10.1109/JRPROC.1948.226245).
- [Su09] Zhongqing Su. *Identification of Damage Using Lamb Waves : From Fundamentals to Applications*. Ed. by Lin Ye. Lecture Notes in Applied and Computational Mechanics. Springer, 2009. doi: [10.1007/978-1-84882-784-4](https://doi.org/10.1007/978-1-84882-784-4).
- [Sun+21] Yongshun Sun et al. "A Lamb Waves Based Ultrasonic System for the Simultaneous Data Communication, Defect Inspection, and Power Transmission". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 68.10 (2021). doi: [10.1109/TUFFC.2021.3087603](https://doi.org/10.1109/TUFFC.2021.3087603).
- [SWH18] Jafar Saniie, Boyang Wang, and Xin Huang. "Information Transmission Through Solids Using Ultrasound Invited Paper". In: *2018 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2018. doi: [10.1109/ULTSYM.2018.8579702](https://doi.org/10.1109/ULTSYM.2018.8579702).
- [TBL18] Victor Farm-Guoo Tseng, Sarah S. Bedair, and Nathan Lazarus. "Acoustic Power Transfer and Communication With a Wireless Sensor Embedded Within Metal". In: *IEEE Sensors* 18.13 (2018). doi: [10.1109/JSEN.2018.2839558](https://doi.org/10.1109/JSEN.2018.2839558).
- [TCP17] L.G. Tran, H.K. Cha, and W.T. Park. "RF power harvesting: a review on designing methodologies and applications". In: *Micro and Nano Systems Letter* 5.14 (2017). doi: [10.1186/s40486-017-0051-0](https://doi.org/10.1186/s40486-017-0051-0).
- [TMÖ22] Xinyao Tang, Soumyajit Mandal, and Tayfun Özdemir. "A CMOS SoC for Wireless Ultrasonic Power/Data Transfer and SHM Measurements on Structures". In: *IEEE Access* 10 (2022). doi: [10.1109/ACCESS.2022.3214231](https://doi.org/10.1109/ACCESS.2022.3214231).
- [TRI2] Stewart J. Thomas and Matthew S. Reynolds. "A 96 Mbit/sec, 15.5 pJ/bit 16-QAM modulator for UHF backscatter communication". In: *2012 IEEE International Conference on RFID (RFID)*. Apr. 2012. doi: [10.1109/RFID.2012.6193049](https://doi.org/10.1109/RFID.2012.6193049).
- [Tse+20] Victor Farm-Guoo Tseng et al. "Ultrasonic Lamb Waves for Wireless Power Transfer". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 67.3 (2020). doi: [10.1109/TUFFC.2019.2949467](https://doi.org/10.1109/TUFFC.2019.2949467).

- [Tse+22] Victor Farm-Guoo Tseng et al. "Self-detachable through-metal acoustic wireless power transfer". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 69.7 (2022).
- [TSM18] Xinyao Tang, Mohammed Sameer, and Soumyajit Mandal. "Acoustic Wireless Power and Data Telemetry for Structural Health Monitoring". In: *2018 IEEE SENSORS*. IEEE, 2018. doi: [10.1109/ICSENS.2018.8589762](https://doi.org/10.1109/ICSENS.2018.8589762).
- [Web+18] Marcus J. Weber et al. "A Miniaturized Single-Transducer Implantable Pressure Sensor With Time-Multiplexed Ultrasonic Data and Power Links". In: *IEEE Journal of Solid-State Circuits* 53.4 (2018). doi: [10.1109/JSSC.2017.2782086](https://doi.org/10.1109/JSSC.2017.2782086).
- [Xu+22] Yunfei Xu et al. "A Lamb waves-based wireless power transmission system for powering IoT sensor nodes". In: *Smart Materials and Structures* 31.10 (Aug. 2022). doi: [10.1088/1361-665X/ac8ada](https://doi.org/10.1088/1361-665X/ac8ada).
- [Yam+14] M. Yamakawa et al. "Wireless Power Transmission into a Space Enclosed by Metal Walls Using Magnetic Resonance Coupling". In: *Wireless Engineering and Technology* 5.1 (2014). doi: [10.4236/wet.2014.51003](https://doi.org/10.4236/wet.2014.51003).
- [Yan+15] Ding-Xin Yang et al. "Through-Metal-Wall Power Delivery and Data Transmission for Enclosed Sensors: A Review". In: *MDPI Sensors* 15.12 (2015). doi: [10.3390/s151229870](https://doi.org/10.3390/s151229870).
- [YCL18] Ming Yuan, Ziping Cao, and Jun Luo. "Characterization the influences of diodes to piezoelectric energy harvester". In: *International Journal of Smart and Nano Materials* 9.3 (2018). doi: [10.1080/19475411.2018.1454532](https://doi.org/10.1080/19475411.2018.1454532).
- [Yea+13] Daniel Yeager et al. "SOCWISP: A 9 μ A, Addressable Gen2 Sensor Tag for Biosignal Acquisition". In: *Wirelessly Powered Sensor Networks and Computational RFID*. Springer, 2013. doi: [10.1007/978-1-4419-6166-2_4](https://doi.org/10.1007/978-1-4419-6166-2_4).
- [YS16] S. Yang and A. C. Singer. "Energy Efficient Ultrasonic Communication on Steel Pipes". In: *2016 IEEE International Workshop on Signal Processing Systems (SIPS)*. 2016. doi: [10.1109/SiPS.2016.59](https://doi.org/10.1109/SiPS.2016.59).
- [Zha+19] Jianing Zhang et al. "Ultrasonic Wireless Communication Through Metal Barriers". In: *Sound & Vibration* 53.2 (2019). doi: [10.32604/sv.2019.03783](https://doi.org/10.32604/sv.2019.03783).
- [Zor22] Jonathan Zorn. "Equalization in a Software-Defined Acoustic Backscatter Receiver". MA thesis. Hamburg University of Technology, 2022.

Index

- Acoustic Impedance, *see* Impedance
- ASIC, 33, 37, 129, 167
- Attenuation, 41

- Backscatter
 - Communication, 31
 - Energy Consumption, 31
 - Modulation, 35
 - Monostatic, 31
- Barker-Sequence, 91
- Baseband, 8
- Batteryless, 29
- Belief Propagation, 9, 67

- Carrier Frequency, 11
 - Offset, 15, 16
- Channel, 11
 - Frequency-selective, 53
 - Guided-Wave, 40
 - Metal, 51
 - Sandwich-Plated, 40
- Channel Code, 60
 - Encoding, 8
 - LDPC, 9, 65, 103
 - Rate, 8
- Coherence Bandwidth, 12, 48, 53, 54
- Constellation, 9
 - ASK, 9
 - Minimum distance, 9, 64
 - OOK, 9, 34, 141
 - PSK, 9, 77
 - QAM, 9
- Constellation Point Adaption, 95, 105
- Delay Spread, 12, 34, 53, 65, 93, 138
- Demodulation, 70, 92
- Dispersion, 42
- Downlink, 31
- Duty Cycle, 27, 33, 114, 159

- Encoding
 - PDM, 34
 - PIE, 34
- Energy Harvesting, 23, 114
- Envelope, 12, 34, 37
- Equalizer, 14, 85, 93
 - DF, 15
 - Linear, 14
 - MMSE, 15, 93
 - Training, 94
- Equivalent Circuit, 61
- Error-Vector Magnitude, 98

- Frequency
 - Intermediate, 152
- Grey code, 65

- Impedance, 36
 - Acoustic, 53, 61
 - Matching, 44
 - Transducer, 46, 51, 53, 72
- Impulse Response, 12, 61
- Inter-Symbol Interference, 12, 34, 38, 42, 53

Index

- Least Mean Square, 96
- Linear Time-Invariant system, 85

- Mapping, 9
- Maximum Power Point Tracking, 26, 44, 115
- Maximum Power Transfer Theorem, 25, 116, 124
- Message, 8
- Modulation, 11
 - Amplitude, 11, 141
 - ASK, 34
 - Frequency, 11, 35
 - Higher-Order, 38
 - OFDM, 11
 - Phase, 11
- Modulation Range, 77

- Narrowband, 10

- OFDM, 42

- Passband, 8
- Passive Device, 151
- Phase Locked Loop, 16, 121, 129
- Piezoelectric
 - Direct Effect, 20
 - Inverse effect, 20
 - Transducer, 20, 50
- Power Management Unit, 27, 37
- Power transfer
 - Efficiency, 24
- Pulse shaping, 9

- Reader, 31
- Rectification, 26, 36
- RFID, 32
 - Channel, 55
 - Computational, 33, 48

- Sample Frequency Offset, 18, 104, 120, 129

- Sandwich-Plated Channel, 46, 54
- Self-Interference, 35, 86
- Signal-to-Noise Ratio, 12, 63, 75
- Steady State, 50, 61
- Symbol interval, 10
- Synchronization
 - Frame, 15, 16, 66
 - Frequency, 15, 153
 - Timing, 18

- Tag, 31, 32
- Thevenin equivalent, 24
- Timing Error, 18
 - Derivative Detector, 19
 - Locked-Loop, 18
 - Square Detector, 18, 93, 100
- Training Sequence, 14, 91, 99
- Transmitter, 23

- Uplink, 31

- Voltage Multiplier, 26
- Voltage Quadrupler, 114

- Wave
 - Bulk, 27
 - Lamb, 27, 52
 - Longitudinal, 27
 - Transversal, 27
- Wideband, 10, 42
- Wireless Power Transfer, 23, 33
- Wireless Sensing, 32
- WISP, 33, 48