



Dipl.-Ing. Bernhard Großwindhager, BSc

Robust, Efficient, and Scalable UWB-based Positioning using Multipath and Quasi-simultaneous Transmissions

DOCTORAL THESIS

to achieve the university degree of
Doktor der technischen Wissenschaften (Dr. techn.)

submitted to
Graz University of Technology

Supervisor and First Examiner:

Univ.-Prof. Dipl.-Inform. Dr.sc.ETH Kay Uwe Römer
Institute of Technical Informatics, TU Graz

Second Examiner:

Prof. Anthony Rowe, PhD
Carnegie Mellon University, Pittsburgh, USA

Third Examiner:

Assoc.Prof. Dipl.-Ing. Dr. Klaus Witrisal
Signal Processing and Speech Communications Laboratory, TU Graz

Graz, June 2020

AFFIDAVIT

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly indicated all material which has been quoted either literally or by content from the sources used. The text document uploaded to TUGRAZonline is identical to the present doctoral thesis.

Date

Signature

Für meine Familie und Freunde

Danksagung

Eine der wohl schwierigsten Entscheidungen in meinem Leben war der Schritt von der Industrie zurück an die Universität, um meine Dissertation zu verfassen. Allerdings habe ich in den letzten mehr als vier Jahren diese Entscheidung niemals bereut. Auch wenn eine Dissertation mit vielen Entbehrungen einhergeht, so überwiegen doch die positiven Eindrücke und Erfahrungen in überragendem Maße. Vor allem durfte ich auf dem Weg zur Dissertation herausragende und wundervolle Persönlichkeiten kennenlernen. Obwohl ich noch wesentlich mehr Menschen zu Dank verpflichtet wäre, werde ich in den nächsten Zeilen einige davon explizit hervorheben.

Zuallererst bedanke ich mich bei meinem Betreuer Prof. Kay Römer für seine tatkräftige Unterstützung. Trotz seiner unzähligen Tätigkeiten an der TU Graz hatte er immer ein offenes Ohr für meine Anliegen und nahm sich dieser bestmöglich an. Nicht nur bei fachlichen Fragestellungen, sondern auch in der Art und Weise wie er als Führungspersönlichkeit auftritt, konnte ich mir viele Dinge anschauen und für mein weiteres berufliches Leben mitnehmen. Gerade durch die weitreichende Unterstützung von Prof. Römer konnte ich mich in den letzten vier Jahren beruflich entfalten und als Persönlichkeit reifen und mich weiterentwickeln. Ich hoffe, dass sich auch in meinem zukünftigen beruflichen Lebensweg unsere Wege des Öfteren kreuzen werden.

Mein weiterer Dank gilt Prof. Anthony Rowe und Prof. Klaus Witrisal, dass sie sich bereit erklärt haben als Gutachter für meine Dissertation zur Verfügung zu stehen. Es ist mir eine große Ehre, gemeinsam mit Prof. Römer, drei renommierte Experten und Wissenschaftler in ihrem Fachgebiet als Gutachter meiner Dissertation bezeichnen zu dürfen. Weiters konnte ich von Prof. Witrisal unglaublich viel im Bereich Ultra-wideband, Signalverarbeitung und Indoor Positionierung lernen. Ohne diese Unterstützung und auch die Kooperation mit seinen PhD Studenten am SPSC wäre diese Arbeit nicht möglich gewesen.

Ich hatte das Glück, dass ich meine Dissertation im Zuge des Lead Projektes "Dependable Internet of Things in Adverse Environments" erarbeiten durfte. Durch die instituts- und fakultätsübergreifende Kooperation konnte ich meinen Horizont wesentlich erweitern und hatte das Vergnügen mit zahlreichen talentierten Wissenschaftlern arbeiten zu dürfen. Ein besonderer Dank gilt dabei meinen Kollegen Mustafa und Michael. Ich bin davon überzeugt, dass wir in unserem Subprojekt tolle Arbeit geleistet haben und ich bin froh darüber, dass ich in euch nicht nur zwei wunderbare Kollegen, sondern Freunde für mein restliches Leben gefunden habe.

Mein Dank gilt dem gesamten Institut für Technische Informatik für vier wunderbare und unvergessliche Jahre. Auch wenn ich jedem Einzelnen sehr dankbar bin, werde ich ein paar Personen speziell hervorheben. Zuallererst möchte ich dabei Carlo erwähnen, da er mich mit unermüdlichem Einsatz auf meinem Weg zu dieser Arbeit unterstützt hat. Mir ist bewusst, dass es dabei nicht immer einfach für ihn war. Die unzähligen Stunden, die er gemeinsam mit mir kurz vor Deadlines verbracht hat, um meinen Arbeiten den letzten Feinschliff zu verpassen, werde ich ihm nie vergessen. Weiters gilt mein Dank Michael Stocker, der mich zu Beginn mit seiner Masterarbeit und in weiterer Folge als PhD Student tatkräftig unterstützt hat. Neben seiner fachlichen Expertise bin ich auch sehr froh in ihm einen wunderbaren Freund gefunden zu haben und ich hoffe, dass wir noch lange in der ein oder anderen Weise beruflich miteinander zu tun haben werden. Hervorheben möchte ich auch die gesamte EAS Gruppe an unserem Institut. Unsere unzähligen gemeinsamen Mittags- und Kaffeepausen waren die perfekte Ablenkung in stressigen und herausfordernden Zeiten während des PhD Studiums. Ich hatte unglaublich viel Spaß mit euch und ihr seid mir alle sehr ans Herz gewachsen.

Einen Absatz möchte ich auch meinem Ex-Kollegen und Freund Marco Cattani widmen. Ohne ihm wäre ich definitiv nicht in der Position diese Dissertation verfassen zu können. Er war mir gerade zu Beginn eine sehr große Stütze und hat mich in jeglicher Hinsicht versucht zu unterstützen und mich dabei motiviert mein Ziel weiterzuverfolgen.

Mein herzlichster Dank gilt meiner Familie, deren unermüdliche Liebe und Unterstützung hat mich zu der Person gemacht, die ich heute bin. Sie haben mich in meinem gesamten Leben immer und überall unterstützt und mir den nötigen Rückhalt gegeben. Ich bin unendlich dankbar eine solch wundervolle Familie zu haben. Selbiges gilt auch für meine Freunde, sowohl in Graz als auch in Ternberg. Ohne die Unterstützung meiner Freunde und Familie hätte ich diese Dissertation nie zu einem Abschluss gebracht. Sobald ich Zeit mit ihnen verbringen durfte, haben sich alle Probleme der Arbeit verflüchtigt, denn neben Gesundheit sind die wohl wichtigsten Dinge im Leben Familie und Freunde. Deswegen auch die Widmung der Dissertation an euch.

Abstract

Satellite-based positioning systems provide a persistent positioning service to electronic devices. Still, in indoor environments the link quality hardly or not at all allows to decode signals sent by the satellites. Researchers hence strive after finding an indoor equivalent to satellite-based systems to seamlessly localize objects in any environment. Due to its superior time resolution and high immunity to multipath fading, ultra-wideband (UWB) technology seems to be best suited to satisfy the particularly high demand on accurate position information in indoor settings. However, existing UWB-based systems focus mainly on accuracy and partially on precision ignoring that equally important a positioning system has to be *robust, efficient, and scalable*.

In this thesis we present solutions dedicated to each of these properties while still maintaining accurate position estimates. First, we present an adaptation scheme that derives at runtime a set of physical layer parameters to optimize energy efficiency and robustness of UWB communication links. This is required to reliably and efficiently share timestamps in an indoor positioning system. To trigger a parameter change we have designed a link state indicator, which exploits the estimated channel impulse response (CIR) provided by UWB transceivers. This allows to extract information about the link quality and to identify the characteristics of the surrounding environment, such as the presence of NLOS conditions and destructive interference.

Second, we present an indoor positioning system named SALMA that cuts down the required infrastructure to a single physical anchor by exploiting multipath propagation. Indeed, the high bandwidth of UWB allows to resolve multipath components (MPCs) from the estimated CIR. To derive the position of a tag in SALMA, we use a known floor plan to model the theoretical multipath propagation and compare it with the estimated CIR. Due to exploitation of the time and angular domain of the MPCs by employing directional antennas, SALMA provides robust position estimates even in obstructed line-of-sight conditions with a 90% error of 30.7 cm, while requiring minimal infrastructure to enable cost- and time-efficient deployments.

Third, we present two schemes to estimate the distance of a tag to multiple neighbors and its position with a single receive operation. To this end, we employ concurrent ranging, where multiple distinct transmitters intentionally inject signal components in the CIR to derive their distances to a node simultaneously. We tackle several key challenges of concurrent ranging and provide a solution that allows the implementation on low-cost UWB devices. Following, we apply concurrent ranging in a TDoA-based positioning system named SnapLoc that allows to obtain responses from multiple anchors simultaneously without requiring a tight synchronization and to enable passive localization of the tags. Furthermore, we present concepts to overcome the limited timestamp resolution of low-cost UWB transceivers to achieve a 90% error of 33.7 cm at theoretical update rates of up to 2.3 kHz and independently of the number of tags.

Kurzfassung

Satellitengestützte Navigationssysteme bieten elektronischen Geräten die Möglichkeit zur kontinuierlichen Ortsbestimmung. Jedoch ist die Verbindungsqualität in Innenräumen in den seltensten Fällen ausreichend, um die gesendeten Signale der Satelliten dekodieren zu können. Deshalb setzen sich Forscher weltweit das Ziel ein Indoor-Äquivalent zu satellitengestützten Systemen zu finden, welches eine lückenlose Lokalisierung von Objekten in jeglichen Umgebungen ermöglichen soll. Aufgrund der guten Zeitauflösung und hohen Immunität gegenüber Mehrwegeausbreitung gilt die Ultra-wideband (UWB) Technologie als am besten geeignet, um die zukünftigen Anforderungen an ein Indoor-Positionierungssystem erfüllen zu können. Der Fokus bestehender UWB-basierter Systeme ist vorwiegend die Genauigkeit sowie teilweise die Präzision der Positionsdaten zu maximieren. Allerdings ist es ebenso wichtig, dass ein Indoor-Positionierungssystem *robust*, *effizient* und *skalierbar* ist.

In dieser Arbeit stellen wir Lösungen zur Optimierung jeder dieser Eigenschaften vor, ohne dabei die Genauigkeit und Präzision der Positionsinformation negativ zu beeinflussen. Zunächst präsentieren wir ein System, das zur Laufzeit einen Parametersatz für die Bitübertragungsschicht bestimmt, welcher die Energieeffizienz sowie Robustheit von UWB Kommunikationsverbindungen optimiert. In einem Positionierungssystem ist dies erforderlich, um Zeitstempeln effizient und zuverlässig zwischen den Knoten austauschen zu können. Zur Triggerung einer notwendigen Adaptierung der Parameter haben wir einen Indikator entwickelt, welcher in Echtzeit die Verbindungsqualität bewertet. Dieser Indikator nützt die von UWB Transceivern bei Empfang eines Paketes estimierte Kanalantwort. Aus Letzterer können Informationen zur Beschreibung der Verbindungsqualität extrahiert werden, sowie Umgebungs-Charakteristika wie eine blockierte Sichtverbindung und destruktive Interferenzen identifiziert werden.

Zweitens präsentieren wir ein Indoor-Positionierungssystem namens SALMA, welches die Mehrwegeausbreitung ausnützt, um dadurch den Infrastrukturbedarf auf einen einzelnen physischen Anker reduzieren zu können. Tatsächlich ermöglicht die hohe Bandbreite von UWB, dass einzelne Multipfadkomponenten aus der estimierten Kanalantwort extrahiert werden können. In SALMA wird nun zur Positionsbestimmung eines Tags anhand eines bekannten Grundrisses der Umgebung die theoretische Mehrwegeausbreitung modelliert und mit der estimierten Kanalantwort verglichen. Durch die Verwendung von Richtantennen nützt SALMA neben der Zeit- noch Winkelinformation aus, wodurch SALMA eine robuste Positionsbestimmung auch bei eingeschränkter Sichtverbindung mit einem 90% Positionsfehler von nur 30,7 cm erreicht.

Drittens präsentieren wir zwei Lösungen, die es ermöglichen die Entfernung eines Tags zu mehreren Nachbarn sowie dessen Position mit einer einzigen Empfangsoperation zu ermitteln. Zu diesem Zwecke wird ein Verfahren namens Concurrent Ranging eingesetzt. Bei Letzterem injizieren mehrere unterschiedliche Sender eine Signalkomponente in die Kanalantwort, wodurch simultan die Distanz eines Empfängers zu all seinen Nachbarn abgeleitet werden kann. Wir lösen einige der Limitierungen von Concurrent Ranging, um eine Implementierung auf kommerziellen UWB Geräten zu ermöglichen. Darauf aufbauend integrieren wir besagtes Konzept in einem TDoA-basierten Positionierungssystem namens SnapLoc, welches die simultane Akquise von mehreren Ankersignalen ausnützt und dadurch eine passive Lokalisierung der Tags ermöglicht, ohne eine präzise Synchronisation zwischen den Ankern zu erfordern. Darüber hinaus präsentieren wir Konzepte, um den Einfluss der begrenzten Zeitauflösung von UWB Transceivern zu reduzieren. Dies erlaubt einen 90% Fehler von 33.7 cm bei theoretischen Positions-Aktualisierungsraten von bis zu 2.3 kHz und unabhängig von der Anzahl der Tags.

Contents

Contents	xiii
List of Figures	xv
List of Tables	xvii
List of Abbreviations	xix
I Main Part	1
1 Introduction	3
1.1 Indoor Positioning	4
1.1.1 Definitions and terms	5
1.1.2 Applications	5
1.1.3 Technologies	6
1.2 Problem Statement and Challenges	8
1.3 Thesis Statement	10
1.4 Contributions	10
1.5 Scientific Impact	12
1.6 Thesis Structure	12
2 Foundations of UWB Technology	13
2.1 Fundamentals	13
2.1.1 Definition of UWB signals	14
2.1.2 UWB spreading: OFDM vs. IR-UWB	14
2.1.3 Ranging performance: UWB vs. narrowband	15
2.2 UWB Regulations and Standardization	16
2.2.1 Worldwide regulations	16
2.2.2 IEEE 802.15.4 standardization	17
2.3 Time-based Ranging Methods using UWB	19
2.3.1 One-way ToA ranging	20
2.3.2 Time difference of arrival (TDoA)	20
2.3.3 Two-way ToA ranging	21
	xiii

2.4	Channel impulse response (CIR)	22
3	Characterization and Runtime Adaptation of UWB PHY Settings	25
3.1	Characterization of UWB PHY Parameters	27
3.1.1	Experimental setup	27
3.1.2	Impact of PHY parameters on robustness and energy efficiency	28
3.1.3	Summary of PHY parameter characterization	31
3.2	Link State Indicator	32
3.2.1	Deriving the received signal power (RSP)	33
3.2.2	Deriving the environmental state	33
3.3	Runtime Adaptation of PHY Parameters	35
3.4	Evaluation	36
3.4.1	Dynamic environments	36
3.4.2	Destructive interference	37
4	Single Anchor Positioning using Multipath Assistance	39
4.1	System Overview	41
4.2	Design Principles	42
4.2.1	Signal model	42
4.2.2	From physical to virtual anchors	43
4.2.3	Positioning algorithm	44
4.3	Implementation on Low-cost Devices	45
4.3.1	Hardware-specific impacts on signal model and positioning algorithm	46
4.3.2	Directional UWB antenna system	46
4.4	SALMA-light: Exploiting Delay Information	47
4.5	SALMA-full: Exploiting Delay and Angular Information	48
4.6	Evaluation	49
4.6.1	Experimental setup	49
4.6.2	Positioning performance	51
4.6.3	Robustness to non-line-of-sight and dynamic environments	54
5	Scalable and Responsive Positioning using Quasi-simultaneous Responses	59
5.1	Concurrent Ranging	61
5.1.1	Open challenges and contributions	63
5.1.2	Reliable response detection	64
5.1.3	Encoding responder ID in the CIR using pulse shaping	65
5.1.4	Mitigating the impact of multipath using response position modulation	66
5.2	SnapLoc: From Concurrent Ranging to Scalable Positioning	67
5.2.1	Limitations of existing time-based positioning approaches	67
5.2.2	Design principles	69
5.2.3	Improving timestamp resolution	74
5.2.4	Evaluation	76
6	Conclusions and Outlook	81

6.1	Future work	82
A	IEEE 802.15.4f and IEEE 802.15.4z	85
A.1	LRP UWB PHY specification (based on IEEE 802.15.4f)	85
A.2	IEEE 802.15.4z draft standard	86
B	Impact of clock variations in two-way ranging (TWR)	87
B.1	Clock variations in single-sided two-way ranging (SS-TWR)	87
B.2	Clock variations in double-sided two-way ranging (DS-TWR)	87
II	Included Publications	89
	Bibliography	140

List of Figures

2.1	UWB emission limits in (a) the United States (Federal Communications Commission (FCC)) and (b) Europe (EC, ETSI).	16
2.2	UWB PHY frame structure. Adapted from [67].	18
2.3	Positioning based on (a) time of arrival (ToA) and (b) time difference of arrival (TDoA) ranging scheme.	20
2.4	Comparison of (a) single-sided two-way ranging with two packet exchanges and (b) double-sided two-way ranging (DS-TWR) with three transmitted messages.	21
2.5	Channel impulse response (CIR) (a) acquired at the receiver (RX) in a rectangular room (b). The multipath components (MPCs) are denoted with $\tau_1 \dots \tau_4$. Adapted from [70].	23
3.1	Sketch of the presented adaptation scheme to derive a set of physical layer (PHY) parameters at runtime maintaining robust and energy-efficient UWB communication. Adapted from [67].	26
3.2	Packet reception rate (packet reception rate (PRR)) and header reception rate (HRR) as a function of the attenuation level as well as in dependence of the PHY parameters (a) preamble symbol repetition (PSR) and (b) pulse repetition frequency (PRF). Adapted from [67].	29
3.3	Packet reception rate (PRR) and header reception rate (HRR) as a function of the attenuation level as well as (a) data rate and (b) carrier frequency/bandwidth. Adapted from [67].	31

3.4	UWB link state indicator employing environmental awareness.	32
3.5	Packet reception rate (PRR) as a function of the received signal power (RSP) when using default settings (blue) and a data rate of 110 kbps (orange). Adapted from [67].	33
3.6	Estimated channel impulse response in the presence (blue solid line) and absence (orange dashed line), respectively, of destructive interference. Adapted from [68]. .	34
3.7	Evaluation of the adaptation scheme. The reliability (PRR) and energy consumption is compared when using static and adaptive PHY settings, respectively. Adapted from [67].	37
3.8	The proposed adaptation scheme detects at runtime situations of destructive interference (grey area) and counteracts these situations successfully. Adapted from [67].	38
4.1	SALMA system overview. The infrastructure of the system consists of a single anchor node connected to a notebook running the positioning engine. Adapted from [70].	41
4.2	Concept of virtual anchors (VAs) in order to relate the tag position \mathbf{p} to the MPC parameters angle of departure ϕ_k and τ_k by mirroring the physical anchor \mathbf{a}_0 . Taken from [70].	44
4.3	Self-made switchable directional UWB antenna system (a) and the measured antenna patterns of the system at 5 GHz in polar coordinates and azimuth plane (b). Taken from [65].	47
4.4	Evaluation setup (2D-plan): (a) An office (Room A) and (b) a stockroom (Room B) are evaluated at 35 evenly distributed evaluation points (red crosses). Adapted from [70].	50
4.5	Evaluation setup (Image): Both rooms are filled with obstacles and scattering objects. The white dashed line marks the height under obstructed line-of-sight (LOS). Adapted from [70].	50
4.6	Solely exploiting the delay domain of MPCs may result in multipath ambiguities (a), which are resolvable by additionally utilizing the angular domain (b). Adapted from [70].	51
4.7	Due to multipath ambiguities SALMA-light sustains an position error below 30 cm only in 70% of the estimates. By exploiting the angular domain, SALMA-full instead exhibits an error below 30 cm in 99% of the cases (Room A). Adapted from [70].	52
4.8	Mean position (blue cross) and 3-fold standard deviation (black ellipses) for each evaluation point in Room A and under clear LOS (a) and obstructed LOS (b). Adapted from [70].	53
4.9	Positioning performance of SALMA-full in clear LOS, obstructed MPC, and obstructed LOS conditions evaluated in Room A (office environment). Adapted from [70].	54
4.10	Positioning performance of SALMA-full in the stockroom (Room B) and in the case of moving obstacles (storage racks are empty or filled). Adapted from [70].	56
4.11	Mean position (blue crosses) and 3-fold standard deviation (black error ellipses) for each evaluation point in Room B with empty racks. Adapted from [70].	56

4.12	Snippet of a 24-hours experiment in a dynamic office environment. The dashed orange line depicts the mean error of 50 position estimates over time, whilst the solid blue line shows the 90% error. The green line indicates the number of people present. Adapted from [70].	57
5.1	Principle of single-sided two-way ranging (SS-TWR) and concurrent ranging. In the latter, responders transmit <i>RESP</i> messages quasi-simultaneously. Adapted from [66].	62
5.2	Principle of the proposed response detection algorithm. Adapted from [66].	63
5.3	Pulse shape $s_{tx}^{(i)}$ for different values of the TC_PGDELAY register. The x-axis marks the sample, whilst the y-axis indicates the normalized amplitude. Adapted from [66].	65
5.4	Estimated channel impulse response (CIR) (a) and matched filter output (b) in the case of two responders at $d_1 = 4m$ and $d_2 = 10m$, respectively, replying with different pulse shapes. Adapted from [66].	66
5.5	Concurrent ranging applied in an anchor-based positioning system. The tag (T) initiates a position estimation by broadcasting an <i>INIT</i> message to the simultaneously responding static anchors ($A_1 \dots A_4$). Taken from [72].	70
5.6	In SnapLoc a reference anchor A_R sends an <i>INIT</i> message, to which all surrounding anchors ($A_1 \dots A_4$) reply quasi-simultaneously with a <i>RESP</i> message. Adapted from [72].	71
5.7	Using response position modulation by assigning an individual delay δ_j to each anchor, the responses are well separated in the estimated CIR. Taken from [72]. . .	72
5.8	Evaluation setup: Two environments with 28 and 14 evaluation points, respectively, four anchors (blue square) and a reference anchor (purple square). Adapted from [72].	76
5.9	Mean position estimate (blue crosses) and standard deviation (black error ellipses) for each evaluation point in Room A without correction of the transmit timestamp (a), with the wireless correction (b), and with the wired correction (c). Adapted from [72].	77
5.10	Performance of SnapLoc depending on the method used to correct the limited transmit timestamp resolution in the two rooms used in our evaluation. Adapted from [72].	78

List of Tables

3.1	DW1000 supported channels.	28
3.2	Default configuration.	28

3.3	Energy consumption for the transmission and reception of a packet depending on the employed PHY settings and broken down into synchronization header (SHR) and payload. Adapted from [67].	30
3.4	Tradeoff between SHR/data reliability and energy efficiency. Adapted from [67]. .	32

List of Abbreviations

ACK	acknowledgement
AGC	automatic gain control
AoA	angle of arrival
AWGN	additive white Gaussian noise
BER	bit error rate
BLE	Bluetooth Low Energy
BPM	burst position modulation
BPSK	binary-phase shift keying
CDF	cumulative distribution function
CIR	channel impulse response
CPU	central processing unit
CRLB	Cramér-Rao lower bound
CSS	chirp spread spectrum
DS-SS	direct-sequence spread spectrum
DS-TWR	double-sided two-way ranging
EC	European Commission
ECC	Electronic Communications Committee
EDGE	Enhanced Data Rates for GSM Evolution
EIRP	equivalent isotropically radiated power
ETSI	European Telecommunications Standards Institute
FCC	Federal Communications Commission
FH	frequency-hopping
GNSS	global navigation satellite system
GPS	global positioning system
GSM	Global System for Mobile Communications
HPBW	half-power beam width
HRP	high-rate pulse repetition frequency

HRR	header reception rate
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
IPS	indoor positioning system
IR	impulse radio
IR-UWB	impulse radio ultra-wideband
ISM	industrial, scientific, and medical
LBS	location-based service
LEIP	location-enhancing information postamble
LOS	line-of-sight
LQI	link quality indicator
LRP	low-rate pulse repetition frequency
LTE	Long Term Evolution
MAC	media access control
ML	maximum likelihood
MPC	multipath component
MSE	mean squared error
NLOS	non-line-of-sight
OFDM	orthogonal frequency-division multiplexing
OOK	on-off keying
PCB	printed circuit board
PHR	physical layer header
PHY	physical layer
PPM	pulse position modulation
PRF	pulse repetition frequency
PRR	packet reception rate
PSD	power spectral density
PSR	preamble symbol repetition
RF	radio frequency
RFID	radio-frequency identification
RPM	response position modulation
RSP	received signal power
RSS	received signal strength
RTT	round-trip time
SFD	start-of-frame delimiter

SHR	synchronization header
SINR	signal-to-interference-plus-noise ratio
SLAM	simultaneous localization and mapping
SNR	signal-to-noise ratio
SPI	Serial Peripheral Interface
SS-TWR	single-sided two-way ranging
STS	scrambled timestamp sequence
TDoA	time difference of arrival
TH	time-hopping
TH-IR	time-hopping impulse radio
ToA	time of arrival
ToF	time of flight
TWR	two-way ranging
UMTS	Universal Mobile Telecommunications System
US	United States
UWB	ultra-wideband
VA	virtual anchor
WPAN	wireless personal area network

Part I

Main Part

Introduction

Global navigation satellite systems (GNSSs) provide electronic devices with 24-hour three-dimensional position, velocity, and time information anywhere on, or near, the surface of the earth [87, 198]. Besides the military use of GNSS, this enabled numerous civilian applications including air and road navigation [196, 209], surveying and mapping [86], environmental monitoring [59, 217], and location-based services [189]. Among others, the longstanding success of GNSSs was facilitated by its global coverage, low number of outages, and high scalability. The latter is enabled by the *passive one-way* communication of major global navigation satellite systems, such as GPS, GLONASS, and Galileo. Instead of actively emitting electromagnetic waves, GNSS-enabled devices purely listen to the signals transmitted by GNSS satellites. However, in outdoor scenarios with reduced view to open sky and indoor environments the use of satellite-based navigation systems is denied due to limited signal reception [87]. Additionally, an achievable accuracy at the meter level is insufficient for many indoor applications and the complexity of GNSS receivers leads to high power consumption [61] and costs [166].

Given that we spend up to 90% of the day in enclosed buildings [109, 199] and around sixteen hours a day at our homes [24], researchers around the world are driven to find an equivalent to GNSS for indoor scenarios to seamlessly position objects in all environments [132]. The investigated technologies stretch from vision-based [21, 101] over sound-based [84, 127] to magnetic systems [20, 80]. However, due to the ability to penetrate obstacles, a long communication range, and cost efficiency, radio frequency (RF)-based technologies are considered most promising. Among others, this includes indoor positioning systems (IPs) based on Wi-Fi [58, 81, 228], Bluetooth [31, 207, 210], and RFID [89, 164, 170]. These technologies have in common that they are narrowband and, thus, highly susceptible to multipath fading and can hardly achieve sub-meter accuracy [132]. Therefore, in 2007, the IEEE 802.15.4 working group presented the IEEE 802.15.4a amendment. The latter specifies additional PHYs dedicated to highly accurate ranging and localization as well as improved robustness [69, 92].

One of the added PHYs is the ultra-wideband (UWB) technology, which spreads the signal power over a wide bandwidth of more than 500 MHz by transmitting pulses of short duration [212]. This results in a high immunity to multipath fading, superior time-domain res-

olution allowing for accurate ranging, as well as noise-like signal properties to limit cross-technology interference [211, 238]. Although Heinrich Hertz already experimented with UWB signals in the 1880s, the theoretical foundation for ultra-wideband was laid in the 1990s and early 2000s [3, 211, 214]. The breakthrough of UWB took even longer, namely, till the introduction of worldwide UWB standards as well as the availability of low-cost UWB transceivers eventually leading to UWB-enabled smartphones and vehicles [44, 118, 215].

The decision to equip billions of electronic devices with ultra-wideband technology was initially triggered by the automotive domain to avoid relay attacks in car access systems by utilizing UWB's ability to securely and accurately estimate the distance between devices [44, 192]. Additionally, this property lifted UWB to become the most promising technology for indoor positioning systems [106, 117, 129, 193]. However, to establish UWB as the long-term indoor equivalent to GNSS, positioning systems based on UWB further have to provide *robust* position estimates. This requires to maintain robust communication links to reliably share timestamps within a network and to precisely estimate the position of electronic devices even under the challenging propagation conditions of indoor environments. Indeed, moving people, walls, furniture, and other surrounding objects result in severe signal attenuation, multipath propagation, scattering of radio waves, and frequent obstructed or blocked LOS situations [126, 224].

The rapid increase and presence of interconnected objects, marketed under the terms Internet of Things (IoT) and cyber-physical systems, enabled numerous applications with high societal and economical relevance and impact such as mobile autonomous systems [68, 73], optimized production and supply chain management [29, 155], as well as assisted living [157, 218]. To enable these use cases, the typically battery-powered devices require context information of its environment including accurate location information [151]. Hence, providing sustainable services requires an *efficient* positioning system bearing the limited energy budget in mind.

Besides the resource efficiency of individual devices, in an IoT network a large number of heterogeneous devices have to co-exist without negatively affecting the overall quality of service. Hence, a positioning system is stringently required to be *scalable* in terms of device density [126, 158]. Indeed, solving the scalability problem enabled the long-lasting success of GNS systems such as GPS. The latter guarantees a constant performance as well as position update rate regardless of the number of users. Suitable indoor positioning systems have to strive for a similar behavior to take the prospective billions of interconnected devices into account.

1.1 Indoor Positioning

The availability of position information provided by satellite-based positioning systems has shifted the research and business focus to indoor environments in which the signal-to-noise ratio (SNR) hardly or not at all allows to decode signals sent by the satellites. Instead, dedicated indoor positioning systems have to complement the GNSSs to enable a location-aware Internet of Things eventually opening up a new dimension called Localization of Things [32, 176]. After defining important terms related to indoor positioning in Sect. 1.1.1, Sect. 1.1.2 discusses application domains enabled by the availability of seamless position information combined with the rapid increase of low-cost connected devices. Due to the diversity of these potential applications, several technologies were investigated to find the best fit for future IPSs (Sect. 1.1.3).

1.1.1 Definitions and terms

In literature, the terms *positioning* and *localization* are commonly used synonymously. However, *positioning* (short for position determination) is the process of obtaining a position (i.e., a set of coordinates) of an object related to a well-defined coordinate system. It is differed between *absolute* and *relative* positioning (i.e., estimating a position based on previous positions) [85]. *Localization* (location determination) instead describes the process of determining the position in terms of topological relations [135] and is often used for low-accuracy systems [132]. Although *positioning* is in most cases the proper term, within this thesis also *localization* will be used at times due to its pervasiveness in the wireless sensor networks community and since positions of objects are estimated occasionally with respect to a given floor plan.

Typically, positioning systems are *anchor-based* meaning that a set of static devices at a known position is used to determine the a-priori unknown coordinates of mobile devices or nodes. The latter can be used to cooperatively refine the position estimates of the mobile devices and extend the coverage [225]. Even though GNSSs do not contain static nodes, still, they employ the *anchor-based* principle as the position of the satellites is known at any point in time [234]. Within this thesis likewise *anchor-based* indoor positioning systems are considered and the static nodes are consistently called *anchors* and the mobile nodes are referred to as *tags*.

Depending on the direction of communication between tag(s) and anchor(s) one differentiates between *client-* and *server-based* positioning. In the latter case tags are transmitting signals to the anchors and the position estimation is performed remotely on a back-end server connected to the anchors [132]. In such a centralized scheme the server holds position information of all tags enabling asset or people *tracking*. Instead, in *navigation* scenarios the position information is required at the mobile devices corresponding to *client-* or *mobile-based* positioning [234]. GPS is a typical example for such a *navigation* system as the satellites are transmitting radio waves and the GPS-enabled devices receive the signals and carry out the position determination locally. In this thesis, among others, a GPS-like *client-based* positioning system is presented.

1.1.2 Applications

The advancement in wireless technologies and the widespread penetration of wireless devices enabled the progression from (relatively) static applications such as remote environmental monitoring (e.g., heat, humidity, or air pollution) in wireless sensor networks [232] towards highly dynamic and more sophisticated use cases. Within this section, the latter will be divided into three categories, namely (i) tracking and monitoring, (ii) navigation and guidance, and (iii) location-based services and analytics. For each category a few use cases are discussed, however, please note that this is not a complete list. Instead, it is expected that further application domains will emerge in the near future due to the continuing progress to equip smartphones, vehicles, and wearables with highly accurate and robust positioning capabilities.

Tracking and monitoring. As discussed in Sect. 1.1.1, server-based systems hold the position information at a central server allowing to track people, animals, and assets. This enables identification and authentication of people [174], medical monitoring of elderly and disabled people to detect emergencies and falls [157,218], and to track employees and workers in automatic safety systems at construction- and production sites [16]. Furthermore, tracking systems are used to

detect if an asset has unauthorizedly left a certain area to prevent thefts or losses of important and expensive instruments and products [7]. In the retail market, positioning systems allow automated inventory tracking and real-time stock information of products [104]. Recently, also fully automated self-checkout systems and cashier-less stores have gained interest [216, 243].

Navigation and guidance. The process of planning and controlling the movement of an object from one place to another is called navigation and guidance and it requires precise position and trajectory determination [85]. In indoor environments this covers the navigation of customers to stores and goods [152], the guidance of visitors through exhibitions and trade shows [19] as well as the support of visually impaired people. Moreover, besides humans, also robots and drones use indoor positioning systems to enable the delivery of goods [142], automatically parking vehicles [205, 210], augmented reality applications [150], and even robot-aided surgeries [208].

Location-based services and analytics. Typically, a navigation system is combined with (wireless) communication capabilities to obtain position-related (geo-)information and enable so called location-based service (LBS) [132]. Hence, an LBS is based on the combination of navigation, information, and communication, where the localization system is a key component that is generally rarely noticed by the user [85]. The previously mentioned applications such as guidance through museums or stores are strictly speaking location-based services which can be extended with proximity-based information and personalized marketing such as push notifications about special offers [191]. Additionally, the collected information might be used to analyze visitor behavior or to optimize industrial processes and flow of visitors [233].

1.1.3 Technologies

Towards identifying an indoor equivalent to GNSSs, numerous different means to estimate the position of a mobile device were investigated. Within this section the most prominent technologies are discussed and categorized into (i) optical technologies, (ii) RF technologies and (iii) alternative technologies that are not employing electromagnetic waves.

Optical technologies. Positioning systems based on infrared and visible light cameras [21, 23] as well as LIDAR [222] achieve accuracies from a few decimeters down to sub-mm enabling sophisticated and demanding applications such as robot-aided surgeries [208]. In such well-defined and limited areas of operation, optical systems outperform other technologies. However, equipping larger areas such as shopping malls, shop floors, museums, and homes with optical systems is impractical due to the limited range and exuberant costs making optical technologies infeasible for most IoT applications. Furthermore, systems based on cameras often conflict with privacy concerns and all systems have in common that they are unusable in non-line-of-sight (NLOS) situations, which further decreases the potential domain of use cases.

RF technologies. Although heavily attenuated in NLOS conditions, RF signals instead are capable of penetrating obstacles resulting in longer communication ranges [166]. Besides the higher cost efficiency compared to optical systems, this ensures a substantially larger coverage [132]. However, due to the longer wavelengths, the accuracy of RF systems is limited to the centimeter to meter(s) range. In the following, the most common wireless technologies for RF-based indoor localization are discussed, namely Wi-Fi, BLE, RFID, and UWB.

IEEE 802.11 (Wi-Fi). Systems employing Wi-Fi signals typically exploit the position dependency of the received signal strength (RSS) combined with fingerprinting [228], in which pre-measured radio maps (fingerprints) are correlated with the present RSS of surrounding Wi-Fi access points. The latter are often readily available in indoor sites allowing for cost-efficient deployments. However, RSS estimates are highly dependent on the propagation environment and show a high time variability [132]. Moreover, the fingerprints have to be acquired in a time-consuming offline phase and kept up-to-date to cope with dynamics. Depending on the density of access points and fingerprints, the achievable accuracy is 2 - 15 m for ranges up to 150 m [97].

Bluetooth Low Energy (BLE). BLE was merged into the Bluetooth Specification v4.0 in 2010 and offers a considerably lower power consumption than classical Bluetooth [69]. The high energy efficiency, low cost, as well as high market penetration makes it a reasonable technology for IPS [240]. To this end, BLE beacons are continuously transmitting signals, whose RSSs are recorded by BLE-capable devices [52]. Like Wi-Fi, BLE-based systems exploit the crowded ISM band and the accuracy is similar. However, dedicated infrastructure is required and the range is limited to 20 - 30 m [12], which demands a denser network of BLE beacons [228].

Radio-frequency identification (RFID). An RFID system uses a reader to identify and track tags. It is differentiated between active RFID systems, in which the tags are equipped with transceivers and internal power supply and passive RFID systems employing inductive coupling making a power supply dispensable. However, passive tags are limited to a detection range of 2 m [12], whereas active tags extend the range up to 30 m [132]. To localize mobile readers or tags, proximity detection or RSS is used. Similar to BLE and Wi-Fi, the limited resolution of provided timestamps makes RFID barely suitable for time-based positioning.

Ultra-wideband (UWB). Although fingerprinting UWB positioning systems exist [184], it is more common to derive the position-related information from the ToA and TDoA of the received UWB signal(s). To this end, the increased bandwidth is utilized to enable accurate and precise timestamping. Moreover, besides accuracies in the cm range even in harsh environments [35, 100], the high bandwidth enables (i) high throughput [61], (ii) high immunity to interference effects [132, 214] and (iii) reduced multipath fading [213]. Compared to Wi-Fi, readily installed infrastructure is not available, however, the decision to equip smartphones with UWB [118, 215] might enable a high market penetration.

Alternative technologies. Thus far, technologies employing electromagnetic waves are discussed, however, mechanical waves such as audible sound and ultrasound are likewise capable of localizing objects [74]. Due to the appealing attribute of the relatively slow speed of sound, the timing requirements are less stringent and cm accuracy and better is achievable [134]. However, the operating range is limited to 10 m and sound systems are highly susceptible to temperature changes, NLOS conditions, and multipath propagation [132]. Systems based on magnetic and electromagnetic fields, instead, have the potential to penetrate obstacles and operate in limited LOS [177]. Depending on the applied approach, however, the range is limited to a few meters [74]. To enhance the accuracy, one can exploit inertial measurement units, consisting of accelerometers, gyroscopes and/or magnetometers. Due to the significant drift of these sensors, it is not advised to use them standalone, however, due to their short term accuracy they are fused with systems providing absolute position information [17].

1.2 Problem Statement and Challenges

Considering the widespread application domain of indoor positioning systems (see Sect. 1.1.2) and the numerous investigated technologies (see Sect. 1.1.3), respectively, it is evident that likewise the requirements and criteria are highly diverse. However, analyzing the literature in recent years shows that the focus in the past was mainly on accuracy and partially on precision¹. Indeed, authors were praising the low median error of the proposed solutions, ignoring that equally important a positioning system has to be, among others, *robust*, *efficient*, and *scalable* [74].

Robust. Indoor environments have the beneficial properties of being less affected by weather conditions and showing small temperature gradients. However, the wave propagation is inherently challenging due to (i) severe multipath fading, (ii) frequent NLOS conditions, (iii) fast temporal changes caused by moving objects and people, and (iv) high attenuation and scattering [132]. Among the technologies discussed in Sect. 1.1.3, UWB seems to be best suited to satisfy the particularly high demand on accurate and precise position information in indoor settings [129], thus, this thesis will focus exclusively on UWB. Although its performance is heavily impacted by NLOS conditions, UWB offers, in contrast to optical technologies and sound systems, the possibility to penetrate obstacles and remain functional in crowded environments. Moreover, due to its high bandwidth and superior time resolution, UWB outperforms other RF technologies such as Wi-Fi, BLE, and RFID due to its higher immunity to multipath fading [4] and the ability to accurately and precisely determine the reception time of packets [35]. However, the latter is not sufficient to provide a robust² indoor positioning system in dynamic and harsh environments. Indeed, the acquired timestamps to estimate distances have to be reliably shared in a network, which requires robust communication links. To this end, currently applied static physical layer settings are incapable of coping with dynamic wireless channels, instead, the parameters shall adapt at runtime to a degrading communication performance, ideally depending on the characteristics of the environment and the wireless channel. To design an effective and robust adaptation mechanism, a detailed understanding of the UWB physical layer settings and their impact on robustness as well as the ability to assess the quality of UWB links is inevitable.

Efficient. To enable the applicability of an indoor positioning system in an ubiquitous fashion, the position information has to be provided efficiently to the users and electronic devices. In this regard, it is differed between time-, cost-, and energy efficiency as they are equally important.

Time-efficient. After comparing the performance of more than 100 indoor positioning systems, Lymberopoulos and Liu [129] concluded that the setup procedure of existing solutions is excessively time-consuming and labor-intensive. This state of affairs represents a serious problem as the success and the market penetration of an IPS highly depends on the deployment effort, meaning the time and man power required to, among others, (i) deploy hardware and infrastructure, (ii) acquire radio maps to support fingerprinting-based solutions, and (iii) train a model-based system. Hence, ideally, a solution decreases the deployment effort required to setup and install the system while still allowing to sustain a high positioning accuracy and robustness.

¹*Accuracy* is defined as the divergence of the estimated position from the true position, whereas the *precision* is defined as the percentage that a certain accuracy is achieved [166].

²Within this thesis, the term *robustness* relates to precise position estimates as well as reliable communication and timestamp acquisition rather than robustness in terms of physical damage.

Cost-efficient. In accordance with time and labor effort, respectively, the costs of an indoor positioning system are raising. Indeed, the more anchors and infrastructure is demanded and has to be deployed, the higher are the overall costs and consequently the less applicable is such a system in large-scale applications. Moreover, positioning systems ask for small and ubiquitous wireless devices with a modest price per unit. Hence, the hardware costs of the infrastructure and the effort to install them as well as the complexity and costs of the nodes have to be minimized.

Energy-efficient. Due to the ubiquitous and at times mobile nature of connected electronic devices, the latter are typically battery driven and, thus, frequent recharging or replacement of these batteries impacts the user experience as well as the availability of the positioning service. Hence besides the positive effect for the environment and reduced costs, it is mandatory to optimize the energy efficiency to maintain a continuous service of the indoor positioning system. This is even more severe considering that UWB radios have a higher current consumption than transceivers based on narrowband technologies such as BLE or Wi-Fi [119] and that a large number of UWB systems require the exchange of several consecutive messages to derive the distance between the mobile tag and multiple anchors to unambiguously determine the position of the tag [106].

Scalable. Although designed in the 70's, still, more than forty years later GPS is probably the most successful positioning solution. Among others, this is due to a persistent performance of the system independent of the number of users utilizing the positioning service. Considering that today we see billions of connected devices with steadily increasing figures [187], it is evident that likewise an indoor positioning system has to provide a similar scalability in terms of tag density, i.e., to support multiple tags without negatively impacting the system performance. However, most of the existing indoor positioning solutions based on UWB technology disregard the scalability property [158]. As a result, current systems typically support only a few tags due to (i) the large number of messages exchanged and (ii) the use of scheduling techniques for collision avoidance. Indeed, besides the high energy consumption, multiple messages exchanged to carry out each distance estimation limits the overall update rate [108] and requires a tag to be heavily involved in the communication. This makes scheduling techniques to avoid collisions between different mobile tags inevitable, hence, such active systems are insufficient in terms of scalability, as an increasing tag density leads to reduced performance or position update rate [132]. To reduce the message overhead and avoid exchanging consecutive messages, a few UWB-based systems employ TDoA techniques. In the tag-initiated TDoA approach a tag broadcasts only one message per position estimate [161]. The message is received from synchronized anchors which compute the TDoA between anchor pairs and communicate back the estimated position to the tag. Whilst this allows to minimize the number of transmissions carried out by a tag and to shift the computational burden to more powerful anchors, one still needs to allocate specific timeslots to each tag in order to avoid collisions. This limits the number of supported tags and, consequently, the scalability of a positioning system. Furthermore, active tag systems are controversial in terms of privacy as a third party could overhear the exchanged packets. In contrast, the anchor-initiated TDoA approach does not need the tag to actively transmit information. However, a tight synchronization in the ns-range between the anchors is still required, which results in a significant overhead [193,227] and is challenging to maintain [223]. Hence, ideally, a positioning system does not need a tight synchronization, but still keeps the tags passive to allow localizing an unlimited number of tags at an update rate of hundreds of Hz.

1.3 Thesis Statement

Exploiting the increased time resolution and fine-grained channel information provided by ultra-wideband transceivers allows to design a robust, efficient, and scalable indoor positioning system that adapts to environmental changes at runtime, minimizes required infrastructure and susceptibility to non-line-of-sight conditions, while maximizing update rate and responsiveness.

1.4 Contributions

The scientific contributions of this thesis in the area of UWB-based indoor positioning are separated in three parts, where each contribution is reflected in a dedicated chapter.

Contribution 1: Characterization and runtime adaptation of physical layer settings. Besides a significantly higher bandwidth than other IoT technologies, IEEE 802.15.4-compliant UWB transceivers provide several physical layer settings to fine-tune the radio sensitivity and energy consumption. This is different from narrowband systems, where most of the knobs to tune the energy efficiency and robustness of communications are at the MAC layer, e.g., duty cycle [241], clear channel assessment threshold [173], and backoff times [25]. Due to the lack of information in the research community, the first contribution of this thesis is to characterize the impact of PHY settings including preamble symbol repetitions, pulse repetition frequency, data rate, bandwidth, and carrier frequency on the performance of UWB communications. It is shown that by tuning the parameters, one can gain up to 8 dB additional link margin. Additionally, we investigate the energy consumption of different PHY configurations to highlight privileged settings in order to optimize link reliability at minimal energy costs.

Based on the characterization and the application requirements, we develop a scheme that derives an optimal set of physical layer settings at runtime. To trigger parameter adaptations, this adaptation algorithm requires real-time information about the link quality. To this end, we design a robust link state indicator. The latter exploits the estimated CIR provided by standard-compliant UWB transceivers to extract information about the characteristics of the surrounding environment and the received signal strength. An extensive experimental evaluation using the Decawave DW1000 transceiver shows the effectiveness of the adaptation scheme to detect and tackle the presence of destructive interference and to maintain a *robust* and *efficient* communication link even under harsh environmental conditions.

Contribution 2: Exploiting instead of mitigating multipath information. In the second contribution of this thesis we present an indoor positioning system that cuts down the required infrastructure to a single physical anchor by exploiting instead of mitigating multipath propagation, i.e., specular reflections originating from static objects. The system solely exploits a floor plan showing the geometry of the building, without requiring a time-consuming setup phase, prior calibration, fingerprint acquisition, or training of a model. Although the high bandwidth of UWB transceivers enables to extract multipath information from the estimated CIR, state-of-the-art UWB-based systems utilize only the leading edge of the CIR, neglecting the valuable position-related information provided by MPCs. In this thesis, however, we exploit the known floor plan and the anchor location to model the theoretical multipath propagation and compare

it with the estimated CIR. This enables to unambiguously determine the position of a tag with only a single physical anchor while still achieving an accuracy comparable to the one achieved by common multi-anchor UWB systems. Still, overlapping MPCs might degrade the performance, thus, we exploit self-made directional antennas to improve the *robustness* of the system by exploring angular information of the MPCs. An experimental evaluation of the system shows that 90% of the position estimates exhibit an error below 20.17 cm, with a median error below 8 cm. Due to the exploitation of time as well as the angular domain of multipath components, even in the presence of obstructed line-of-sight 90% of the position estimates exhibit less than 30.7 cm error and a median error below 15 cm. Thus, the presented system provides *robust* position estimates even in non-line-of-sight conditions while requiring minimal infrastructure to enable *cost-* and *time-efficient* deployments.

Contribution 3: Scalable and responsive positioning using quasi-simultaneous responses.

The third contribution of this thesis are two schemes to estimate the distance of a tag to multiple neighboring nodes as well as its position with a single receive operation. This contribution is separated in the first building block scalable ranging and the second one scalable positioning.

Scalable ranging. Estimating the distance between two nodes in a non-synchronized network requires a two-way ranging scheme, i.e., a pair-wise exchange of at least two messages. Hence, to estimate the distance between N nodes, one typically needs to schedule the exchange of $N \cdot (N - 1)$ messages, which is time-consuming and energy-inefficient. Similar to contribution 2 where individual multipath components are extracted, one can extract quasi-simultaneous responses of an arbitrary number of responders from the CIR. This concept is called concurrent ranging and was experimentally shown in [33], but without addressing several key challenges and hindering a practical solution. These challenges include (i) to design an algorithm to reliably detect multiple responses in the CIR even in the case of overlapping signal components, (ii) to identify responders, i.e., to associate a distance estimate to a specific responder, and (iii) to mitigate the impact of multipath components potentially leading to misclassification of MPCs as responses. Within this thesis, we tackle all these challenges and provide a solution that allows the practical implementation of concurrent ranging on off-the-shelf UWB devices. To this end, among others, we present novel techniques called pulse shaping and response position modulation that allow to associate a distance measurement to a specific responder and to prevent the overlap of responses and strong multipath components from other responders.

Scalable positioning. The concept of concurrent ranging enables highly responsive, i.e., high update rate and low latency, positioning systems. However, due to the active tag initiating a concurrent ranging, collision avoidance techniques are still required. Instead, within this thesis, we apply the concept to a TDoA-based anchor-initiated approach to obtain responses from multiple anchors quasi-simultaneously without requiring the anchors to be tightly synchronized. Moreover, due to the initialization message sent by a reference anchor, the tag does not require to actively transmit messages and in fact its radio-on time is reduced to a single receive operation. The system hence allows passive and privacy-preserving positioning at theoretical update rates of up to 2.3 kHz. Despite the limited transmit timestamp resolution of low-cost UWB transceivers, an evaluation of the system shows that a 90% error of 33.7 cm and a median error of 18.4 cm is achievable. Thus, the performance is comparable to systems based on TWR, but at significantly higher position update rates and independently of the number of tags.

1.5 Scientific Impact

In total, fourteen peer-reviewed publications were authored and co-authored by the writer of this doctoral thesis, out of which four are an integral part of the thesis. These publications have been published in top-tier conferences, namely IEEE Intern. Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM), IEEE Intern. Conference on Distributed Computing Systems (ICDCS), ACM Conference on Embedded Networked Sensor Systems (SenSys), and ACM/IEEE Intern. Conference on Information Processing in Sensor Networks (IPSN).

The paper "*SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags*" was awarded with the **Best Paper Award** at the IPSN 2019 conference.

The demo "*UWB-based Single-anchor Low-cost Indoor Localization System*" was awarded with the **Best Demo Award** at the SenSys 2017 conference.

The poster "*Switchable Directional Antenna System for UWB-based Internet of Things Applications*" was awarded with the **Best Poster Award** at the EWSN 2017 conference.

Besides the scientific contributions as well as a book chapter published in Springer International Publishing, the work discussed in this thesis resulted in two patent applications regarding a single anchor positioning system using multipath assistance and a highly scalable and quasi-simultaneous indoor positioning system, respectively.

1.6 Thesis Structure

The remainder of this dissertation is organized as follows. Chapter 2 discusses the fundamentals of the ultra-wideband technology, its regulations and standardization, as well as applicable time-based ranging methods. Chapter 3 characterizes the PHY settings of IEEE 802.15.4-compliant UWB transceivers and proposes a scheme to adapt these settings at runtime to maximize robustness and energy efficiency of UWB communication links. It is followed in Chapter 4 by an indoor positioning system that exploits instead of mitigating multipath propagation. This allows to reduce the required infrastructure to a single physical anchor. In Chapter 5, the high bandwidth of UWB transceivers is exploited to present a scheme to acquire responses and distance estimates from multiple neighboring nodes within a single receive operation. This scheme is extended to a scalable and responsive indoor positioning system where multiple anchors are responding quasi-simultaneously. Chapter 6 concludes with a summary of the contributions and an outlook on potential future work. A collection of the publications is attached in Part II.

Foundations of UWB Technology

This chapter covers the fundamentals and history of ultra-wideband technology in Sect. 2.1 discussing the definition of UWB signals (Sect. 2.1.1), the differences between OFDM and IR-UWB (Sect. 2.1.2), as well as the benefits of UWB over narrowband technologies for indoor positioning (Sect. 2.1.3). Following the fundamentals, Sect. 2.2 gives an overview of the worldwide UWB regulations (Sect. 2.2.1) and standardization (Sect. 2.2.2). In Sect. 2.3 time-based ranging methods using UWB, namely, one-way ToA ranging (Sect. 2.3.1), TDoA (Sect. 2.3.2), and two-way ToA ranging (Sect. 2.3.3) are analyzed. This chapter closes with a discussion of the channel impulse response including the introduction of a signal model in Sect. 2.4.

2.1 Fundamentals

When Heinrich Hertz generated the first man-made electromagnetic waves in the 1880s, he was experimenting with sparks and was thus, generally speaking, working with ultra-wideband signals. Likewise, Guglielmo Marconi was using UWB signals to transmit Morse Code messages over a few miles and presumably also across the Atlantic Ocean [41, 166]. However, UWB was banned in the 1920s due to the large occupied spectrum and was primarily limited to military applications [245]. Consequently, it was mostly replaced by narrowband carrier-based systems. Thereupon, it took until 1989, when the U.S. Department of Defense coined the term ultra-wideband and until the 1990s to establish an ultra-wideband research community and to lay the theoretical foundations including the introduction of time-hopping impulse radios by Win and Scholtz [168, 212]. One of the major milestones for UWB wireless communications followed in 2002, when the US frequency regulator FCC allocated the 3.1-10.6 GHz frequency band for unlicensed operation of UWB devices with a maximum equivalent isotropically radiated power (EIRP) of -41.3 dBm/MHz. The latter is also the limit for unintentional radiators (e.g., TVs and monitors) [53]. Gradually, also other countries - with slight differences to the FCC spectrum mask - published their own UWB regulations [69, 238].

2.1.1 Definition of UWB signals

According to the definition by the FCC and International Telecommunication Union Radio-communication Sector (ITU-R), UWB signals have an absolute bandwidth $B \geq 500$ MHz or a fractional (relative) bandwidth B_{frac} larger than 0.2 [98], where B_{frac} is defined as:

$$B_{frac} = \frac{B}{f_c} = \frac{f_u - f_l}{f_c}, \quad (2.1)$$

where f_u and f_l denote the upper and lower frequencies, respectively, at which the power spectral density (PSD) is 10 dB below its maximum and f_c is the center frequency given by

$$f_c = \frac{f_u + f_l}{2}. \quad (2.2)$$

Hence, the definition of UWB neither defines applications nor the modulation: it rather entails that the components of a UWB system have to be capable of handling a wide spectrum [245].

2.1.2 UWB spreading: OFDM vs. IR-UWB

There are multiple ways to spread signals to the large bandwidth defined in Sect. 2.1.1, including the classical approaches of frequency-hopping (FH) and direct-sequence spread spectrum (DS-SS). In FH, different carrier frequencies are used sequentially according to a user-specific hopping sequence. This enables fairly simple transceivers consisting of narrowband modulators and dynamic oscillators. However, FH possibly leads to significant interference to legacy systems, as the latter might be exposed to the maximum power of the UWB signal at a given time [238]. In DS-SS based systems, each bit of a transmit signal is multiplied with a high-rate spreading sequence effectively increasing the bandwidth approximately to that of the spreading signal [139]. Despite its success in cellular communications and its early use in high data rate UWB systems [45, 159], DS-SS is less pervasive for low-power communication networks mainly due to the high-rate processing required at the receiver [238]. Thus, the approaches considered as most suitable to exploit the UWB spectrum are orthogonal frequency-division multiplexing (OFDM) and time-hopping impulse radio (TH-IR). In OFDM, a UWB signal is generated by combining parallel narrowband signals (sub-carriers) occupying different portions of the UWB spectrum. Due to the orthogonality of different sub-carriers their spectra may overlap leading to a high spectral efficiency. In time-hopping impulse radio instead, a data symbol is represented by a sequence of ultra-short pulses with a duration of sub-nanoseconds to nanoseconds. This corresponds to UWB signals in the frequency domain. Analogous to FH in the frequency domain, TH-IR pulses are hopping in time according to a user-specific sequence [211]. Interferers may distort the pulse shape preventing, in the worst-case, a receiver from correctly receiving the UWB signal. OFDM UWB systems, instead, enable blacklisting of certain sub-carriers in case of a degrading and interfered link. However, in contrast to OFDM, the impulse radio principle allows for low-complexity and low-power transmitter architectures, as it is possible to generate basic pulse shapes using simple analog components [211, 245]. Due to this property, TH-IR is highly appealing for low-power networks [238] and consequently for highly-efficient positioning system emphasized by the decision to employ TH-IR UWB in the IEEE 802.15.4 standards (see Sect. 2.2). Hence, this thesis focuses on TH-IR UWB systems.

2.1.3 Ranging performance: UWB vs. narrowband

Exploiting a bandwidth of several hundred megahertz corresponds to a time resolution in the order of a few nanoseconds. In the spatial domain this allows to distinguish multipath components (MPCs) from reflecting surfaces separated by only a few decimeters [146]. This enables to dissociate the first signal path or LOS component from the multipath, which allows to accurately and precisely determine its time of arrival (ToA) and consequently the time of flight of the signal and hence the distance to the source. Instead, when using narrowband systems, MPCs will interfere constructively and destructively with the LOS component, which will introduce fading and pulse distortions and thus increases the uncertainty of the ToA estimation.

To mathematically quantify the impact of bandwidth on the ranging performance, one may employ the Cramér-Rao lower bound (CRLB)¹ [105]. The latter defines a theoretical lower limit on the variance of the estimated time of arrival $\hat{\tau}_a$ of the first path component and the mean squared error (MSE) in case of an unbiased estimate $\hat{\tau}_a$,

$$\text{var}\{\hat{\tau}_a\} = \mathbb{E}\{(\hat{\tau}_a - \tau_a)^2\} \geq \text{CRLB}, \quad (2.3)$$

where CRLB is defined in ideal conditions, hence, ignoring multipath and assuming additive white Gaussian noise (AWGN) with zero mean and a spectral density of $N_0/2$ by [35, 200]

$$\text{CRLB} = \frac{1}{8\pi^2\beta^2\text{SNR}} \quad (2.4)$$

with β being the effective signal bandwidth, $\text{SNR} = E/N_0$, and E the average received signal energy. Please note that the lower bound decreases inversely proportional to the squared bandwidth β^2 and the SNR. This indicates that an increased signal energy and/or bandwidth improves the ranging performance, which is not necessarily the case in RSS-based systems [166]. Multiplying the ToA estimate $\hat{\tau}_a$ with the speed of light in air c results in the achievable standard deviation of a range estimate \hat{d} as follows (assuming synchronized nodes, see Sect. 2.3.1):

$$\text{std}\{\hat{d}\} = \sqrt{\text{var}\{\hat{d}\}} = c \cdot \sqrt{\text{var}\{\hat{\tau}_a\}} \geq \frac{c}{2\sqrt{2}\pi\beta\sqrt{\text{SNR}}}. \quad (2.5)$$

This equation shows the dependency of the ranging performance on the utilized bandwidth and highlights the superior capabilities of UWB-based systems compared to narrowband technologies. Considering a SNR of 5 dB, the lower limit of the achievable standard deviation of a range estimate for a UWB systems with a bandwidth $\beta = 500$ MHz is about 3.8 cm and for $\beta = 1$ GHz about 1.9 cm. In narrowband systems, instead, this bound is significantly higher. Considering a Wi-Fi based system with a bandwidth of 20 MHz, the lower limit increases to 94 cm and with a system employing BLE and a bandwidth of 2 MHz even to 9.5 m. This emphasizes the impracticality of time-based positioning for narrowband technologies. Instead, it is more common to employ the less robust and less accurate distance estimation based on RSS. In the presence of multipath, the impact of bandwidth on the ranging performance is even more severe, as the interfering multipath also scales reciprocally with bandwidth [69]. This leads to an inverse super-linear scaling of the square root of the CRLB with the bandwidth [219]. For the explicit derivation of the CRLB in presence of multipath please refer to [35, 219, 237].

¹There are also other bounds such as the Ziv-Zakai bound [244] and Barankin bound [133], which are more accurate at low and moderate SNR [35]. However, they are out of the scope of this doctoral thesis.

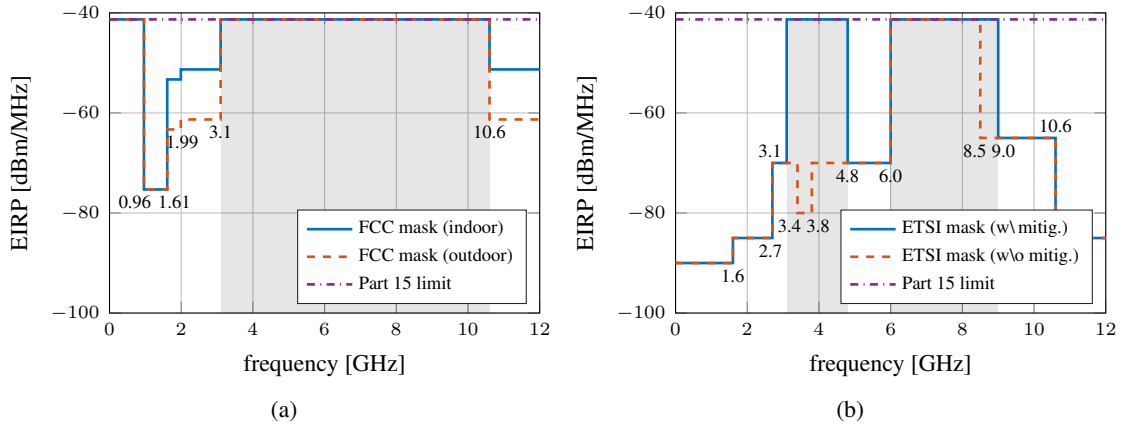


Figure 2.1: UWB emission limits in (a) the United States (FCC) and (b) Europe (EC, ETSI).

2.2 UWB Regulations and Standardization

Enabling interoperability of UWB devices and the co-existence with other radio technologies requires worldwide regulations and standardization. Sect. 2.2.1 discusses the regionally dependent frequency ranges, power limits, and emission masks for UWB communications to confine interference to other systems. It is followed in Sect. 2.2.2 by an analysis of the IEEE 802.15.4 standards containing the definitions of the UWB physical layer and media access control (MAC) layer to enable low-power location-enabled wireless personal area networks (WPANs).

2.2.1 Worldwide regulations

Due to the large bandwidth of UWB systems and the overlap of the spectrum with other radio technologies, it is required to strictly regulate the utilization of UWB signals. Hence, regulatory bodies worldwide have released rules to govern the frequency ranges, maximum emission levels, as well as the area of application, i.e., indoor/outdoor, portable/fixed-installed, for UWB [245]. As discussed in Sect. 2.1, in 2002 the US Federal Communications Commission (FCC) was the first agency worldwide to publish national regulations for the use of UWB technology [53] and allocated the 3.1-10.6 GHz frequency band for unlicensed operation of UWB devices. Fig. 2.1a shows the emission limits of the power spectral density (PSD) in the United States (US), where the grey area marks the in-band emission limits and the rest of the figure the out-of-band limits. The maximum equivalent isotropically radiated power (EIRP) in any direction may not exceed -41.3 dBm/MHz over a 1 ms integration time, which is identical to the FCC Part 15 limit for unintentional radiators (dot-dashed purple line) [82]. Hence, the total transmitted power for the FCC is restricted to 0.56 mW. In contrast to indoor environments (solid blue line), UWB-based outdoor applications are required to operate without a fixed infrastructure and, furthermore, the out-of-band emissions are more restricted (dashed orange line) [181]. Following the United States, UWB regulations have been released by numerous other regions and countries including China, Japan, Russia, Canada, South Korea, and Australia [37]. In early 2007 also the European

Commission (EC) followed studies and recommendations by the Electronic Communications Committee (ECC) of the European Conference of Postal and Telecommunications Administrations (CEPT) and made the final decision to publish UWB regulations in Europe [50]. Particularly, the ETSI EN 302 065-1 standard defines the emission limit for UWB applications² (see Fig. 2.1b) [47]. It is evident that the European regulations are - although heavily influenced by the FCC's decision - more restrictive than the US counterpart. Besides leaving out the frequency band between 4.8 to 6.0 GHz entirely, the bands below 4.8 GHz and above 8.5 GHz require mitigation techniques such as low duty cycle and/or detect and avoid. Similarly to the United States, in Europe fixed UWB infrastructure is restricted to indoor environments.

2.2.2 IEEE 802.15.4 standardization

Following the FCC release to allocate the 3.1-10.6 GHz band for unlicensed operation of UWB devices, major chip manufacturers including Texas Instruments and Intel expressed interest in UWB technology. This led to the formation of the IEEE 802.15.3a task group. However, the members could not agree on the underlying physical layer, with OFDM [13] and DS-SS [159] under discussion, eventually causing the dissolution of the task group in 2006. Meanwhile, the IEEE 802.15.4a task group recognized the potential of UWB for accurate ranging and published in 2007 an amendment to support low-rate low-power WPAN with localization capabilities using ultra-wideband technology [35, 92], which was finally merged into the IEEE 802.15.4-2011 standard [93]. The latter followed the amendment IEEE 802.15.4f-2012 incorporating an additional UWB PHY definition, namely low-rate pulse repetition frequency (LRP) UWB. To differ between LRP and the one based on the IEEE 802.15.4a-2007 amendment, the latter is accordingly named high-rate pulse repetition frequency (HRP) in the latest version of the IEEE 802.15.4 standard published in 2015 [95]. Since the main UWB chip manufacturers such as Decawave, NXP, and Apple aim for the HRP UWB PHY definition, this thesis focuses on the latter with Sect. 2.2.2.1 covering its specification. Still, for the sake of completeness, Appendix A.1 discusses the LRP definitions and Appendix A.2 presents details of the drafted IEEE 802.15.4z standard enhancing existing UWB PHYs with increased integrity and security support.

2.2.2.1 HRP UWB PHY specification (based on IEEE 802.15.4a)

The popularity of the IEEE 802.15.4 standard was mainly driven by the provision of low-cost and low-power wireless technology for the deployment of wireless sensor networks [160]. However, it was evident that enabling accurate ranging and positioning capabilities could expand the field of applications and market opportunities significantly [103]. To this end, the IEEE 802.15.4a task group was formed with the aim to define alternative PHYs supporting, among others, accurate and precise distance estimation and enhanced robustness. Besides a chirp spread spectrum (CSS) based physical layer, the task group eventually framed the first UWB-based standard for low-rate low-power wireless communications [35, 92]. The standardization process triggered semiconductor companies to resume the design of UWB hardware with the Decawave DW1000 being the first low-cost IEEE 802.15.4-compliant UWB transceiver available on the market [39].

²In fact, the EN 302 065-1 defines the requirements for generic UWB applications. The requirements for UWB location tracking are defined in EN 302 065-2 [48] and for road and rail vehicles in EN 302 065-3 [49], respectively.




Figure 2.2: UWB PHY frame structure. Adapted from [67].

UWB band allocation. The IEEE 802.15.4 standard allocates three bands for HRP UWB communications: the sub-gigahertz band (250-750 MHz), the low band (3244-4742 MHz), and the high band (5944-10234 MHz). A standard-compliant UWB device should support at least one of the three specified bands [95]. Within these three bands, the standard suggests 16 different channels for UWB, where one for each frequency band is mandatory (these are channel 0, 3, and 9 for sub-gigahertz, low band, and high band, respectively). Twelve channels support a bandwidth of 499.2 MHz and four make use of an increased bandwidth of up to 1354.97 MHz.

Frame format and modulation. As discussed in Sect. 2.1.2 and illustrated in Fig. 2.2, the IEEE 802.15.4 standard employs time-hopping impulse radio (TH-IR) UWB, thus, a standard-compliant UWB frame is represented by a sequence of pulses with very short duration and each frame contains two main blocks: a SHR and a data portion.

Synchronization header. The SHR is composed of a preamble (used for AGC tuning, signal detection, frame synchronization as well as channel estimation) and a start-of-frame delimiter (SFD) indicating the end of the SHR and beginning of the data portion. Each code symbol in the SHR is drawn from a ternary alphabet $\{-1,0,1\}$ representing a negative pulse, no pulse, or a positive pulse, respectively, and is transmitted in slots of fixed duration $t_p = 1/f_p$ where f_p corresponds to the highest frequency at which standard-compliant UWB transceivers are allowed to transmit pulses, i.e., 499.2 MHz. The duration of the SHR depends on two settings: the *pulse repetition frequency (PRF)* and the *preamble symbol repetitions (PSRs)*. The latter defines the number of preamble symbols sent in the SHR, with the standard suggesting four possible PSRs: 16, 64, 1024, and 4096. Thus, the transmitter sends repeatedly preamble symbols with a duration

of $t_{ps} = N/f_p$ each, where N is dependent on the PRF defining the number of transmitted pulses within a certain period: N is 496 or 508 for a (mean) PRF of 15.6 and 62.4 MHz, respectively. The sequence of the pulses within a preamble symbol is defined by the preamble code. The latter is either a length-31 or length-127 ternary sequence (excluding inserted zeros to match N) and features the perfect periodic autocorrelation property, i.e., the autocorrelation function consists of a single non-zero element [56]. This allows to estimate a CIR by cross-correlating the received signal with the preamble code [238]. A standard-compliant UWB device shall support at least two length-31 preamble codes per channel, whereby the combination of a channel and a preamble code is termed *complex channel*. The low cross correlation between preamble codes of the same channel enables, in theory, the independent operation of complex channels in adjacent networks. Similar to the preamble, the SFD consists of multiple preamble symbols.

Data portion. The data portion of a UWB frame consists of a physical layer header (PHR) as well as a payload. It is, in contrast to the synchronization header, sent as bursts using BPM/BPSK modulation, i.e., the information is encoded in the position of the burst (burst position modulation or BPM) and in the phase of the burst (binary-phase shift keying or BPSK). This enables both coherent and noncoherent receivers to be supported by the standard. Noncoherent receivers such as the energy detector and autocorrelation receiver multiply the received signal with itself followed by an integration and thus can only obtain the envelope of the signal [220]. This causes the loss of phase information and only the information encoded in the position of the burst is serviceable. Coherent receivers such as the rake receiver, instead, are able to exploit the phase information of the received signal and hence also the information encoded in the phase of the burst can be used to provide higher coding gain and to improve robustness [238]. The physical layer header is 19 bits long and contains information required to successfully decode the packet such as the length of the payload and the data rate. The latter can be 110 kbps, 851 kbps, 6.81 Mbps, and 27.24 Mbps, where 851 kbps is mandatory and the others are optional. Varying data rates are achieved through the use of variable-length bursts [103].

2.3 Time-based Ranging Methods using UWB

Estimating the position of a mobile device or tag is commonly a two-way process: In the *acquisition phase* information related to the position of the tag is acquired and in the *computation phase* this position-related information as well as prior knowledge such as the position of the anchors or floor plans are utilized to compute the position of the tag. Although UWB-based systems that employ the position-related angle of arrival (AoA) or received signal strength (RSS) information exist [204, 239], the most common methods are exploiting the accurate and precise timestamping capabilities of UWB transceivers to estimate the time of flight (ToF) $\tau_f = d/c$ of a signal, with d being the distance between two nodes and c the speed of electromagnetic waves in air [35, 166]. These time-based techniques differ by what means they acquire the distance information required to estimate the position of a mobile device. In one-way time of arrival (ToA) ranging, the distance between two devices is directly derived from a single message (Sect. 2.3.1), whereas two-way ToA ranging employs multiple packet exchanges (Sect. 2.3.3). In the third method, time difference of arrival (TDoA) between multiple anchors is utilized to relax the requirements on the clock synchronization between tags (Sect. 2.3.2).




Figure 2.3: Positioning based on (a) ToA and (b) TDoA ranging scheme.

2.3.1 One-way ToA ranging

To estimate the distance between two transceivers in one-way ToA ranging, a transmitter sends a message to the receiver containing the transmission timestamp t_0 . At timestamp t_1 the packet arrives at the receiving node, which estimates the time of arrival t_1 by detecting the first arriving signal path or leading edge. Based on the acquired information, the receiver is capable of estimating the time of flight $\tau_f = t_1 - t_0$ and distance d to the transmitter. However, this requires a tight clock synchronization between participating nodes, which results in additional overhead [223], complexity [227] and ranging error [35]. Each distance estimate to a node with known position, i.e., an anchor or a previously localized tag, defines a range circle with radius d , where at least three distance estimates are required to estimate the 2-D position located at the intersection of the obtained range circles using trilateration (see Fig. 2.3a) [235]. Due to noisy range estimates, it is highly unlikely to observe a single intersection of the range circles. Hence, an optimization algorithm estimating the most probable position of the tag is required [229].

2.3.2 Time difference of arrival (TDoA)

TDoA-based systems do not use the absolute ToF between anchor nodes and tags, but rather the time difference of arrival between multiple anchors. This addresses a drawback of one-way ToA ranging, namely, the need for synchronization between anchor clocks and tag clock. Instead, in TDoA only the anchor clocks need to be synchronized wirelessly or via a wired backbone connection [96, 235]. In contrast to the range circle in ToA-based systems, estimating the TDoA corresponds geometrically to a hyperbolic function, as illustrated in Fig. 2.3b [76]. To estimate the tag's position in a coplanar scenario, at least two hyperbolas and TDoA measurements, respectively, are necessary. As discussed in Sect. 1.2, it is differentiated between *tag-initiated* and *anchor-initiated* TDoA. In the latter, each of the synchronized anchors broadcasts sequentially messages to the mobile devices, which estimate the TDoA between different anchor pairs to derive the required position-related information. In *tag-initiated* TDoA the tag broadcasts a blink or reference signal to the surrounding anchors, which share the ToA to estimate the TDoA [35].

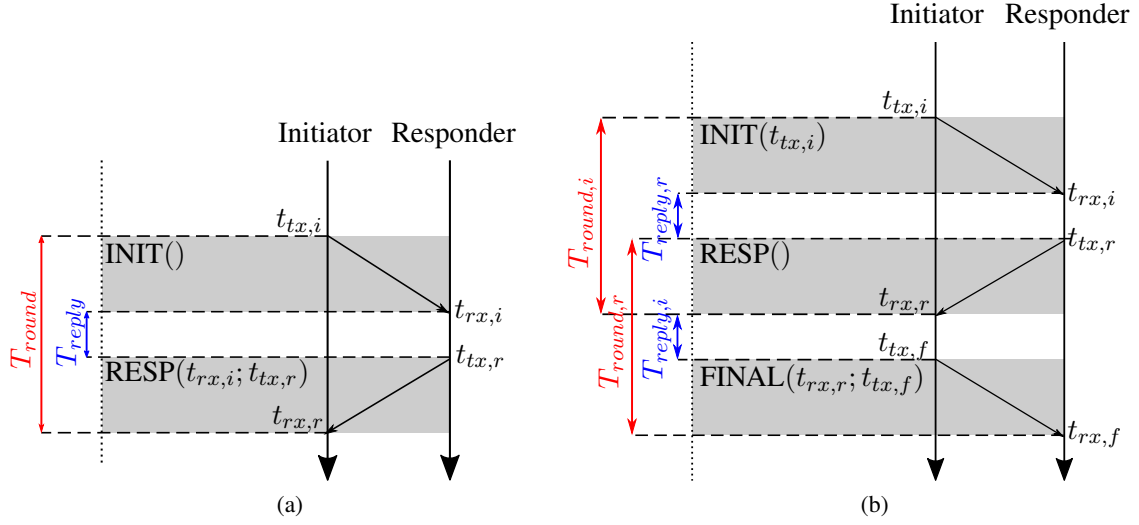


Figure 2.4: Comparison of (a) single-sided two-way ranging with two packet exchanges and (b) double-sided two-way ranging (DS-TWR) with three transmitted messages.

2.3.3 Two-way ToA ranging

Instead of estimating the time of flight based on a single message, in two-way ToA ranging the round-trip time (RTT) of a packet is measured, avoiding the necessity of a common time base. However, it requires multiple packet exchanges between the participating nodes, which are typically referred to as *initiator* and *responder* (see Fig. 2.4). The role assignment is subject to the structure of the positioning system, i.e., client-based, server-based or cooperative (see Sect. 1.1.1). Depending on the number of packets sent, it is differed between single-sided two-way ranging (SS-TWR) and double-sided two-way ranging (DS-TWR).

SS-TWR. The operation of SS-TWR is presented in Fig. 2.4a. To initiate the distance estimation, the initiator sends an initialization message (*INIT*) to the responder. The latter responds after T_{reply} with a *RESP* message containing the reception timestamp $t_{rx,i}$ and transmission time stamp $t_{tx,r}$ to complete the packet exchange. After determining the reception time of the *RESP* message $t_{rx,r}$, the initiator is capable of estimating the time of flight $\hat{\tau}_{f,ss}$ as

$$\hat{\tau}_{f,ss} = \frac{\hat{T}_{round} - \hat{T}_{reply}}{2} = \frac{(\hat{t}_{rx,r} - \hat{t}_{tx,i}) - (\hat{t}_{tx,r} - \hat{t}_{rx,i})}{2}. \quad (2.6)$$

It is hence necessary that the UWB transmitter is capable of embedding the timestamps used to estimate T_{reply} in the *RESP* message, which requires an accurate prediction of $t_{tx,r}$. Transceivers such as the Decawave DW1000 support this feature. If this is not the case, an additional message is necessary [96]. Since T_{round} and T_{reply} are measured independently by the internal clock of the initiator and responder, respectively, no time synchronization is required. However, the individual clock drift e_i and e_r causes an estimation error, as discussed in Appendix B.1.

DS-TWR. In double-sided two-way ranging, a second round-trip time measurement is employed; thus, an additional *FINAL* message is required, as illustrated in Fig. 2.4b. In this configuration the *RESP* message is used to finalize the first RTT measurement and to initialize a second one. Hence analogously to Eq. 2.6, two definitions for the ToF are derived as:

$$\tau_f = \frac{1}{2}(T_{round,i} - T_{reply,r}) \quad \tau_f = \frac{1}{2}(T_{round,r} - T_{reply,i}), \quad (2.7)$$

combining these definitions results in the true ToF $\tau_{f,ds}$ using DS-TWR

$$\tau_{f,ds} = \frac{1}{4}(T_{round,i} - T_{reply,r} + T_{round,r} - T_{reply,i}). \quad (2.8)$$

Please note that in the scheme shown in Fig. 2.4b, the time of flight $\hat{\tau}_{f,ds}$ is estimated at the responder. In case the distance information is required at the initiator, an additional message has to be transmitted. As derived in Appendix B.2, the estimation error due to clock drifts $\epsilon_{c,ds}$ is minimal when the reply times at initiator and responder, respectively, are equal. However, if the application or transceiver cannot ensure that the reply times remain similar, an alternative approach was suggested by Neiryneck et al. [144]. Instead of summing up the definitions for the time of flight in Eq. 2.7, they are multiplied, which results in:

$$\begin{aligned} T_{round,i}T_{round,r} &= (2\tau_{f,ds} + T_{reply,r})(2\tau_{f,ds} + T_{reply,i}) \\ &= 2\tau_{f,ds} \underbrace{(2\tau_{f,ds} + T_{reply,r} + T_{reply,i})}_{=T_{round,i}} + T_{reply,i}T_{reply,r}. \end{aligned} \quad (2.9)$$

Rearranging Eq. 2.9 yields an alternative equation for the ToF $\tau_{f,ds}$ in DS-TWR

$$\tau_{f,ds} = \frac{T_{round,i}T_{round,r} - T_{reply,i}T_{reply,r}}{2(T_{round,i} + T_{reply,i})}. \quad (2.10)$$

Please note that Eq. 2.10 implies that $T_{round,i} + T_{reply,i} = T_{round,r} + T_{reply,r}$ [144]. To differ between the DS-TWR approaches defined in Eq. 2.8 and Eq. 2.10, respectively, the first one is referred to as symmetric DS-TWR and the second one as asymmetric DS-TWR. According to the analysis in Appendix B.2, the resulting estimation error due to clock uncertainties in asymmetric DS-TWR solely depends on the ToF and not on the reply times. It is thus the preferred DS-TWR method used in the remainder of this thesis.

2.4 Channel impulse response (CIR)

The time-based ranging methods discussed in Sect. 2.3 require to measure the time of arrival (ToA) of an incoming packet. To this end, a coherent UWB transceiver such as the Decawave DW1000 estimates the channel impulse response (CIR) by cross-correlating the received signal with the preamble code featuring the perfect periodic autocorrelation property [238] (see Sect. 2.2.2.1) and employs the estimated CIR to precisely determine the ToA by detecting the first path or leading edge of the CIR [39]. Fig. 2.5a exemplarily shows an estimated channel impulse response acquired in a rectangular room. It contains information about the first path




Figure 2.5: Channel impulse response (CIR) (a) acquired at the receiver (RX) in a rectangular room (b). The multipath components (MPCs) are denoted with $\tau_1 \dots \tau_4$. Adapted from [70].

component (τ_0) as well as the multipath components (MPCs) ($\tau_1 \dots \tau_4$), i.e., signals arriving at the receiver (RX) via different paths by being reflected at or diffracted by walls and other objects in the environment. The path delays τ_k of the signal components relate to the path lengths via $d_k = \tau_k \cdot c$, with c being the speed of light in air. Note that, in LOS conditions, the first path corresponds to the direct path between transmitter (TX) and receiver (RX). In NLOS conditions, instead, the first path might origin from a reflected or diffracted signal. According to the tapped delay line model, the CIR describing the multipath propagation can be modeled as [69, 139]

$$h(t) = \sum_{k=0}^K \alpha_k(t) \delta(t - \tau_k) + \nu(t), \quad (2.11)$$

where α_k and τ_k denote the complex-valued amplitude and path delay of the k^{th} MPC, respectively, and $\delta(t)$ the Dirac delta function. Theoretically, infinitely many signal paths might affect the propagation channel between transmitter and receiver. However, in the remainder of this thesis solely K MPCs originating at reflective surfaces are considered. The diffuse multipath $\nu(t)$ originating at small objects and rough surfaces is treated as additive noise [113]. Resolving MPCs, as illustrated in Fig. 2.5a, is a unique feature of UWB transceivers enabled by their high bandwidth. In contrast, narrowband receivers cannot distinguish between multipath components, instead, MPCs will overlap with a high probability and interfere constructively or destructively with each other causing small-scale fading. The number of resolvable MPCs is given by the maximum excess delay τ_{max} and the bandwidth B as τ_{max}/B [139].

Within this thesis, the CIR is exploited (i) to gain knowledge about the environmental conditions of a communication channel to guide a PHY parameter adaptation (Chapter 3), (ii) to minimize the required infrastructure of a positioning system to a single anchor (Chapter 4), and (iii) to acquire the response signals of multiple transmitters simultaneously (Chapter 5).

Characterization and Runtime Adaptation of UWB PHY Settings

The ranging schemes discussed in Sect. 2.3 require to precisely estimate the time of arrival of packets and to share the acquired timestamps in a wireless network. Thus, providing robust and efficient communication links in dynamic and harsh indoor environments is inevitable. To this end, standard-compatible UWB transceivers provide several PHY settings to tune the communication performance while maintaining energy-efficient operation. Within this chapter, we study the characteristics of these parameters and present a UWB link state indicator. The latter employs the estimated CIR to accurately measure the link quality and to extract relevant information about the surrounding environment. Based on the characterization and the link state indicator, we present a scheme that adapts the UWB PHY settings at runtime (see Fig. 3.1).

Related work. Tuning the radio sensitivity and energy consumption requires a deep understanding and characterization of the relevant parameters. However, recent research activities on UWB technology largely ignored the PHY parameters and, instead, focused on algorithms to increase the positioning accuracy [4]; for instance by employing inertial sensors [110] and exploiting antenna diversity [106]. We bridge this gap and characterize the impact of PHY parameters on the communication performance. Among the observed parameters is the preamble symbol repetition (PSR) and pulse repetition frequency (PRF) as well as data rate, bandwidth, and carrier frequency. Even though previous works have investigated the impact of various settings on the ranging accuracy [62, 137, 231], none of them experimentally quantified the impact of the settings on the robustness and energy efficiency. This experimental characterization of the PHY parameters lays the foundation to adapt them at runtime and to select the most appropriate setting for a certain link quality and environmental condition. This is of great importance, as mobile transceivers are typically designed for the worst-case channel conditions and, hence, many users may experience unnecessarily high signal quality [5], which hinders an efficient use of system resources. Therefore, in wireless communications, adaptation techniques to optimize radio transmission and reception are employed. Indeed, it is common to adapt, among others,




Figure 3.1: Sketch of the presented adaptation scheme to derive a set of PHY parameters at runtime maintaining robust and energy-efficient UWB communication. Adapted from [67].

the transmit power, data rate, and modulation scheme in cellular systems such as GSM [43], EDGE [57], UMTS [124, 149], and LTE [111, 179] to increase the operating time of batteries [139] and the spectral efficiency [143]. Similarly, in low-power wireless networks several works have tuned PHY and MAC protocol parameters to optimize link reliability and energy efficiency [42, 173, 180, 241]. In contrast, UWB systems make still use of static PHY parameter, i.e., the settings are fixed at the deployment [194], which makes the systems unable to cope with dynamic environments. Thus, to the best of the author’s knowledge, the proposed scheme represents the first runtime adaptation of UWB parameters to improve communication performance. To trigger such an adaptation of the parameters, a continuous estimation of the link quality is required [143]. Ideally, one would have information about the bit error rate (BER) or PRR available. However, reliable estimates of these parameters require several measurements, which would cause delays in the adaptation. Hence, other types of link quality measurements that are related to the PRR, such as SNR and signal-to-interference-plus-noise ratio (SINR), are preferred [5, 10]. Still, due their low complexity and availability in narrowband transceivers, RSS and link quality indicators (LQIs) are most commonly used for low-power wireless networks [8, 183]. However, the link quality information derived from low-cost low-power UWB transceivers is largely unexplored as concepts applicable for narrowband transceivers cannot be easily applied to UWB due to fundamental differences in the PHY, such as the low power spectral density [153, 238]. Instead, the high bandwidth in UWB facilitates new opportunities such as the estimation of the CIR (see Sect. 2.4). The latter contains information about the multipath propagation and hence about the environment, which supports the selection of a new parameter set in the case of fluctuating channel conditions. Indeed, utilizing the CIR information allows to identify NLOS conditions and destructive interference. We exploit these capabilities to present the first environment-aware UWB link state indicator. By continuously analyzing the estimated CIR, the link state indicator provides real-time information about the link quality and environmental conditions. In case the link quality drops below a threshold, the adaptation logic decides for a new PHY configuration. Such a threshold-based controller is a common technique also in adaptation algorithms using other wireless technologies [90].

Structure. The remainder of this chapter is structured as follows. In Sect. 3.1 we characterize the impact of PHY parameters on the robustness and energy efficiency of UWB communication. Based on this experimental characterization as well as the application requirements, a ranking of PHY settings is derived, which privileges parameters assisting to increase the communication performance while meeting the requirements (see Fig. 3.1). Following, Sect. 3.2 presents our proposed UWB link state indicator consisting of a received signal power estimator (Sect. 3.2.1) and an environmental state estimator (Sect. 3.2.2). The characterization of the PHY parameters as well as the link state indicator are integral components of the presented adaptation scheme to derive an optimal set of PHY parameters at runtime, which is discussed in Sect. 3.3. Evaluations in Sect. 3.4 illustrate the effectiveness of the proposed adaptation scheme.

3.1 Characterization of UWB PHY Parameters

This section discusses the experimental characterization of the PHY parameters provided by IEEE 802.15.4 compatible UWB transceivers. Sect. 3.1.1 defines the experimental setup to perform the measurements in a reproducible manner as well as the employed hardware. Sect. 3.1.2 presents the impact on the robustness and energy efficiency of UWB communication individually for each of the parameters: preamble symbol repetition (PSR) (Sect. 3.1.2.1), pulse repetition frequency (PRF) (Sect. 3.1.2.2), data rate (Sect. 3.1.2.3), channel and bandwidth (Sect. 3.1.2.4).

3.1.1 Experimental setup

To experimentally characterize the UWB physical layer parameters in terms of energy efficiency and reliability, we use the IEEE 802.15.4-compliant UWB transceiver Decawave DW1000 [38]. In particular, we employ the EVB1000 board consisting of a DW1000 radio, an ARM Cortex M3 processor, and an omnidirectional UWB antenna. The DW1000 is a coherent transceiver, thus, it exploits the phase as well as position of the bursts and it operates on six channels with center frequencies between 3.5 to 6.5 GHz and a bandwidth of 499.2 or 900 MHz¹ (see Table 3.1).

Wired setup. In order to maximize reproducibility and generalizability by avoiding environmental influences, temporal fluctuations, and reflecting obstacles, we connect two EVB1000 boards via SMA cables with a programmable attenuator in between. The latter allows to control the level of attenuation in 0.25 dB steps and hence to finely explore the impact of each parameter on the transceiver sensitivity, i.e., a measure for link reliability. For each attenuation step, the two EVB1000 boards carry out 1000 asymmetric DS-TWR operations (see Sect. 2.3.3).

Measuring sensitivity and energy efficiency. Initially, the DW1000 transceivers operate with a pre-defined default configuration (see Table 3.2). Subsequently each PHY parameter is varied individually to quantify its impact on the link reliability. For each 0.25 dB attenuation step we measure the PRR and, additionally, we derive a HRR as the ratio between the number of packets for which at least the SHR was successfully decoded and the number of packets sent. The average current consumption for transmission and reception of SHR as well as payload is measured using a Keysight MSO-S 254A oscilloscope.

¹The restriction to a maximum bandwidth of 900 MHz is due to an internal limitation of the Decawave DW1000 and is not in agreement with the standardization discussed in Sect. 2.2.2.1.

Table 3.1: DW1000 supported channels.

Ch.	f_c [MHz]	Bandw. [MHz]
1	3494.4	499.2
2	3993.6	499.2
3	4492.8	499.2
4	3993.6	1331.2 (900)
5	6489.6	499.2
7	6489.6	1081.6 (900)

Table 3.2: Default configuration.

PHY Setting	Value
Channel	5
Pulse repetition frequency	16 MHz
Preamble symbol repetition	128
Data rate	6.8 Mbps
Payload size	12 Bytes

3.1.2 Impact of PHY parameters on robustness and energy efficiency

As discussed in Sect. 2.2.2.1, an IEEE 802.15.4-compatible frame is divided in two main blocks, namely, a synchronization header (SHR) and a data portion consisting of a physical layer header (PHR) and payload. While the SHR consists of a sequence of single pulses, the data portion is constructed by UWB bursts using BPM/BPSK modulation, as illustrated in Fig. 2.2. Due to the differences in the modulation schemes, we introduce HRR to allow an isolated analysis of the parameter's impact on SHR and the data portion, respectively. The impact of preamble symbol repetition is discussed in Sect. 3.1.2.1 followed by an analysis of the pulse repetition frequency in terms of energy efficiency and link reliability in Sect. 3.1.2.2. This section concludes with an investigation of the data rate as the most effective parameter to increase the sensitivity of the data transmission (Sect. 3.1.2.3) and of the carrier frequency as well as bandwidth (Sect. 3.1.2.4).

3.1.2.1 Preamble symbol repetition (PSR)

The preamble is an integral part of the synchronization header and its main purpose is to synchronize to an incoming frame and to derive the channel impulse response as discussed in Sect. 2.2.2.1 and Sect. 2.4. The preamble symbol repetition (PSR) defines the number of preamble symbols sent within one preamble. Since the DW1000 transceiver does not allow to select the lowest standard-compatible PSR of 16, the characterization focuses on PSR values of 64, 128, 1024, and 4096. Fig. 3.2a shows the PRR and HRR as a function of the attenuation level for the four different PSR values. Apparently, the PSR does not impact the packet reception rate, as it solely tunes the SHR. Besides increasing the ranging accuracy [166], a higher PSR results in a correct SHR reception at a significantly lower SNR, e.g., increasing the PSR from 64 to 4096 gives an additional 6 dB gain. However, a higher PSR results in longer packets and thus decreases the energy efficiency as shown in Table 3.3. The latter represents the energy consumption of transmitting and receiving a UWB packet with a 12 bytes payload. The values are derived from the measured average current consumption times the packet's over-the-air time and the supply voltage (3.3 V). Exemplarily, increasing the PSR from 64 to 4096 results approximately in a 39 times higher total energy consumption, thus, as long as the propagation environment tolerates it, a low PSR is preferable. Note that Table 3.3 takes the active transmission and reception into account, however, also listening for a preamble results in a continuous current draw of about 130 mA. Additionally, it is evident in Fig. 3.2a that the SHR can be successfully detected at a higher attenuation level than the one at which the first packets are lost. This observation enables

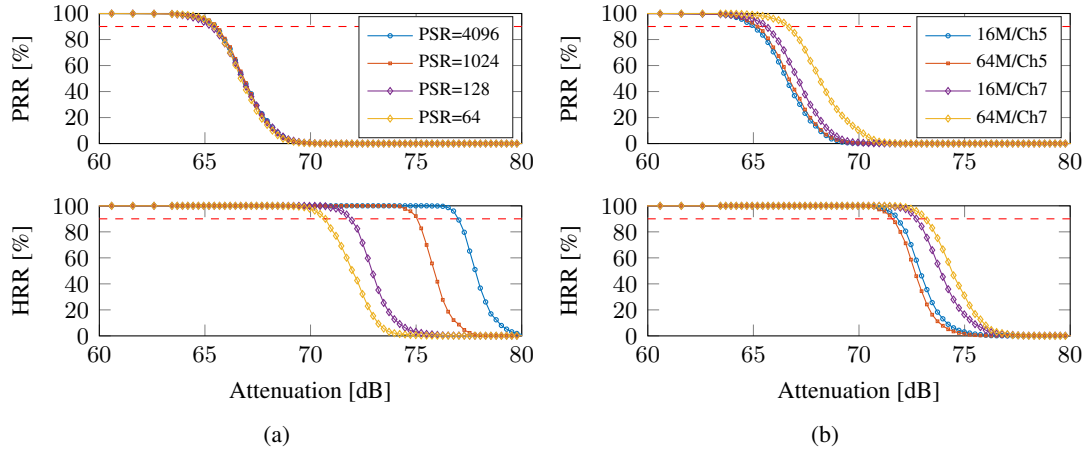


Figure 3.2: Packet reception rate (PRR) and header reception rate (HRR) as a function of the attenuation level as well as in dependence of the PHY parameters (a) preamble symbol repetition (PSR) and (b) pulse repetition frequency (PRF). Adapted from [67].

two appealing features, which we discuss in more detail in [67]. First, even if the payload is corrupted, the receiver might still be capable of extracting channel information from the SHR. This enables the anticipation of a degrading channel by analyzing the environmental state as described in Sect. 3.2.2. Second, the successful reception of the SHR and preamble might act as a robust binary acknowledgement (ACK) message to confirm the reception of a packet and to agree upon new channel settings as exploited in the adaptation scheme presented in Sect. 3.3.

3.1.2.2 Pulse repetition frequency (PRF)

The PRF defines the number of UWB pulses sent within a certain period and influences the synchronization header as well as the data portion. The highest possible (peak) PRF is characterized in the standard as 499.2 MHz, which defines the slot duration in the SHR (see Fig. 2.2) and affects the burst duration in a data symbol. However, the PRF is typically stated as the mean PRF defined as the total number of pulses transmitted during a (preamble or data) symbol period divided by the duration of the symbol [95]. The IEEE 802.15.4 standard defines three (mean) PRFs, namely, 3.90 MHz, 15.6 MHz, and 62.4 MHz. Due to the missing support of 3.90 MHz in the DW1000, the differences for a PRF of 15.6 MHz and 62.4 MHz, respectively, are investigated. In addition to the transmission of more UWB pulses in a defined period, a higher PRF occurs in a longer preamble symbol (see Sect. 2.2.2.1), effectively increasing the energy consumption by about 9 % (see Table 3.3). However, sending more pulses typically results in a more robust link as shown in Fig. 3.2b. In general, lower PRFs are better suited for environments with high delay spread, i.e., multipath-rich conditions, instead, in low delay spread environments and when using coherent receivers a high PRF is preferable [166].

Table 3.3: Energy consumption for the transmission and reception of a packet depending on the employed PHY settings and broken down into SHR and payload. Adapted from [67].

PHY Setting	Transmission		Reception		Transm. + Recept.	
	SHR [μJ]	Payl. [μJ]	SHR [μJ]	Payl. [μJ]	Tot. [μJ]	Δ [%]
Default	31.23	6.198	66.34	18.03	121.8	-
PSR = 4096	2918	6.198	2001.9	18.03	2969	2337
PSR = 1024	237.0	6.198	503.4	18.03	764.6	527.8
PSR = 64	16.54	6.198	35.12	18.03	75.9	-37.7
Channel 1	28.39	5.177	59.91	16.03	109.5	-10.10
Channel 2	28.96	5.329	63.35	17.09	114.7	-5.80
Channel 3	28.37	5.323	62.79	17.16	113.6	-6.701
Channel 4	33.35	5.710	62.16	17.04	118.3	-2.901
Channel 7	30.83	6.233	67.72	19.06	123.8	1.676
PRF = 62.4 MHz	38.65	6.567	69.11	18.03	132.3	8.648
DR = 850 kbps	31.45	22.71	67.78	66.29	188.2	54.55
DR = 110 kbps	44.46	176.2	92.48	554.7	867.8	612.5

3.1.2.3 Data rate

The data rate solely influences the packet reception rate, as confirmed in Fig. 3.3a. The investigated data rates are restricted to 110 kbps, 850 kbps, and 6.8 Mbps, as the highest data rate of 27 Mbps is not supported by the DW1000. Apparently, reducing the data rate results in the highest gain in terms of robustness: indeed, using a data rate of 850 and 110 kbps raises the link margin by 2.9 and 5.5 dB, respectively. However, using a lower data rate in an adaptation scheme such as the one presented in Sect. 3.3 should be chosen wisely, as a lower data rate increases the packet's over-the-air time and hence the energy consumption. Indeed, transmitting and receiving a payload of 12 bytes (ignoring the SHR) with a data rate of 850 and 110 kbps consumes respectively up to 3.7 and 31 times more energy compared to the default settings.

3.1.2.4 Channel and bandwidth

This section discusses the impact of different channels, i.e., center frequencies f_c (see Table 3.1) and bandwidth on reliability and energy costs. However, evaluating the dependency of the reliability on f_c in a cable-based setup such as the one described in Sect. 3.1.1 is challenging as one has to take the free-space path loss or other UWB channel models into account [138, 140]. Still, when employing the described setup, Fig. 3.3b shows that lowering f_c (Ch. 4 (4 GHz) vs. Ch. 7 (6.5 GHz) and Ch. 2 (4 GHz) vs. Ch. 5 (6.5 GHz), see Table 3.2) increases the robustness of both SHR and data portion. This observation would be emphasized by employing the Friis equation as the path loss is reduced at lower frequencies [139]. Similarly, increasing the bandwidth improves the robustness due to the higher total transmit power (Ch. 2 (500 MHz) vs. Ch. 4 (900 MHz) and Ch. 5 (500 MHz) vs. Ch. 7 (900 MHz)). In terms of energy consumption, it appears that the channels in the low band (below 5 GHz) are slightly more efficient, whilst the channels in the high band (above 5 GHz) consume more energy, as shown in Table 3.3.

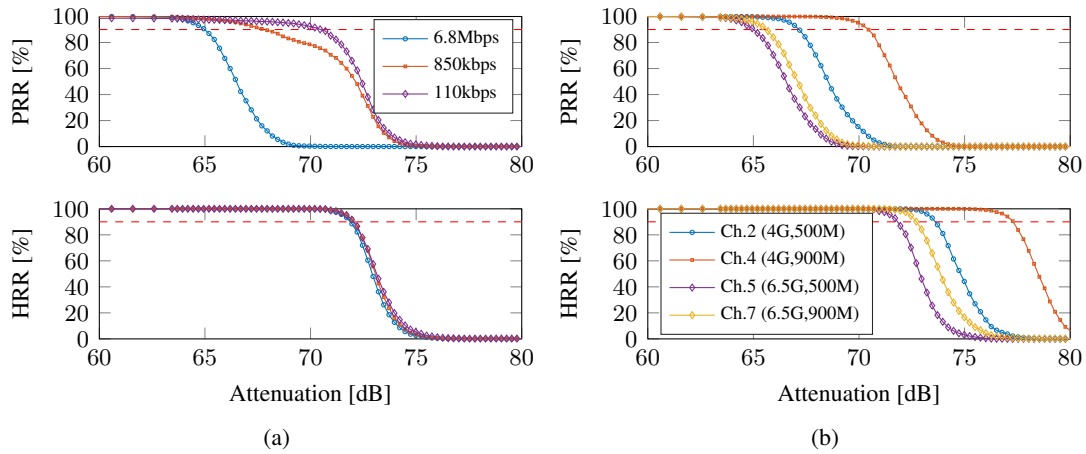


Figure 3.3: Packet reception rate (PRR) and header reception rate (HRR) as a function of the attenuation level as well as (a) data rate and (b) carrier frequency/bandwidth. Adapted from [67].

3.1.3 Summary of PHY parameter characterization

The characterization of the PHY settings discussed in Sect. 3.1.2, allows to quantify the impact of the parameters in terms of reliability and energy efficiency as well as to assess which PHY settings are most effective in increasing the robustness of a link while retaining the energy consumption at a minimum. Table 3.4 sums up the findings and illustrates the tradeoff between optimizing the link reliability and the energy efficiency. Exemplarily, a higher preamble symbol repetition ensures a significantly increased reliability in synchronizing to an incoming frame and in deriving the channel impulse response; however, the energy consumption is much higher due to longer packets. Thus, an adaptive scheme as the one presented in Sect. 3.3 should merely utilize a longer PSR as the last resort to optimize link reliability. That being said, in Chapter 4 a multipath-assisted positioning system is presented, which exploits the channel impulse response to derive delay and angular information of multipath components. In such systems it might be reasonable to use a higher PSR to increase the SNR of the CIR estimate. Similarly to the impact of the PSR on the reliability of the synchronization header, a lower data rate is the most effective option to increase the reliability of the data portion. However, this comes at the cost of a considerably higher energy consumption due to the longer packet's over-the-air time. If the energy budget is not limited as it might be the case for anchor nodes connected to the electrical grid, lowering the data rate should be the prioritized option to preserve a robust communication link. Thus, the order in which an adaptation scheme tunes the configuration is highly depending on the requirements and constraints. To this end, besides the characterization of the PHY settings, the adaptation scheme in Sect. 3.3 takes the application requirements into account to derive a ranking of the PHY settings. Changing the pulse repetition frequency, channel as well as bandwidth has a more moderate impact on the reliability and energy efficiency and is thus very valuable to fine tune the performance of the adaptation algorithm.

Table 3.4: Tradeoff between SHR/data reliability and energy efficiency. Adapted from [67].

Change in PHY setting	Reliability data portion	Reliability SHR	Energy efficiency
Higher PSR	-	↑↑↑↑	↓↓↓↓
Higher PRF	↑	↑	↓↓
Lower data rate	↑↑↑↑	-	↓↓↓↓
Lower carrier frequency	↑↑↑	↑↑↑	↑
Higher bandwidth	↑↑	↑↑	↓

3.2 Link State Indicator

To adapt PHY parameters as a countermeasure in case of a degrading channel, it is crucial to assess the link quality at runtime, as this highly affects the choice of which parameter to tune in order to maintain robust communication. To this end, we present a link state indicator that is also environment-aware meaning that besides indicating the received signal strength (Sect. 3.2.1), it predicts if an environmental impact such as destructive interference between the first path and a strong multipath component has caused the degradation of the link (Sect. 3.2.2). Indeed, although UWB technology provides a high bandwidth and hence high immunity to multipath fading, still, in multipath-rich indoor environments the communication performance is highly influenced by overlapping multipath components as we have shown experimentally in [67]. In the case of destructive interference, neither increasing the transmit power nor lowering the data rate achieves a significant improvement, instead, the only viable option is to adapt the center frequency f_c . Thus, the environmental awareness of the link state indicator assists in deriving the optimal adaptation strategy. Combining the link quality estimation based on the RSP and the environmental state indicator, results in the first UWB link state indicator presented in Fig. 3.4. The link quality is a continuous value between -128 and 0 given in dBm, instead, the environmental state are binary values indicating the occurrence of a certain environmental condition, such as the aforementioned destructive interference or NLOS propagation. Both the link quality and environmental state exploit the channel impulse response estimate provided by standard-compliant ultra-wideband transceivers (see Sect. 2.4).




Figure 3.4: UWB link state indicator employing environmental awareness.




Figure 3.5: Packet reception rate (PRR) as a function of the received signal power (RSP) when using default settings (blue) and a data rate of 110 kbps (orange). Adapted from [67].

3.2.1 Deriving the received signal power (RSP)

Since the channel impulse response contains information about the received LOS signal as well as reflected signal components, its integral can be used as an estimate of the RSP defined in dBm [39]. Nonetheless, the purpose of an adaptation scheme is not to guarantee a high RSP but a high PRR. Hence, initially one has to characterize the relation between the estimated received signal power and PRR. To this end, we performed a measurement campaign in a multipath-rich hallway at the University campus. Fig. 3.5 depicts the packet reception rate as a function of the estimated RSP (mean \pm standard deviation) when using the default settings (blue) and the lowest data rate of 110 kbps (orange), respectively. Each plotted data point corresponds to 1000 individual measurements and the interpolated dotted lines illustrate the trend of the relation between PRR and RSP. The figure clearly indicates that using a more robust setting, i.e., the lower data rate of 110 kbps (orange), requires less signal strength to achieve the same packet reception rate. In fact, the orange curve is shifted by about 5-6 dB compared to the blue line, which corresponds to the observations in the cable-based parameter characterization discussed in Sect. 3.1.2.3 and illustrated in Fig. 3.3a. These results motivate the applicability of the RSP to estimate the packet reception rate and hence the reliability of a communication link. In the evaluation discussed in Sect. 3.4, a link is considered as reliable if the PRR is above 90 % (indicated by the dashed red line in Fig. 3.5). Thus, the characterization of the relation between PRR and RSP allows to derive a threshold $P_{th,r}$ to trigger a parameter adaptation. In the illustrated example the threshold using the default settings is $P_{th,r} = -93.3 \text{ dBm}$.

3.2.2 Deriving the environmental state

Sect. 3.2.1 shows that the estimated received signal power is a valuable input to the adaptation scheme to trigger a parameter change. However, combining this insight with the environmental information provided by the CIR estimate could possibly lead to an improved adaptation strategy. For instance, we discuss two environmental states, which highly influence the performance of UWB communication and ultimately the performance of an indoor positioning system requiring to precisely estimate the ToA of a packet and to share these timestamps in a network.




Figure 3.6: Estimated channel impulse response in the presence (blue solid line) and absence (orange dashed line), respectively, of destructive interference. Adapted from [68].

LOS/NLOS. Time-based ranging methods as the ones discussed in Sect. 2.3 and used in the positioning systems presented in Chapter 4 and Chapter 5 inherently assume LOS signal propagation. Indeed, situations of blocked LOS lead to a positive range bias and substantial positioning errors [235]. Likewise, the communication performance is highly affected by NLOS conditions due to a reduced received signal strength. To address this issue and to mitigate the impact of NLOS scenarios, received signal features such as time of arrival, delay spread, and kurtosis [77, 78, 202] as well as machine learning techniques are utilized [83, 131, 226]. Such methods are capable of deriving a binary information about the presence of NLOS conditions. This knowledge can be included in the link state indicator shown in Fig. 3.4 to adjust the parameter settings, re-route packets via links providing line-of-sight, or steering directional antennas as the one presented in Sect. 4.3.2 towards the direction of non-blocked multipath components.

Destructive interference. Despite the high immunity to multipath fading, still, UWB communication suffers from overlapping MPCs due to the finite bandwidth of UWB transceivers. In [67] we show experimentally that such destructive interference of the LOS component with a strong reflection highly degrades a communication link, resulting in a PRR close to zero. As the communicating nodes move further apart from each other to escape the area of destructive interference, the PRR rises again and a high link quality is recovered. Detecting such areas of destructive interference in advance helps to avoid deep fades by triggering appropriate countermeasures. To this end, the estimated CIR provides valuable information, which can be exploited to detect such environmental changes and react accordingly. In fact, the best and most viable option is to change the center frequency, as it is highly unlikely that differing wavelengths experience destructive interference at the same position. Fig. 3.6 shows two recorded channel impulse responses, one was acquired while a highly reliable link was obtainable and the other one inside a deep fade due to destructive interference of the LOS component with a strong multipath signal. It is evident that the interference causes a significant drop of the LOS signal's amplitude α_0 . Instead, the amplitudes of the MPCs are less affected as indicated by the strong multipath component at about 25 ns in Fig. 3.6. Based on these observations, an efficient mechanism to detect destructive interference at runtime is proposed. The algorithm continuously tracks the ratio ER between the LOS signal energy and the energy of multipath components as follows:

$$ER = \frac{|\hat{\alpha}_0|^2}{\sum_{k=1}^K |\hat{\alpha}_k|^2}. \quad (3.1)$$

Please note that Eq. 3.1 requires knowledge of the number of significant multipath components K and their estimated amplitude $\hat{\alpha}_k$. This increases the complexity of the method since additional information is required, for instance a floor plan as utilized in the multipath-assisted positioning system presented in Chapter 4. Hence, the total signal energy following the LOS component is considered as multipath energy and even with this simplification reasonable results are achievable as shown in the evaluation in Sect. 3.4. Furthermore, we derive a threshold of the energy ratio ER_{th} empirically and set it to $ER_{th} = 1.5$ [67].

3.3 Runtime Adaptation of PHY Parameters

The overall structure of the presented adaptation scheme is illustrated in Fig. 3.1. Apart from the discussed experimental characterization and the link state indicator, it contains an *adaptation logic* block. The latter determines the runtime behavior of the adaptation scheme by exploiting the estimated received signal power to detect a degrading communication link. Subsequently, the adaptation logic decides on the ideal countermeasure taking the ranking of the PHY parameters, the application requirements, as well as the environmental state into account.

Detection of a degrading link. The link state indicator triggers a change of the PHY settings in case of a degrading channel, where the channel is considered as degrading as soon as the received signal power falls below the threshold $P_{th,r} = -93.3$ dBm derived in Sect. 3.2.1. However, to determine circumstances allowing to select a more energy-efficient setting, we derived a second threshold experimentally and set it to $P_{th,ee} = -88.8$ dBm as it showed a reasonable tradeoff between link robustness and energy efficiency. As soon as this threshold is exceeded, a more efficient setting is utilized in the UWB transceiver. In case the energy efficiency is a less stringent requirement of the application one can raise the threshold $P_{th,ee}$.

Ranking of UWB PHY parameters. The ranking of the UWB PHY parameters is defined by the experimental characterization discussed in Sect. 3.1. However, to keep the adaptation scheme generic, it additionally takes the application requirements into account. We demonstrate the impact of the requirements in Sect. 3.4 by defining two adaptation policies, namely, an *energy-conservative* and an *aggressive* approach. The latter maximizes the link reliability regardless of the energy expenditures, thus, it always selects the settings providing the largest link margin even if more energy-efficient configurations would be available. In the energy-conservative policy, instead, the most efficient settings are chosen while still maintaining a reasonable link reliability. According to Table 3.3, the ranking of the PHY settings employing an energy-conservative approach is: (1) lower carrier frequency, (2) higher bandwidth, (3) higher PRF, (4) lower data rate. The aggressive approach, instead, gives a lower data rate the highest priority [67].

Influence of environmental state on adaptation. The estimated received signal power is the first instance to trigger an adaptation of the PHY parameters, however, the decision on the parameter can be overruled by the environmental state. Indeed, the presence of NLOS conditions

or destructive interference might prioritize different settings. For instance, if the LOS component is destructively interfering with a multipath signal, changing the channel and center frequency, respectively, should always be privileged as shown in the evaluation in Sect. 3.4.

Implementation. The overhead of the adaptation is constituted by two bytes for each exchanged packet. The first byte encloses the PHY configuration to be used for future transmissions and the second byte contains the link state indicator consisting of 7 bits RSP and 1 bit indicating the presence of destructive interference. A NLOS identification technique was not included in the evaluation. If one of the nodes detects a degradation of the link it proposes a parameter change to its neighboring nodes, which have to acknowledge the parameter change to become effective. To avoid disagreements, the higher robustness of the SHR is exploited by accepting the successful reception of the preamble as a binary ACK, as discussed in Sect. 3.1.2. Note that the current implementation considers one global PHY configuration throughout a network. Defining clusters to determine local optimal settings is subject of future work (see Chapter 6).

3.4 Evaluation

This section evaluates the performance of the presented adaptation scheme. The first part compares the reliability and energy efficiency in dynamic environments when using static versus adaptive PHY settings (Sect. 3.4.1). The second evaluation shows the effectiveness of the proposed approach to identify and counteract destructive interference at runtime (Sect. 3.4.2).

3.4.1 Dynamic environments

To assess the performance of the adaptation scheme, we mount a transmitter and receiver in a hallway at a fixed distance of 5 m with the transmitter sending every 20 ms a packet to the receiver. The latter continuously estimates the link quality by employing the presented link state indicator (see Sect. 3.2). To mimic harsh and dynamic environments in a reproducible manner, a programmable attenuator is inserted before the transmitter’s antenna. This allows to run a defined sequence of signal strength variations such as the one shown in Fig. 3.7 (top). The first part is a triangle-shaped sequence (0 to 130 s) to push the transceiver to the edge of its sensitivity range and the second part is a random sequence emulating irregular variations (130 to 360 s).

When employing this attenuation sequence, initially, we use static settings (the default ones in Table 3.2) and observe the PRR, link state indicator as well as the energy consumption. Following, we repeat the experiment when using the presented adaptation scheme with an energy-conservative and an aggressive adaptation policy, respectively. Fig. 3.7 (second from top) shows that the PRR collapses when using the static settings (blue solid line), thus, a robust communication and in the case of a positioning system reliably sharing timestamps is prevented. Instead, when using the adaptation scheme, the PRR is always higher than 95 and 98 % when using the energy-conservative (orange dashed line) and the aggressive (purple dotted line) policy, respectively. Additionally, Fig. 3.7 (third from top) shows the estimated RSP over time (blue solid line). This information is used to trigger parameter changes. A more robust setting is used when the RSP is below the threshold $P_{th,r} = -93.3$ dBm (red dashed line) and a more efficient setting when the estimated RSP is above the threshold $P_{th,ee} = -88.8$ dBm (red dotted line).




Figure 3.7: Evaluation of the adaptation scheme. The reliability (PRR) and energy consumption is compared when using static and adaptive PHY settings, respectively. Adapted from [67].

Sustaining a high link reliability by applying the adaptation scheme comes at the cost of a slightly higher energy consumption (total energy for transmission and reception) as shown in Fig. 3.7 (bottom). For the given attenuation sequence, the energy costs are 5 and 29 % higher for the energy-conservative and the aggressive policy, respectively. However, the higher energy consumption is mainly required to overcome phases of high attenuation and challenging link conditions (marked with Ph. 2 and Ph. 4 in Fig. 3.7). In situations with higher link margin, the adaptation scheme consumes significantly less energy (Ph. 1 and Ph. 3). Please note in the fourth phase (Ph. 4) the difference in the energy consumption between the energy-conservative and the aggressive approach. Since the latter uses a lower data rate as the first instance to increase link reliability, which comes at an increased energy consumption.

3.4.2 Destructive interference

The measurements in the hallway highlighted areas where the LOS component was destructively interfering with a strong multipath component. The aim of this evaluation is to determine if the algorithm presented in Sect. 3.2.2 is capable of detecting these areas at runtime and if it is able to counteract these environmental conditions fast enough to avoid dropping a significant amount of packets. To this end, the transmitter is kept at a static position, whereas the receiver is moved around eventually entering the area of destructive interference. Fig. 3.8 shows the recording of the packet reception rate as well as the energy ratio ER defined in Eq. 3.1. After 30 s the




Figure 3.8: The proposed adaptation scheme detects at runtime situations of destructive interference (grey area) and counteracts these situations successfully. Adapted from [67].

receiver enters the area of destructive interference where it remains for another 30 s (marked with light grey area in Fig. 3.8). When using static PHY settings, the PRR drops to almost zero (blue solid line). Essentially the same happens when employing the adaptation scheme with an aggressive policy and without exploiting the environmental state information provided by the link state indicator (orange dashed line) as lowering the data rate does not help significantly in the presence of destructive interference. However, continuously tracking the energy ratio ER enables to identify destructive interference as shown in Fig. 3.8 (bottom). As soon as the power ratio falls below the experimentally determined threshold $ER_{th} = 1.5$ (see Sect. 3.2.2), the link state indicator pinpoints the presence of destructive interference. Correspondingly, the ranking of PHY settings is overruled and a more appropriate parameter change is triggered. Indeed, directly changing the carrier frequency enables to overcome the situation of destructive interference and to sustain a high PRR as demonstrated in Fig. 3.8 (purple dotted line).

Single Anchor Positioning using Multipath Assistance

The adaptation of physical layer parameters to reliably acquire and share timestamps in a network (see Chapter 3) lays the foundation for building robust and efficient UWB-based positioning systems. The latter typically derive the ToA and hence the required timestamps by detecting the first path component or leading edge of the estimated channel impulse response, as discussed in Sect. 2.4. However, in contrast to narrowband technologies, the high bandwidth of UWB transceivers further allows to resolve MPCs, i.e., specular reflections originating from walls and other objects. Still, typical UWB-based systems neglect this valuable multipath information. Thus, in this chapter we present a system called SALMA, which utilizes the first path component as well as the multipath information given by the estimated CIR. That additional information enables to cut down the required infrastructure to a single physical anchor by exploiting instead of mitigating multipath propagation. This significantly reduces the deployment effort, a stringent requirement to make indoor positioning systems *time-* and *cost-efficient*.

Related Work. Indoor positioning systems based on Wi-Fi [58], Bluetooth [11, 31], or Zig-Bee [116, 147] inherently suffer from multipath fading due to the low bandwidth and hence hardly achieve an accuracy below one meter. Instead, an increased bandwidth enhances the time resolution and hence the resilience to multipath propagation leading to decimeter-level accuracy of UWB-based positioning systems [129]. Thus, as discussed in Sect. 1.2, the large bandwidth resulting in a high resilience against multipath fading and a high time resolution combined with its ability to propagate through certain materials makes ultra-wideband increasingly attractive in indoor positioning systems using low-cost transceivers [106, 107, 117, 193]. However, these state of the art systems have in common that they require a high amount of anchors. For instance, the system presented in [193] employs eight and the system in [107] even fifteen anchors to cover a single room. SALMA, instead, requires just a single physical anchor by exploiting the reflections from static objects and walls, while still competing with multi-anchor systems in terms of accuracy and precision, as shown in the evaluation in Sect. 4.6.

Similarly, recent works have tried to reduce the employed infrastructure of positioning systems to a single anchor. Indeed, systems using quasi-static magnetic fields [6], dual-tone radio interferometric [123], or even single beacons for underwater navigation [26, 36] have been proposed to allow cost-efficient deployments. To this end, the most common technique to reduce the number of anchors are hybrid approaches combining ToA and AoA techniques using radio waves [182, 190, 203, 230]. In contrast to single anchor systems based on mechanical waves such as ultrasound [145], RF-based systems have the potential to penetrate obstacles and are less susceptible to multipath propagation [132]. Still, hybrid approaches using narrowband technologies suffer from the reduced signal bandwidth and high number of required antennas [228]. To increase the utilized bandwidth, an interesting concept was presented by Vashist et al. [201], there system named Chronos emulates a wideband radio on commodity Wi-Fi hardware, still, due to the lack of multipath assistance SALMA outperforms Chronos in terms of accuracy.

Correspondingly to the error bounds for time of arrival estimation discussed in Sect. 2.1.3, prior works have derived performance bounds for multipath-assisted indoor positioning [60, 88, 221] and even proved these observations using expensive, bulky, and wired-synchronized equipment [120, 136]. SALMA, instead, has built upon these algorithms and theoretical observations to provide the first low-cost multipath-assisted single anchor UWB-based indoor positioning system. Recently, the research on multipath-assisted positioning in the ultra-wideband spectrum is transferred to 5G and the mmWave spectrum to enable single anchor positioning at higher center frequencies and even higher bandwidth [2, 28, 102].

Even though UWB technology allows to resolve MPCs, still, it appears that multipath signals overlap and interfere with each other as discussed in Chapter 3. Hence, also the basic principle of SALMA - which is based on exploiting solely the delay information of the MPCs - suffers from a performance degradation in case of overlapping reflections. To tackle this issue, we extend the basic version with capabilities to exploit also the angular information of the MPCs. To this end, we utilize a self-made switchable directional antenna system. Similar approaches are presented using narrowband technologies and electronically steerable or switchable antenna systems to enable single anchor positioning [30, 63, 148, 163]. However, SALMA outperforms these systems due to the exploitation of the position-related information of the multipath components. Likewise, UWB-based systems exploiting directional antenna arrays exists, still, they do not exploit multipath information [188, 206] and/or are solely based on simulations [154, 236].

Structure. The remainder of this chapter is structured as follows. Sect. 4.1 gives a comprehensive overview of the SALMA and its main components followed by the design principles in Sect. 4.2. The latter includes the definition of the signal model, which extends the model discussed in Sect. 2.4 to support multiple directional antennas. Furthermore, Sect. 4.2 discusses the concept of virtual anchors and the positioning algorithm in detail. Sect. 4.3 describes the implementation of SALMA on low-cost UWB hardware based on the DW1000 and introduces the self-made directional antenna system. Sect. 4.4 discusses the basic version SALMA-light, while Sect. 4.5 describes the extended version SALMA-full and how it tackles multipath ambiguities. The chapter closes with an evaluation of SALMA in Sect. 4.6 demonstrating its resilience to NLOS conditions. Furthermore, a 24-hours measurement campaign demonstrates the robustness of SALMA in dynamic and crowded environments. As mentioned in Part II, note that two authors have contributed equally to most of the material discussed in this chapter.

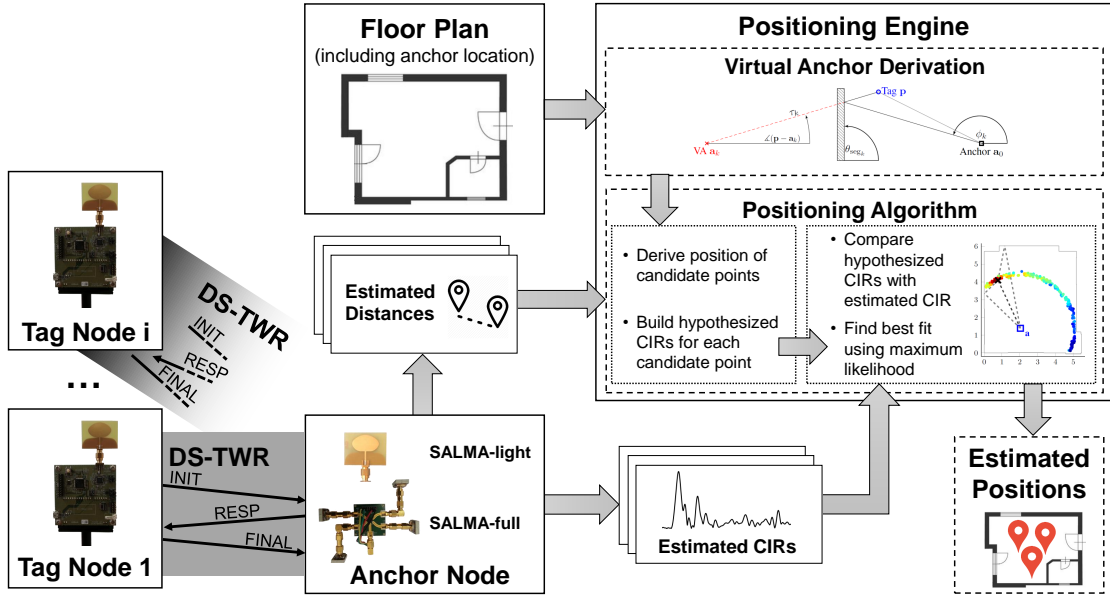


Figure 4.1: SALMA system overview. The infrastructure of the system consists of a single anchor node connected to a notebook running the positioning engine. Adapted from [70].

4.1 System Overview

Fig. 4.1 illustrates the overall structure of the SALMA positioning system. It consists of a single physical anchor node positioned at a known and static position \mathbf{a} as well as multiple mobile and battery-powered tags placed at unknown positions \mathbf{p}_i with $i = 1 \dots N_t$ and N_t being the number of tags. The anchor either employs a single omnidirectional antenna or an antenna system consisting of four switchable self-made directional UWB antennas. Depending on the applied antenna, we differ between two versions of SALMA: (i) *SALMA-light*, which utilizes an omnidirectional antenna and thus solely exploits the delay information of the MPCs and (ii) *SALMA-full*, which additionally exploits the angular domain and utilizes the directional antenna system. The static anchor is connected to and powered by a notebook running the *positioning engine* to compute the estimated positions of the tag. Since the notebook as a central entity holds the position information of all mobile devices, SALMA employs a server-based scheme as described in Sect. 1.1.1. However, the position estimation is initiated by the tags via a double-sided two-way ranging (DS-TWR) exchange with the anchor node to estimate the distance between anchor and tag $d_0 = \|\mathbf{p}_i - \mathbf{a}\|$. Besides the estimated distance \hat{d}_0 , the anchor further records the estimated channel impulse response provided by the UWB transceiver and forwards this information to the positioning engine. The latter also requires a floor plan showing the geometry of the building in which the system is installed as well as the anchor position and orientation. Please note that the map solely includes static objects such as walls and windows and there is no need to keep track of whether tables, shelves, or other furniture were moved. Starting from the floor plan, SALMA models the theoretical multipath propagation by employing the concept of virtual anchors, i.e., mirroring the anchor position on reflecting surfaces [136]. Based on the

position of the virtual anchors, one can build a hypothesized CIR for a given position. The positioning engine then compares the hypothesized channel impulse response with the estimated CIR derived by the anchor node in the DS-TWR exchange. To limit the computational complexity of SALMA and to enable its real-time applicability, the hypothesized CIR is solely built for a set of candidate points, which are selected on a circle of radius \hat{d}_0 centered in the anchor position \mathbf{a} . The positioning engine selects the candidate point attaining the best fit between the hypothesized CIR and the acquired one as the estimated tag position $\hat{\mathbf{p}}$.

4.2 Design Principles

This section discusses the mathematical foundations of the SALMA positioning system. In Sect. 4.2.1 we define a signal model aiming to represent the estimated CIR via the parameters of multipath components (MPCs). Next, it is shown in Sect. 4.2.2 how to use the known position of the anchor and the floor plan to determine virtual sources of the MPCs, so-called virtual anchors. The latter relate the tag position with MPC parameters embedded in the estimated CIR to introduce additional position-related information. Based on this information, Sect. 4.2.3 presents an algorithm to efficiently estimate the tag's position using a maximum likelihood estimator.

4.2.1 Signal model

Exploiting multipath information provided in the estimated channel impulse response requires its proper modeling¹. Thus, in this section we extend the model defined in Sect. 2.4 and Eq. 2.11 to introduce a signal model relating the estimated CIR $r(t)$ and the parameters of multipath components. To derive the CIR, the receiver excites the propagation channel using the transmitted UWB pulse $s_{tx}(t)$. For simplification, the latter also contains de-spreading and filtering at the receiver [139]. Thus, the estimated or observed CIR $r(t)$ can be modeled as the convolution of the definition in Eq. 2.11 with the transmitted pulse $s_{tx}(t)$ and follows as

$$r(t) = \sum_{k=0}^K \alpha_k(t) s_{tx}(t - \tau_k) + w(t). \quad (4.1)$$

The first term on the right-hand side describes K specular MPCs, i.e., reflections of the transmitted signal $s_{tx}(t)$ at flat surfaces such as walls, windows, and doors. These MPCs can be described using the geometric model introduced in Sect. 4.2.2. As discussed in Sect. 2.4, each MPC is characterized by its complex-valued amplitude α_k and delay τ_k . To simplify the positioning algorithm described in Sect. 4.2.3 and to allow for efficient computations, we assume the diffuse multipath to be zero-mean white Gaussian noise $w(t)$ with variance σ_w^2 .

The signal model introduced in Eq. 4.1 encloses the delay domain of the MPCs, however, SALMA is additionally capable of exploiting the angular domain by utilizing M antennas at the anchor (SALMA-full). Each of these antennas with index $m = 1, \dots, M$ is characterized by its

¹The estimated CIR is also referred to as observed CIR and corresponds to the effective system impulse response. It is derived by the anchor node from the last frame sent in the DS-TWR exchange from the tag to the anchor.

beampattern $b_m(\phi)$. Extending the signal model in Eq. 4.1 with antenna support allows to model the observed CIR $r_m(t)$ between a tag and the anchor's m^{th} antenna as [70]

$$r_m(t) = \sum_{k=0}^K \alpha_k b_m(\phi_k) s_{tx}(t - \tau_k) + w_m(t), \quad (4.2)$$

where, beside the amplitude α_k and delay τ_k , we introduce an additional parameter of the MPC, namely, the angle of departure ϕ_k of the k^{th} multipath component.

To compactly describe the signal model in Eq. 4.2, vector notation is used. With $f_s = 1/T_s$ being the sampling frequency of the observed signal $r_m(t)$ and N_s the number of acquired samples, the vector notation of the signal model in Eq. 4.2 follows as [70, 113]

$$\mathbf{r} = \mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha} + \mathbf{w} \quad (4.3)$$

with

$$\begin{bmatrix} \mathbf{r}_1 \\ \vdots \\ \mathbf{r}_M \end{bmatrix} = \begin{bmatrix} \mathbf{X}_1(\boldsymbol{\tau}, \boldsymbol{\phi}) \\ \vdots \\ \mathbf{X}_M(\boldsymbol{\tau}, \boldsymbol{\phi}) \end{bmatrix} \boldsymbol{\alpha} + \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_M \end{bmatrix} \quad (4.4)$$

and

$$\begin{aligned} \mathbf{r}_m &= [r_m(0 \cdot T_s), \dots, r_m([N_s - 1] \cdot T_s)]^T \\ \mathbf{X}_m(\boldsymbol{\tau}, \boldsymbol{\phi}) &= [b_m(\phi_0)\mathbf{s}(\tau_0) \dots b_m(\phi_K)\mathbf{s}(\tau_K)] \\ \mathbf{s}(\tau_k) &= [s_{tx}(0 \cdot T_s - \tau_k), \dots, s_{tx}([N_s - 1] \cdot T_s - \tau_k)]^T \\ \mathbf{w}_m &= [w_m(0 \cdot T_s), \dots, w_m([N_s - 1] \cdot T_s)]^T \\ \boldsymbol{\phi} &= [\phi_0, \dots, \phi_K]^T; \boldsymbol{\tau} = [\tau_0, \dots, \tau_K]^T; \boldsymbol{\alpha} = [\alpha_0, \dots, \alpha_K]^T. \end{aligned}$$

The proposed signal model represents the observed signal, i.e., the estimated CIR, via the MPC parameters, amplitude α_k , delay τ_k , and angle of departure ϕ_k .

4.2.2 From physical to virtual anchors

Since it is well-observed that reflections from flat surfaces, hence specular MPCs, contain position-related information regarding the tag position [114, 218], it is desired to derive relations between the latter and MPC parameters. To this end, MPCs are interpreted as signals originating from virtual sources, so called virtual anchors (VAs) [136]. The positions of these VAs \mathbf{a}_k ($k > 0$) are derived by mirroring the single physical anchor $\mathbf{a}_0 \triangleq \mathbf{a}$ at each reflective surface of the given floor plan. Fig. 4.2 illustrates the concept of virtual anchors using a single reflective wall resulting in a specular MPC (black solid line). The latter is assigned to a virtual anchor enabling efficient calculation of the MPC parameters delay τ_k and angle of departure ϕ_k . The delay of the k^{th} multipath component τ_k follows as the Euclidean distance between tag position \mathbf{p} and VA position \mathbf{a}_k , divided by the speed of light in air c , according to [70]

$$\tau_k = \frac{1}{c} \|\mathbf{p} - \mathbf{a}_k\| \quad (4.5)$$




Figure 4.2: Concept of virtual anchors (VAs) in order to relate the tag position \mathbf{p} to the MPC parameters angle of departure ϕ_k and τ_k by mirroring the physical anchor \mathbf{a}_0 . Taken from [70].

and as shown in Fig. 4.2 the angle of departure ϕ_k is derived as

$$\phi_k = 2\theta_{\text{seg}_k} - \angle(\mathbf{p} - \mathbf{a}_k), \quad (4.6)$$

where θ_{seg_k} denotes the angle of the reflective surface at which the physical anchor was mirrored to generate the VA and $\angle(\mathbf{p} - \mathbf{a}_k)$ indicates the azimuth angle between tag and virtual anchor. While the delay τ_k and angle of departure ϕ_k can be related directly to the tag position \mathbf{p} using the known virtual anchors, a proper model for the MPC amplitudes α_k is difficult to obtain due to inevitable path overlaps in indoor environments [112]. Hence, we propose to treat the amplitude α_k as nuisance parameter, estimated directly from the observation \mathbf{r} .

Please note that in the described concept the signal originates from the anchor and is transmitted to the tag. In SALMA instead the observed signal is derived from the *FINAL* message in the DS-TWR exchange, hence, from the tag to the anchor (see Fig. 4.1). Still, the applied concepts hold due to their independence of the direction of wave propagation. Furthermore, solely *single-order* reflections, i.e., signals that were reflected by a single reflective object in the signal path, are modeled and considered in SALMA. Thus, the number of reflecting surfaces in the floor plan determines the number of used MPCs K and hence the number of VAs. In the rectangular-shaped room shown in Fig. 2.5b $K = 4$ and hence also four VAs are obtained. The reasoning for focusing solely on single-order reflections is that higher-order reflections are strongly attenuated due to their increased path length and additional reflection losses and hence are hardly resolvable by low-cost transceivers. Additionally, it is required to take the visibility of the VAs at the tag position \mathbf{p} into account and to check the direct path from \mathbf{p} to the VA position \mathbf{a}_k for intersections with any obstacles or wall segments. Indeed, we consider only MPCs with a single involved intersection with the correct wall segment in the signal model [70].

4.2.3 Positioning algorithm

The signal model presented in Sect. 4.2.1 and the geometric model discussed in Sect. 4.2.2 allow to estimate the position of a tag \mathbf{p} . To this end, we employ a maximum likelihood (ML) estimator derived from the signal model in Eq. 4.3. As discussed in Sect. 4.2.1, zero-mean white Gaussian measurement noise \mathbf{w} is assumed to allow for efficient computations. Thus, the likelihood $p(\mathbf{r}|\mathbf{p})$ of observation \mathbf{r} conditioned on the tag position \mathbf{p} follows as [70, 105]:

$$p(\mathbf{r}|\mathbf{p}) = \left(\frac{1}{2\pi\sigma_w^2}\right)^{MN_s} \exp\left\{-\frac{1}{2\sigma_w^2}\|\mathbf{r} - \mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha}\|^2\right\}. \quad (4.7)$$

Taking the logarithm of Eq. 4.7 results in the log-likelihood function as

$$\log p(\mathbf{r}|\mathbf{p}) = -MN_s \log(2\pi\sigma_w^2) - \frac{1}{2\sigma_w^2} \|\mathbf{r} - \mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha}\|^2. \quad (4.8)$$

The geometric model presented in Sect. 4.2.2 allows to relate the parameters $\boldsymbol{\tau}$ and $\boldsymbol{\phi}$ to the tag position via Eq. 4.5 and Eq. 4.6, respectively. However, as discussed in Sect. 4.2.2, the amplitudes $\boldsymbol{\alpha}$ have to be derived directly from the observation \mathbf{r} . To this end, $\boldsymbol{\alpha}$ is estimated using the least squares solution according to [112]

$$\hat{\boldsymbol{\alpha}} = (\mathbf{X}^H(\boldsymbol{\tau}, \boldsymbol{\phi})\mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi}))^{-1} \mathbf{X}^H(\boldsymbol{\tau}, \boldsymbol{\phi})\mathbf{r} \quad (4.9)$$

with the superscript $(\cdot)^H$ denoting the Hermitian transpose. Thus, the amplitudes $\boldsymbol{\alpha}$ are substituted in Eq. 4.8 with the estimated amplitudes $\hat{\boldsymbol{\alpha}}$. Consequently, the estimated position of the tag $\hat{\mathbf{p}}$ maximizing the log-likelihood in Eq. 4.8 can be formulated as

$$\hat{\mathbf{p}} = \arg \max_{\mathbf{p} \in \mathcal{P}} \log p(\mathbf{r}|\mathbf{p}) = \arg \min_{\mathbf{p} \in \mathcal{P}} \|\mathbf{r} - \mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi})\hat{\boldsymbol{\alpha}}\|^2. \quad (4.10)$$

Employing $\boldsymbol{\tau}$ and $\boldsymbol{\phi}$, derived from the geometric model in Sect. 4.2.2, as well as the estimated amplitudes $\hat{\boldsymbol{\alpha}}$ allows to build the hypothesized CIR $\mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi})\hat{\boldsymbol{\alpha}}$. The latter is compared with the observed CIR for all possible solutions \mathcal{P} . The position for which the hypothesized CIR comes closest to the observed CIR (and thus maximizes the likelihood in Eq. 4.10) defines the position estimate $\hat{\mathbf{p}}$. In fact, the set of possible positions \mathcal{P} contains all positions within the communication range of the single physical anchor. Thus, finding a global maximum for the maximum likelihood estimator would result in an enormous computational complexity as it requires to evaluate Eq. 4.10 at each feasible tag position within communication range. This would hinder an efficient implementation on low-cost low-power hardware. Thus, in SALMA the set of possible tag positions is limited to a subset of N_C candidate points $\mathcal{P} = \{\mathbf{p}^{(j)}\}_{j=1}^{N_C}$. The latter are arranged on a circle around the anchor \mathbf{a} with Gaussian distributed radius $d^{(j)} \sim \mathcal{N}(\hat{d}_0, \sigma^2)$ and uniformly distributed angle $\phi^{(j)} \sim \mathcal{U}(0, 2\pi)$ [113], where \hat{d}_0 defines the estimated distance between tag and anchor derived from the DS-TWR exchange and σ^2 the range variance of the employed hardware (see Sect. 4.3.1). Hence, the position of the j^{th} candidate point follows as

$$\mathbf{p}^{(j)} = [d^{(j)} \cos(\phi^{(j)}), d^{(j)} \sin(\phi^{(j)})]^T + \mathbf{a}. \quad (4.11)$$

We discard candidate points lying outside of the area of interest determined by the floor plan.

4.3 Implementation on Low-cost Devices

The design principles discussed in Sect. 4.2 depict the theoretical foundation of SALMA, which are inherently formulated to allow its implementation on low-cost devices. Additionally, Sect. 4.3.1 describes hardware-specific impacts on and properties of the signal model and positioning algorithm, namely, the transmitted pulse shape as well as the selection of candidate points. Sect. 4.3.2 discusses in detail the employed self-made directional UWB antenna system.

4.3.1 Hardware-specific impacts on signal model and positioning algorithm


Identically to the characterization of the PHY parameters described in Chapter 3, we utilize the EVB1000 boards employing the Decawave DW1000 UWB transceiver in SALMA for the tags and anchor. Whereas the tags can be moved around freely, the single physical anchor is mounted at a fixed position \mathbf{a} and is connected to a notebook running the positioning engine in MATLAB.

Transmitted pulse. The signal model in Eq. 4.2 requires knowledge of the pulse shape $s_{tx}(t)$ transmitted by the EVB1000 UWB modules. Thereto, the IEEE 802.15.4 standard gives freedom to the developers by allowing an arbitrary pulse shape, as long as it fulfills certain requirements on its cross-correlation with a standard reference signal, a root-raised-cosine pulse with a roll-off factor of $\beta = 0.5$ [95]. However, Decawave does not provide information regarding the transmitted signal of the DW1000. Hence, we identified the transmitted signal $s_{tx}(t)$ in a measurement campaign. To this end, we acquired 1000 CIRs in clear LOS and considered the average LOS component over these signals as the transmitted pulse shape $s_{tx}(t)$ [70].


Selection of candidate points. To efficiently implement the maximum likelihood estimator described in Sect. 4.2.3, Eq. 4.10 is solely evaluated on a circle around the anchor position \mathbf{a} , where the Gaussian distributed radius $d^{(j)} \sim \mathcal{N}(\hat{d}_0, \sigma^2)$ depends on the estimated distance between tag and anchor \hat{d}_0 as well as the range variance σ^2 . To account for hardware variability, it is recommended to calibrate the distance estimation when employing new hardware. To this end, we performed 5000 DS-TWR exchanges between anchor and tag in clear LOS and at a distance of $d_0 = 2m$. The derived mean and variance of the difference between the reported distance \hat{d}_{DW} and the true distance $d_0 = 2m$ is $\mu = 0.26 m$ and $\sigma^2 = (0.054 m)^2$, respectively. The latter defines the variance of the Gaussian distribution of the candidate points and the mean of the Gaussian distribution is defined by the calibrated distance $\hat{d}_0 = \hat{d}_{DW} - \mu$.

4.3.2 Directional UWB antenna system

As indicated in Sect. 4.1, we differ between the basic version SALMA-light exploiting the delay information of MPCs and SALMA-full utilizing the delay as well as angular domain. To this end, at the anchor either a single linearly polarized omnidirectional dipole antenna or multiple directional antennas are employed. Due to the lack of directional antennas for UWB-based systems, a compact microstrip patch UWB antenna is designed in-house by Bakr et al. [9]. To fulfill the requirements of an efficient positioning system, the design process focused on low cost as well as small form factor while still maintaining high directionality and constant electrical performance over the large frequency range of UWB systems. Due to the planar structure of the antenna and low-cost FR4 substrate, the price of one antenna is potentially less than one dollar in mass production at a size of just 25 x 23 x 4 mm [65]. We have combined four of these antennas to a switchable directional antenna system using a high-speed RF switching network (see Fig. 4.3a). Antenna pattern measurements in an anechoic chamber recorded in the azimuth plane (elevation = 90°) with 2° resolution show that the four antennas properly cover the 360° horizon (see Fig. 4.3b). The average gain over all four antennas is 3.028 dB and the half-power beam width (HPBW) is 98.5° [65]. The results are exemplarily shown at 5 GHz, however, the system supports center frequencies up to 8 GHz and therefore covers all radio channels supported by the DW1000 transceiver (see Table 3.1). Please note that the HPBW of 98.5° of the self-



(a)



(b)

Figure 4.3: Self-made switchable directional UWB antenna system (a) and the measured antenna patterns of the system at 5 GHz in polar coordinates and azimuth plane (b). Taken from [65].

designed directional UWB antenna was measured in an anechoic chamber. However, these results were not experienced when combining the directional antenna system with the Decawave EVB1000 UWB modules. Indeed, the effective HPBW is increased to about 150° , which the author assumes is due to internal saturation effects when estimating the CIR in the transceiver. Despite this reduced directionality when using low-cost transceivers, we show in Sect. 4.6.2.1 the effectiveness of the directional antenna system to mitigate multipath ambiguities.

4.4 SALMA-light: Exploiting Delay Information

The signal model in Eq. 4.1 defines the estimated channel impulse response when employing an ideal omnidirectional antenna at the anchor node (see Fig. 4.1), i.e., when solely exploiting the delay information of the multipath components τ while ignoring the angle of departure ϕ . This implementation is referred to as *SALMA-light* and defines the basic version of SALMA.

As discussed in Sec. 4.2.2, the MPC amplitudes α required to estimate the position of a tag are challenging to model, hence, we directly derive them from the observation \mathbf{r} according to Eq. 4.9. However, the latter requires a computationally expensive matrix inversion, and, in the case of overlapping MPCs, the matrix might not even be invertible at all. Hence, to ensure an efficient implementation of SALMA-light on low-cost hardware it is proposed to evaluate the log-likelihood function in Eq. 4.10 iteratively [70, 113]:

$$\text{init} : \mathbf{r}^{(0)} = \mathbf{r}$$

$$\text{for } k = 1 \dots K : \hat{\alpha}_k = \mathbf{s}^H(\tau_k) \mathbf{r}^{(k-1)} \quad (4.12)$$

$$\mathbf{r}^{(k)} = \mathbf{r}^{(k-1)} - \hat{\alpha}_k \mathbf{s}(\tau_k) \quad (4.13)$$

Essentially, this iterative approach takes the observed CIR (*init*) and sequentially subtracts sub-hypotheses ($\hat{\alpha}_k \mathbf{s}(\tau_k)$). The latter are composed of the transmitted pulse shape $s_{tx}(t)$ shifted

to the respective delay τ_k and weighted by the amplitude estimate $\hat{\alpha}_k$. After K iterations the residual signal $\mathbf{r}^{(K)}$ corresponds to the similarity of the hypothesized and observed CIR. Thus, it appears as an approximation of the log-likelihood function in Eq. 4.10. This employed iterative algorithm is repeated for each candidate point and the one with the highest value and thus the highest similarity to the observed CIR is chosen to be the tag position estimate $\hat{\mathbf{p}}$.

4.5 SALMA-full: Exploiting Delay and Angular Information

While the method described in Sect. 4.4 allows an efficient implementation of the log-likelihood function, the single omnidirectional antenna and ignoring its beampattern restricts the algorithm to solely exploit the delay information of the multipath components. Hence, this restriction makes the algorithm susceptible to overlapping MPCs as it is not capable of distinguishing between MPCs arriving with the same delay but from different directions as shown in Sect. 4.6.2. In *SALMA-full* in addition to the delay domain, we exploit the angular information of multipath components to reduce the sensitivity to overlapping MPCs and ambiguities. To this end, the anchor is equipped with the self-made switchable directional UWB antenna system (see Sect. 4.3.2) and the positioning algorithm utilizes the full signal model from Eq. 4.2 and Eq. 4.3, respectively. Within the signal model each antenna m is characterized by its beampattern $b_m(\phi)$ covering one sector of the azimuth plane as shown in Fig. 4.3b.

Non-phase-coherent amplitude estimates. The directional UWB antenna system described in Sect. 4.3.2 employs four self-made directional antennas, which are sequentially selected to observe the estimated CIR between the tag and each antenna. To combine these individual antenna measurements one would require phase-coherency, i.e., a stable phase-lock between consecutive measurements, which is in contrast to dedicated measurement equipment such as vector network analyzers not given by low-cost transceivers [112]. Instead, the non-coherent clocks between consecutive measurements are perceived as an unknown phase offset φ affecting the complex-valued MPC amplitude $\alpha_{k,m} \approx e^{j\varphi} \alpha_{k,m'}$. Hence, the unknown phase offset φ would be required in the amplitude estimation in Eq. 4.9 as well as in the log-likelihood function in Eq. 4.10. To avoid the necessity of phase-coherent measurements while preserving the performance gain provided by the angular domain, we follow an approach presented by Kulmer et al. [112]. To this end, non-overlapping multipath components, i.e., $(\mathbf{s}(\tau_1)^H \mathbf{s}(\tau_2) \approx 0)$ for any $\tau_1 \neq \tau_2$, are assumed. This allows to estimate an MPC amplitude α_k independently as projection of the normalized signal $\frac{\mathbf{s}^H(\tau_k)}{\mathbf{s}^H(\tau_k)\mathbf{s}(\tau_k)}$ onto the m -th measurement \mathbf{r}_m according to [70]

$$\hat{\alpha}_{k,m} = \frac{\mathbf{s}^H(\tau_k) \mathbf{r}_m}{\mathbf{s}^H(\tau_k) \mathbf{s}(\tau_k)}. \quad (4.14)$$

Furthermore, to take the beampattern into account one can define the absolute-valued average [141] and hence an estimate of the k^{th} MPC amplitude α_k^{avg} according to

$$\hat{\alpha}_k^{\text{avg}} = \frac{\sum_{m=1}^M |\alpha_{k,m}| \cdot |b_m(\phi_k)|^2}{\sum_{m'=1}^M |b_{m'}(\phi_k)|^2}. \quad (4.15)$$

Thus, the amplitude estimate $\hat{\alpha}_{k,m}^{\text{avg}}$ of the m -th antenna and k -th MPC follows as

$$\hat{\alpha}_{k,m}^{\text{avg}} = \hat{\alpha}_k^{\text{avg}} \exp(j\angle\alpha_{k,m}), \quad (4.16)$$

where the remaining phase $\angle\alpha_{k,m}$ is extracted from the individual antenna measurement. This approximates the amplitude estimation in Eq. 4.9 by combining MPC amplitudes from non-phase-coherent measurements and by taking the directivity of M antennas into account.

Positioning algorithm using directional antennas. To consider the physical separation of the antennas employed in SALMA-full (see Fig. 4.3a), the mean value of all ranges between the tag and each antenna is utilized to create the set of candidate points around the anchor position \mathbf{a} . Equivalent to SALMA-light, for each of these candidate points the MPC parameters $\boldsymbol{\tau}$ and ϕ are derived using the geometric relations presented in Eq. 4.5 and Eq. 4.6, respectively. However since SALMA-full collects observations from M directional antennas, the iterative approach of SALMA-light to estimate the tag's position \mathbf{p} (see Sect. 4.4) has to be adapted to utilize multiple observed CIRs $\mathbf{r} = [\mathbf{r}_1^T, \dots, \mathbf{r}_M^T]^T$. Furthermore, we substitute the amplitude estimate $\hat{\alpha}_k$ in the iteration step in Eq. 4.12 by the non-phase-coherent amplitude estimates from Eq. 4.16 and replace the subtraction step in Eq. 4.13 with [70]

$$\mathbf{r}_m^{(k)} = \mathbf{r}_m^{(k-1)} - b_m(\phi_k) \hat{\alpha}_{k,m}^{\text{avg}} \mathbf{s}(\tau_k).$$

This results in a stacked residual $\mathbf{r}^{(K)} = [(\mathbf{r}_1^{(K)})^T, \dots, (\mathbf{r}_M^{(K)})^T]^T$, which we use as an approximation of the log-likelihood function also taking the beampattern into account. This iterative approach is performed in SALMA-full for each candidate point to obtain the estimate of the tag's position $\hat{\mathbf{p}}$ by picking the candidate point with the maximum value.

4.6 Evaluation

This section presents an evaluation of SALMA illustrating its capabilities in real-world deployments. In Sect. 4.6.1 the experimental setup is described, thereby we challenged SALMA in two indoor environments: an office and a stockroom. Sect. 4.6.2 discusses the positioning performance with a focus on the comparison between SALMA-light and SALMA-full (Sect. 4.6.2.1) followed by a detailed analysis of SALMA-full (Sect. 4.6.2.2). Sect. 4.6.3 highlights that SALMA excels even in non-line-of-sight conditions (Sect. 4.6.3.1), changing as well as crowded environments (Sect. 4.6.3.2 and Sect. 4.6.3.3), respectively. Note that this section presents the evaluation in a condensed fashion, for additional measurements please refer to [70].

4.6.1 Experimental setup

We perform the initial experiments discussed in Sect. 4.6.2 in the office environment shown in Fig. 4.4a and Fig. 4.5a (Room A), which contains several obstacles and scattering objects. The tag and the anchor are mounted on a tripod at a height of 1.50 m, i.e., above various obstacles to ensure LOS conditions. Instead, we evaluate the performance of SALMA in NLOS situations separately in Sect. 4.6.3 in Room A and in a stockroom (Room B, see Fig. 4.4b and Fig. 4.5b). The latter is significantly larger than Room A (46.7 m² vs. 31.6 m²) and contains storage racks

and several other metal obstacles. Initially, we derived for each room the dimensions of the four walls used as virtual signal sources. The anchor is placed at a static position (blue square in Fig. 4.4), whereas the tag is placed in each room at $N_{EP} = 35$ evenly distributed evaluation points. Both the tag and anchor use the same PHY settings, namely: data rate of 6.8 Mbps, PRF of 64 MHz, and a PSR of 1024. The latter is chosen to ensure a high SNR of the employed CIR estimate as discussed in Sect. 3.1.3. Furthermore, we utilize Channel 7 of the DW1000 due to its high bandwidth (900 MHz, see Table 3.1). At each evaluation point 100 position estimates are observed and the absolute error of each trial is calculated as the Euclidean distance between the position of the evaluation point \mathbf{p}_{EP} and the i^{th} position estimate $\hat{\mathbf{p}}_i$: $\text{Err}_i = \|\hat{\mathbf{p}}_i - \mathbf{p}_{EP}\|$. To derive statistically meaningful observations, we investigate the cumulative distribution function (CDF) over the errors Err_i of all $N_{EP} \cdot 100 = 3500$ position estimates.

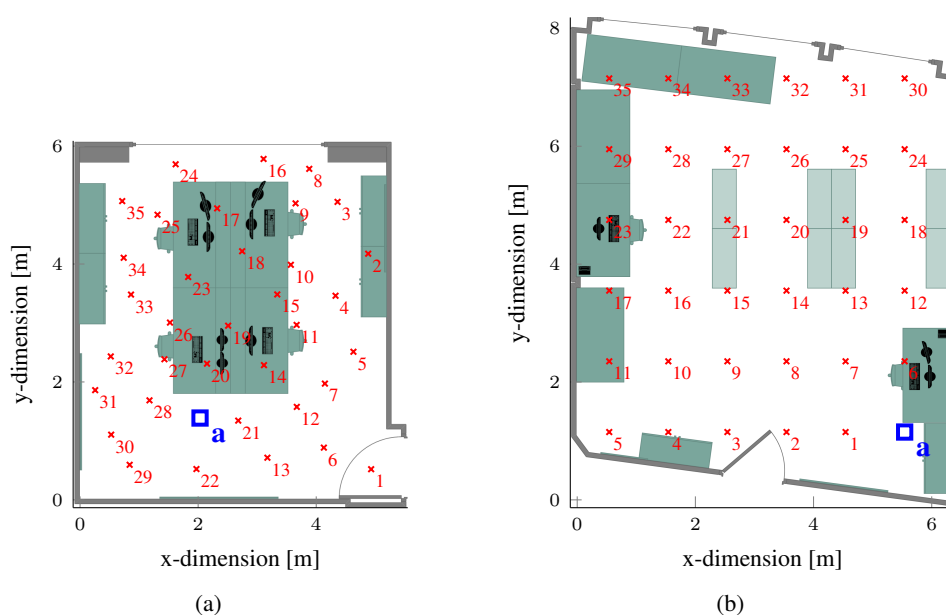


Figure 4.4: Evaluation setup (2D-plan): (a) An office (Room A) and (b) a stockroom (Room B) are evaluated at 35 evenly distributed evaluation points (red crosses). Adapted from [70].



Figure 4.5: Evaluation setup (Image): Both rooms are filled with obstacles and scattering objects. The white dashed line marks the height under obstructed LOS. Adapted from [70].

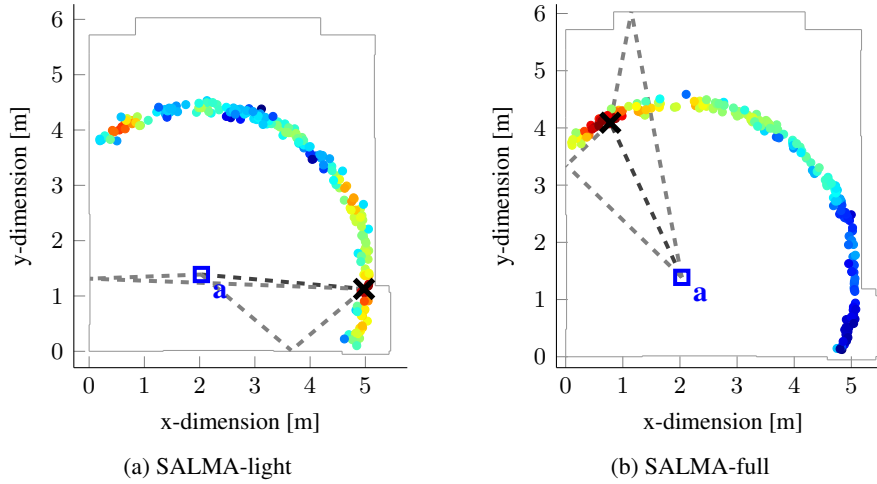


Figure 4.6: Solely exploiting the delay domain of MPCs may result in multipath ambiguities (a), which are resolvable by additionally utilizing the angular domain (b). Adapted from [70].

4.6.2 Positioning performance

This section discusses the positioning performance of SALMA in an office environment and in LOS conditions. In Sect. 4.6.2.1 we compare SALMA-light and SALMA-full qualitatively and quantitatively showing that SALMA-full is capable of mitigating outliers caused by multipath ambiguities. Sect. 4.6.2.2 studies the performance of SALMA-full individually for each evaluation point and discusses the role of candidate points and the antenna.

4.6.2.1 SALMA-light vs. SALMA-full

Two different implementations of SALMA are presented, namely, SALMA-light (see Sect. 4.4) and SALMA-full (see Sect. 4.5). This section compares them qualitatively by investigating a single position estimate and its likelihood values and quantitatively by inspecting the cumulative distribution function of the position error, respectively.

Qualitative comparison. SALMA-light solely exploits the position-related information provided by the path delay of MPCs, however, this may result in a highly multimodal likelihood for the candidate positions, i.e., multiple regions seem to fit the observed signal. This is illustrated in Fig. 4.6, which shows the position estimate (black cross) using SALMA-light (Fig. 4.6a) and SALMA-full (Fig. 4.6b), respectively, for a tag positioned at evaluation point #34. The colored dots indicate the likelihood value of each candidate point (red=high, blue=low). Despite the fact that the evaluation point #34 is in the top left corner of the room (cf. Fig. 4.4a), SALMA-light displays three regions with similarly high likelihood values (red and orange dots). This is due to ambiguous MPC paths: Although the MPCs indicated with grey dashed lines in Fig. 4.6a and Fig. 4.6b are arriving from different directions, SALMA-light cannot distinguish between them due to the equal path lengths and the employed omnidirectional antenna. SALMA-full instead




Figure 4.7: Due to multipath ambiguities SALMA-light sustains an position error below 30 cm only in 70% of the estimates. By exploiting the angular domain, SALMA-full instead exhibits an error below 30 cm in 99% of the cases (Room A). Adapted from [70].

takes the angle of departure of the MPCs into account and hence explores the angle domain on top of the delay information. This allows to resolve the multipath ambiguities and to find the true position of the tag due to the combined directional antenna observations.

Quantitative comparison. The previous discussion shows that using directional antennas mitigates multipath ambiguities and hence avoids a multimodal likelihood. However, to derive statistically meaningful observations and to compare SALMA-light and SALMA-full quantitatively, we present a CDF of the absolute position error taking all 3500 estimates into account in Fig. 4.7. It depicts the CDF for SALMA-light (dashed orange line) and SALMA-full (solid blue line). Using SALMA-light, 67.3% of all evaluations exhibit a position error below 25 cm. However, the tail of the CDF shows that 21.7% of the estimates have an error above one meter. These outliers are due to ambiguities in the delay domain of the MPCs. By additionally exploiting the angular information of the MPCs and using the directional antenna system, SALMA-full mitigates these outliers and eventually ensures that 90% of all evaluations exhibit an error below 20.17 cm and the error of 99% of the evaluations is still below 29.72 cm. These results show that SALMA-full clearly outperforms SALMA-light and fulfills the requirements for a robust positioning system. Hence, the following evaluations will focus on SALMA-full only.

4.6.2.2 Detailed analysis of SALMA-full

The results in Sect. 4.6.2.1 show that towards a robust single anchor positioning system it is required to exploit the delay as well as angular information of multipath components. Hence, in this section SALMA-full is evaluated in more detail and individually for each evaluation point. To this end, Fig. 4.8a shows the mean (blue cross) and the 3-fold standard deviation (black error ellipse) for the 100 estimates at each evaluation position of Room A. The mean indicates the estimation bias (distance to true position marked with red crosses), which is modest for all evaluation points. The error ellipses instead give insights in the radial or ranging (facing the LOS) as well as the tangential or angle (perpendicular to the LOS) error. The latter is determined by the information gained from the MPCs and is higher than the radial error due to the good ranging precision of UWB. An asset of SALMA is that its performance does not degrade at higher ranges. Indeed, for instance evaluation points #8, #16, and #24 are placed more than four

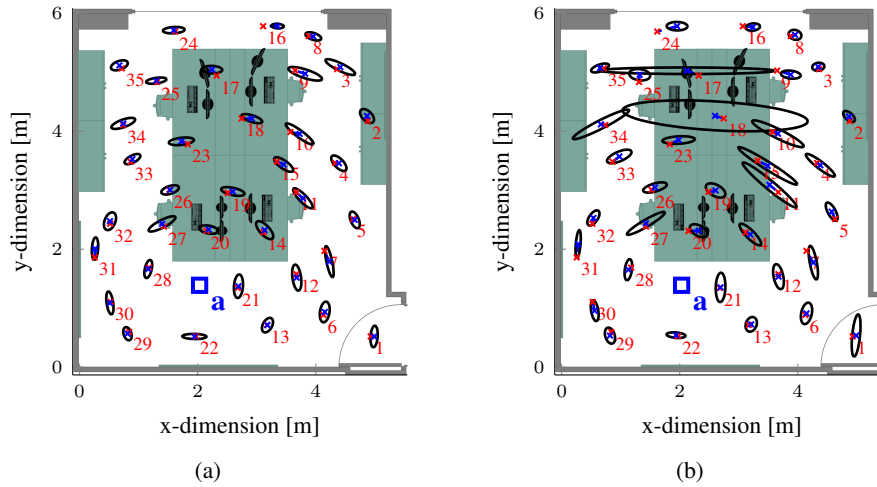


Figure 4.8: Mean position (blue cross) and 3-fold standard deviation (black ellipses) for each evaluation point in Room A and under clear LOS (a) and obstructed LOS (b). Adapted from [70].

meters apart from the anchor, remarkably, their estimates are as accurate as the ones obtained at much closer distances to the single physical anchor (blue square) [70]. This is in contrast to indoor positioning systems making use of angle of arrival techniques such as visual or camera based systems, where the precision decreases at larger distances [128, 132].

Role of candidate points. The simplification to solely evaluate the log-likelihood function at a set of candidate points arranged on a circle around the anchor position a influences the positioning performance of SALMA. The more fine-grained the sampling and hence the higher the number of candidate points N_C is, the better is the performance. However, higher number of candidate points results in an increased computational complexity. Indeed, the computation time increases linearly with the number of candidate points. We have shown in simulations of the positioning performance and the computational complexity that $N_C = 200$ is a reasonable trade-off between positioning performance and computational expenses as further increasing N_C does not result in significant improvements [70]. Hence, $N_C = 200$ is the preferred setup and used for all evaluations discussed within this chapter. Note that the current implementation evaluates the log-likelihood function for each candidate point, future works might employ algorithms to optimize the search for the candidate point showing the maximum likelihood (see Sect. 6.1).

Role of directional antennas. Besides the number of candidate points N_C , the number of employed directional antennas as well as their beam pattern influences SALMA-full and its performance. We have shown in simulations that already increasing the number of antennas from four to six and improving the half-power beam width to 90° decreases the 90% error by about 10 cm and 99% of the position estimates would achieve an error below 20 cm [70]. This highly motivates future work on directional UWB antennas and to extend the presented directional antenna system to an antenna array with a significantly higher amount of antennas.




Figure 4.9: Positioning performance of SALMA-full in clear LOS, obstructed MPC, and obstructed LOS conditions evaluated in Room A (office environment). Adapted from [70].

4.6.3 Robustness to non-line-of-sight and dynamic environments

The evaluations in Sect. 4.6.2 highlight the potential of exploiting multipath signals to provide robust and efficient position information. Although we performed the measurements in a challenging indoor environments, still, there is always clear line-of-sight between the single anchor and the tag. However, the success and value of a positioning system strongly depend on (i) its performance in NLOS conditions as well as (ii) its robustness in dynamic environments and crowded settings. Hence, this section investigates the positioning performance of SALMA under NLOS conditions (Sect. 4.6.3.1), in different environments and when obstacles are moved (Sect. 4.6.3.2), as well as in highly-dynamic and crowded surroundings (Sect. 4.6.3.3).

4.6.3.1 Performance in NLOS conditions

Positioning systems based on optical technologies inherently fail in case the direct view between devices is restricted. RF-based systems instead possibly penetrate obstacles and remain functional in NLOS conditions, however, blocking the LOS in time-based systems leads to a positive range bias [131]. This is either caused by the lower propagation speed in case the signal propagates through the obstacle, or, in case of a fully blocked LOS, due to the misinterpretation of a reflection as the direct path. Indeed, we have shown in experiments with low-cost UWB transceivers that already two humans blocking the LOS cause a range bias of more than 40 cm [70]. Since SALMA requires the distance estimate between the anchor and tag \hat{d}_0 to distribute the candidate points on a circle around the physical anchor (see Sect. 4.2.3) it is affected by a blocked LOS as the positive range bias increases the radius of the circle.

Quantitative analysis. However, SALMA is not solely utilizing a single path, namely the LOS component, but rather more paths due to the exploitation of multipath propagation. Thus, we evaluate if this enables a higher robustness to obstructed LOS situations and, additionally, we analyze the impact of blocked MPCs on the performance of SALMA. To this end, we repeat the evaluation in Room A described in Sect. 4.6.2 with the tag and the anchor at a height of 1.20 m, corresponding to the height of several obstacles such as monitors and shelves (see Fig. 4.5a). This results in obstructed LOS for twelve evaluation points, thus in total for 1200 positioning trials. Note that for all these evaluation points there were also specular MPCs blocked by ob-

stacles. For twenty evaluation points, instead, the LOS was still clear but specular MPCs were obstructed. The three remaining evaluation points are in clear LOS with no blocked MPCs, thus, they are ignored for this evaluation. Fig. 4.9 shows the CDF of the position error using SALMA-full in clear LOS, obstructed MPC, and obstructed LOS conditions. The observations in clear LOS (blue solid line) are taken from Sect. 4.6.2.1 with a 90% error of 20.17 cm. The orange dashed line indicates the measurements where significant MPCs are blocked by obstacles. The median error is at 10.25 cm and the error for 90% of the estimates is below 30.52 cm. Even under blocked LOS (purple dotted line) the 90% error is still below 30.7 cm. These results emphasize that SALMA remains robust even in the case of obstructed MPCs and LOS, respectively. This is accomplished due to the exploitation of multiple signal paths between the anchor and tag as well as by employing directional antennas to utilize angular information in SALMA-full.

Qualitative analysis. Fig. 4.8b shows the positioning results individually for each evaluation point in Room A and for obstructed MPC and LOS conditions, respectively. While the position bias (distance between red and blue crosses) is still low and did not increase significantly compared to the clear LOS case (see Fig. 4.8a), the larger error ellipses indicate that the tangential or angle deviation (perpendicular to the LOS) increases for most of the evaluation points. This is caused by the increased range bias as well as modeled MPCs that are suppressed by obstacles. Particularly interesting are the evaluation points #17 and #18 due to their significantly higher standard deviation. The reason is the position of the physical anchor a for these two evaluation points. Indeed, the monitors and obstacles at the left and right wall block the respective MPCs. Thus, the only two reflections contributing position-related information, the LOS and the reflection from the window, are arriving at the same angle at the anchor and are hence difficult to resolve in the angular domain [70]. This results in multipath ambiguities similar to the ones obtained when using SALMA-light (see Sect. 4.6.2.1). Selecting a different position for the anchor would avoid this unfortunate geometric configuration. In general, it is preferred to put the anchor closer to the center of a room to benefit from the beampatterns in all directions. In SALMA-light, instead, it might be beneficial to mount the anchor in the corner of a room as this results in a smaller circle of candidate points and hence a reduced risk of ambiguities.

4.6.3.2 Performance in changing environments

Until now all the evaluations were performed in an office environment (Room A). However, to challenge SALMA in a different environment and in the presence of moving obstacles, we evaluate the system additionally in a stockroom (Room B). The latter is significantly larger than Room A (46.7 m^2 vs. 31.6 m^2) and it depicts an even more challenging environment cluttered with desks, storage racks (bright rectangles in Fig. 4.4b) and several other metal obstacles (see Fig. 4.5b). Mounting the anchor and tag at a height of 1.20 m, Fig. 4.10 (solid blue line) shows the CDF of all evaluation trials in Room B. The median error is 18.6 cm and 90% of all position estimates obtain an error below 44.5 cm. Thus, the larger size of the room with more metal obstacles results in a slightly worse performance compared to Room A. Furthermore, the wall materials and its reflective properties influence the performance of SALMA due to the exploitation of multipath propagation. Indeed, materials such as glass and metal are reflecting electromagnetic waves well and including them in the geometric model enhances the position




Figure 4.10: Positioning performance of SALMA-full in the stockroom (Room B) and in the case of moving obstacles (storage racks are empty or filled). Adapted from [70].

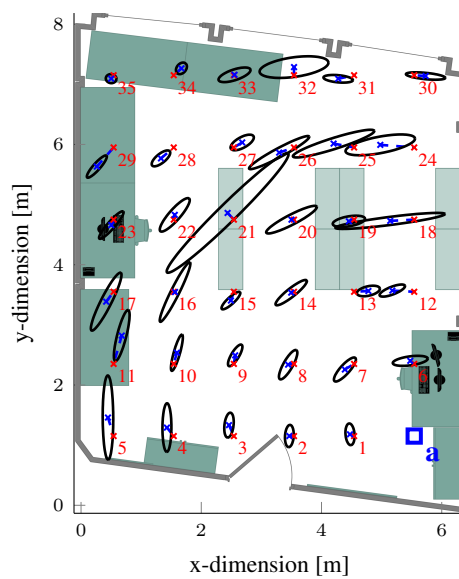


Figure 4.11: Mean position (blue crosses) and 3-fold standard deviation (black error ellipses) for each evaluation point in Room B with empty racks. Adapted from [70].

estimation. Instead, plaster boards (such as the eastern wall in Room B) or wooden surfaces give little to no contribution in terms of specular reflections [70]. Thus, such materials should not be included in the model. Fig. 4.11 shows the mean (blue cross) and three-fold standard deviation (black error ellipses) for the individual evaluation points in Room B. Similar to the observations in Room A (see Sect. 4.6.3.1), some evaluation points (e.g., #4, #5, #12, #18, #24) suffer from an unfavorable anchor placement as the LOS is arriving from a similar angle as strong MPCs. Others, such as #21, #25, #26, #32 are affected by obstructed LOS and MPCs, respectively.

Please note that the current evaluations were performed and discussed with empty racks. However, we evaluate SALMA also when the storage racks in Room B are stocked up with full beer crates and other objects (see Fig. 4.5b) to simulate a changing environment. Due to the well-reflecting nature of the added objects, it might be beneficial to add them to the geometric model




Figure 4.12: Snippet of a 24-hours experiment in a dynamic office environment. The dashed orange line depicts the mean error of 50 position estimates over time, whilst the solid blue line shows the 90% error. The green line indicates the number of people present. Adapted from [70].

to derive additional virtual sources. However, this is not practical in real-world environments, which are continuously changing due to moving obstacles or furniture. Fig. 4.10 (dashed orange line) shows that the performance of SALMA is affected by the added obstacles due to a higher range bias, still, the position error stays in reasonable bounds. Positioning systems based on RSS profiling and fingerprinting instead are highly affected by such adaptations of the environment and often require to update or repeat measurements.

4.6.3.3 Performance in crowded environments

Several indoor positioning applications discussed in Sect. 1.1.2 involve dynamic environments due to the movement of people, hence, the latter frequently block the LOS in a dynamic manner. Thus, we perform the final evaluation of SALMA in a crowded office environment. To this end, we install SALMA-full in the office environment (Room A) to perform a 24-hours endurance test, while the system was exposed to the usual ongoing work flow. Hence, colleagues recurrently block the LOS or specular reflections between the anchor and three tags placed at representative evaluation points (#3, #15, #33). The latter are intentionally chosen to be in obstructed LOS conditions. During the 24 hours, every five seconds the position of the three tags is estimated by SALMA-full resulting in more than 50000 position estimates. Fig. 4.12 shows the positioning performance, where the dashed orange line represents the mean error and the solid blue line the 90% error over 50 position estimates. Additionally, we continuously track the number of present people in the room (green shaded step function). It is evident that the usual working environment (with two employees present) does not impair the performance of SALMA as it sustains an average error below 11.2 cm. Please note in Fig. 4.12 two prominent events: (i) a meeting with five people at 11:00 o'clock and (ii) at 14:00 o'clock a live presentation of SALMA with in total 15 people attending. We encouraged the audience to move around in the room freely during the presentation, thus, the LOS and the MPCs were obstructed in a dynamic fashion. Fig. 4.12 indicates that the position error increases during these events, still, even during the live presentation, the average error was below 34 cm and the 90% error below 79 cm [70]. These results demonstrate that SALMA is robust also in a highly dynamic environment and reinforce its resilience to NLOS situations. It is worth to mention that SALMA was also showcased at numerous conferences and public events achieving similar results [71].

Scalable and Responsive Positioning using Quasi-simultaneous Responses

Chapter 3 presents a real-time adaptation scheme to share data and timestamps between UWB devices in a robust and energy-efficient fashion, while Chapter 4 comprises a positioning system that exploits multipath propagation to cut down the required infrastructure to a single physical anchor, hence, to provide position information time- as well as cost-efficiently. Towards a robust, efficient, and scalable positioning system, this chapter, instead, provides solutions tackling the *scalability* property. The latter is of high relevance for future technologies to cope with billions of connected devices and the steadily increasing figures [187] emphasizing the importance of positioning systems that provide scalable operation, i.e., to support an increasing number of devices utilizing the positioning service without suffering from a performance degradation. Although designed in the 70's, still, GPS meets the requirements of a massively connected world due to a relentless performance that is independent of the number of users. Likewise, future indoor positioning systems should be designed with an enormous increase of connected devices in mind. However, most of the existing solutions based on UWB technology disregard the scalability property [158]. This chapter, instead, introduces an efficient *concurrent ranging* technique, which eventually leads to a highly scalable and responsive positioning system named *SnapLoc*.

Related Work. The multipath-assisted positioning system presented in Chapter 4 exploits multipath components present in the estimated CIR. Instead, in concurrent ranging and SnapLoc, respectively, significant peaks in the CIR do not necessarily originate from reflecting objects but rather from multiple distinct transmitters, which intentionally inject signal components in the CIR estimate by responding in a synchronized fashion to an initialization message. Conceptually, the principle may look similar to the synchronous transmissions proposed in narrowband radios [54, 55, 115, 242], which aim to exploit constructive interference between the transmitted signals of multiple devices to improve robustness and efficiency [242]. Instead, in the presented concepts, it is desired that the responses arrive quasi-simultaneously at the receiver, i.e., well enough separated in time to avoid interference and to allow resolving the responses in the es-

estimated CIR. Furthermore, the above-mentioned related works focus on communication, while the presented ideas aim to estimate the distance between a tag and its neighboring nodes as well as its position with a single read operation. To this end, instead of scheduling ranging operations between an initiator and multiple responders, the latter reply quasi-simultaneously to a single broadcast message. At the receiver these response signals are visible in the estimated channel impulse response and encode the relative distances between the responders. This principle and its feasibility was experimentally shown in [33]. In this thesis, instead, we address the remaining key challenges to make concurrent ranging applicable in real-world deployments.

As discussed in Chapter 4, several implementations of indoor positioning systems using low-cost UWB radios exist [79, 106, 117, 178]. However, the main objective of these systems is to achieve a high positioning accuracy and precision neglecting critical properties such as the position update rate and scalability in terms of tag density. For instance, Kempke et al. [106] achieve a 99% error in 3D of 53 cm with an update rate of 12 Hz. Silva et al. [178] report average errors between 5 and 40 cm in 2D at an update rate of 10 Hz, while the work by Hartmann et al. [79] attain an average error of 27 cm in 2D and an update rate of 50 Hz. These position update rates seem to be comparable to the responsiveness of GPS systems. However, the listed numbers refer to single-tag positioning, whereas the actual update rates decrease linearly with more tags. Hence, to derive the effective update rate, the given figures have to be divided by the number of tags. Instead, SnapLoc, the positioning system presented in this chapter, achieves a similar performance (90% error of 33.7 cm) at much higher update rates and with the ability to support an unlimited number of tags. This is enabled thanks to the ability of SnapLoc to carry out passive self-localization, i.e., tags are not actively transmitting a signal and remain anonymous.

These features are also provided by GPS systems and one of the key reasons for their long-lasting and enduring success [87]. Still, GPS requires the use of atomic clocks to maintain synchronization of the satellites and is not available indoors. To find passive alternatives for GPS-denied areas, positioning systems are often combined with device-free techniques, which do not need to attach tags to the tracked entities. However, due to the utilization of RSS measurements, these systems hardly achieve sub-meter accuracy [125, 165, 172] and they allow to track just a small number of objects [22]. Thus, to provide GPS-like scalable and accurate self-localization it is advised to make use of passive tags. To this end, systems based on ultrasound [46, 169], mobile crowdsourcing [27], passive RFID [130, 197], and optical sensors [156, 186] have been presented. Still, none of them provides decimeter-level accuracy as well as robustness in multipath-rich environments and NLOS conditions.

A UWB-based positioning system comparable to GPS has been introduced in [117], however, it requires a tight synchronization at the anchors and clock skew correction at the tag due to the use of sequential messages transmitted by the anchors. SnapLoc, instead, removes the need for tight synchronization and does not need a clock correction at the tag. Chorus [34] takes a similar approach as SnapLoc and exploits the concept of quasi-simultaneously transmitting anchor nodes to enable passive self-localization of the tags. However, Chorus focuses on modeling the uncertainty of low-cost UWB transceivers, while SnapLoc proposes techniques to overcome these limitations to enable its implementation on highly-constrained devices. Thereby, SnapLoc is the first positioning system employing quasi-simultaneous transmissions while achieving decimeter-level accuracy in real-world deployments.

Structure. The remainder of this chapter is structured as follows. Sect. 5.1 introduces in detail the concept of concurrent ranging as well as its open challenges and proposed countermeasures (Sect. 5.1.1). The open challenges include (i) to design an algorithm to reliably detect multiple responses in the CIR even in the case of overlapping signal components (Sect. 5.1.2), (ii) to identify responders, i.e., to associate a distance estimate to a specific responder (Sect. 5.1.3), and (iii) to mitigate the impact of multipath signals potentially leading to misclassification of MPCs as anchor responses (Sect. 5.1.4). Following, the concept of quasi-simultaneous responses is exploited to provide a highly scalable and fast indoor positioning system in Sect. 5.2. This includes a discussion of the limitations of current approaches (Sect. 5.2.1), a description of the design principles (Sect. 5.2.2) as well as techniques to overcome the limited timestamp resolution of low-cost UWB transceivers (Sect. 5.2.3) and an extensive evaluation (Sect. 5.2.4).

5.1 Concurrent Ranging

As discussed in Sect. 2.3, estimating the distance between two nodes in a non-synchronized network requires a two-way ranging (TWR) exchange by utilizing multiple packet transmissions. Indeed, in single-sided two-way ranging (SS-TWR), for instance, two messages (*INIT*, *RESP*) are exchanged between initiator and responder (see Fig. 2.4a and Fig. 5.1). Thus, estimating the pairwise distance between N nodes requires to schedule $N \cdot (N - 1)$ messages. Given the high current draw of UWB transceivers (see Chapter 3), it is a stringent requirement to reduce this high number of messages to make UWB feasible for energy-efficient location-aware IoT applications. Furthermore, scheduling the distance estimation to each neighbor results in a significant timing overhead, which increases channel utilization and traffic load. In highly dynamic settings, moreover, consecutively estimating the distances may result in inaccuracies.

The capability of UWB technology to resolve individual multipath components allows to build multipath-assisted positioning systems such as the one presented in Chapter 4. However, the high bandwidth of UWB transceivers, also, enables to extract the simultaneous responses of multiple transmitters from the estimated CIR. This principle is called *concurrent ranging* and its feasibility was shown experimentally by Corbalán and Picco [33]. Instead of scheduling several two-way ranging exchanges between an initiator node and other neighbors (responders), in concurrent ranging the initiator broadcasts an *INIT* message to all neighbors, which reply quasi-simultaneously with a *RESP* message after the time of flight of the packet and a pre-defined constant delay Δ_R , as illustrated in Fig. 5.1. Consequently, the *RESP* messages (containing the timestamps $t_{rx,j}$ and $t_{tx,j}$ in the payload) sent by the responders are overlapping in time. This allows to observe the signal components associated to each responder in the estimated CIR at the initiator. Thus, in contrast to SALMA (see Chapter 4), the peaks in the CIR do not necessarily originate from reflecting or scattering objects but from distinct transmitters. By detecting these intentionally injected pulses, it is possible to estimate the path delay and distance not only to a single neighbor, but to all responders concurrently. This implies that concurrent ranging reduces the total number of messages required to estimate distances between all neighbors from $N \cdot (N - 1)$ to N . In fact, the initiator has to broadcast just a single message and, due to the high current consumption during reception more importantly, to perform only a single read operation to aggregate all responses. This is especially relevant for battery-powered initiators.

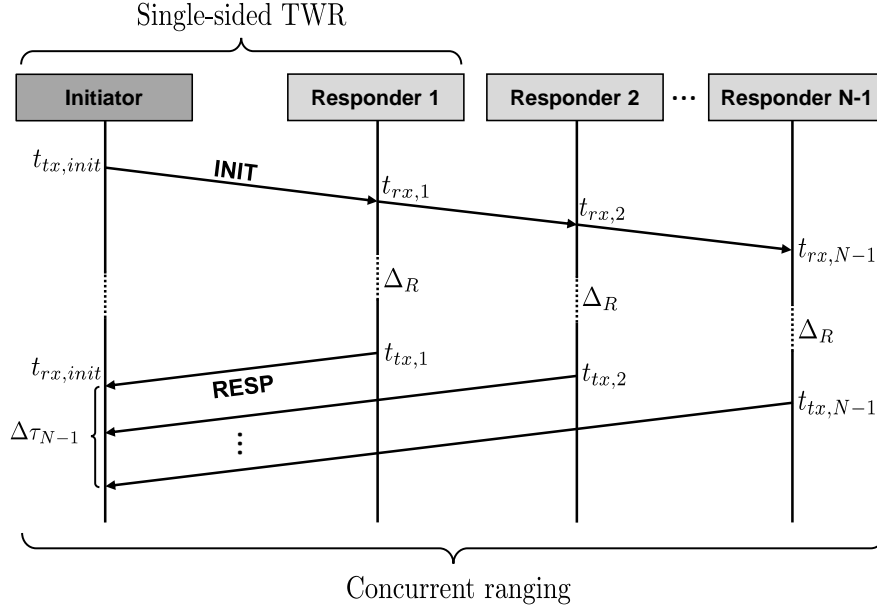


Figure 5.1: Principle of single-sided two-way ranging (SS-TWR) and concurrent ranging. In the latter, responders transmit *RESP* messages quasi-simultaneously. Adapted from [66].

Fig. 5.2a shows an acquired (normalized) channel impulse response when three neighbors are responding quasi-simultaneously. The responders are placed in a hallway at a distance from the initiator of $d_1 = 3m$, $d_2 = 6m$, and $d_3 = 10m$, respectively. Three significant peaks are visible in the CIR, representing the strongest signal component of each neighbor and encoding the relative distance between the responders. However, the absolute distance between the initiator and the closest neighbor is derived from the SS-TWR scheme (see Fig. 5.1), as it is still possible to decode one of the concurrently transmitted payloads containing the required timestamps [33]. The equation to calculate the distance between initiator and responder 1 follows as:

$$\hat{d}_1 = \frac{(t_{rx,init} - t_{tx,init}) - (t_{tx,1} - t_{rx,1})}{2} \cdot c \quad (5.1)$$

with c denoting the propagation speed in air. The distance between the initiator and the remaining responders instead is derived from the CIR. Due to the larger distance to the initiator node, the responding peaks of responder 2 and responder 3, respectively, arrive with a delay of $\Delta\tau_2 = 2 \cdot (\tau_2 - \tau_1)$ and $\Delta\tau_3 = 2 \cdot (\tau_3 - \tau_1)$ at the initiator, with τ_j ($j \in \{1, \dots, 3\}$) denoting the path delay between the initiator and each responder. The resulting delays $\Delta\tau_2$ and $\Delta\tau_3$ are due to the transmission of the *INIT* as well as the *RESP* message: thus, they have to be halved to correctly derive the distances. The estimated distances between initiator and responders, hence, follow as $\hat{d}_2 = \hat{d}_1 + \frac{c \cdot \Delta\tau_2}{2} = 6m$ and $\hat{d}_3 = \hat{d}_1 + \frac{c \cdot \Delta\tau_3}{2} = 10m$, respectively.




Figure 5.2: Principle of the proposed response detection algorithm. Adapted from [66].

5.1.1 Open challenges and contributions

Although the feasibility and potential of concurrent ranging was shown experimentally in [33], within this thesis the still open key challenges are addressed to make concurrent ranging applicable in real-world deployments. Within this section the open challenges are discussed, whereas each contribution to tackle these limitations are presented in one of the following sections.

Reliable detection of multiple responses. To practically implement concurrent ranging, it is required to process the estimated CIR at run-time and reliably detect the signal peaks associated to the different neighboring nodes. To this end, we propose a method to let initiators efficiently detect responses in the channel impulse response at run-time. Furthermore, the discussed method performs reasonably well in the case of overlapping responses (Sect. 5.1.2).

Identifying responders. A key challenge hampering the versatility of concurrent ranging is the impossibility to associate a signal peak in the CIR to a specific responder, i.e., the anonymity of the derived distance information. In previous works investigating the feasibility of concurrent ranging, all nodes are placed in a line topology [33], which gives the initiator prior knowledge about the order in which the responses are received in the CIR. This is impractical in real-world settings, as the relative locations of nodes are typically unknown especially in mobile setups. To overcome this limitation, we present a technique based on *pulse shaping* that allows to associate a distance estimate to a specific responder, so that ranging is no longer anonymous (Sect. 5.1.3).

Mitigating the impact of multipath reflections. Another critical challenge is to differentiate between a response and a strong multipath component from another responder. Corbalán et al. [33] suggest to use power boundaries based on the Friis equation to differentiate between responses in the channel impulse response and disturbing MPCs. However, this principle is not applicable in real-world applications due to three main issues [66]: (i) the Friis equation is idealized and does not hold true in typical UWB operational areas, (ii) in the case of an attenuated direct path, it is likely that multipath signals have higher amplitudes than the LOS component, (iii) the amplitude of the peaks in a CIR derived from low-cost UWB transceivers is highly varying. Thus, we propose the concept called *response position modulation*, which prevents the overlap of responses and strong multipath component from other responders (Sect. 5.1.4).

5.1.2 Reliable response detection

To make concurrent ranging feasible in real-world systems, it is essential to detect responses reliably in the CIR. To this end, we propose a scheme based on the *search and subtract* algorithm [51]. The latter employs a matched filter computing the correlation between the estimated CIR and a transmitted UWB pulse template with duration T_p . The pulse shape $s_{tx}(t)$ is derived from the previous chapter, where it is required in the signal model (see Eq. 4.2.1). To detect the $N - 1$ strongest responses in the CIR, we propose the following procedure [66]:

1. Upsample the CIR in order to improve the time granularity for further processing. Furthermore, to correct for the unknown time offset of the CIR derived from low-cost transceivers such as the Decawave DW1000 [112], the CIR is aligned with the response of the closest neighbor and hence the distance estimate \hat{d}_1 (see Eq. 5.1 and Fig. 5.2).
2. The time-discrete impulse response of the matched filter follows as the time-reversed pulse template: $\mathbf{h}_{MF} = [s_{tx}((N_p - 1) \cdot T_s), s_{tx}((N_p - 2) \cdot T_s), \dots, s_{tx}(0 \cdot T_s)]$ with T_s defining the sampling period, and $N_p = T_p/T_s$ the number of samples of the pulse. The output of the matched filter \mathbf{y} is computed as the discrete convolution ($*$) between the impulse response of the matched filter \mathbf{h}_{MF} and the observed CIR denoted as \mathbf{r}

$$\mathbf{y} = \mathbf{h}_{MF} * \mathbf{r}. \quad (5.2)$$

3. Identify the sample corresponding to the maximum of the matched filter output \mathbf{y} , indicating the index of the strongest path l_j . The latter relates to the path delay with $\tau_j = l_j \cdot T_s$ and to the path length with $d_j = \tau_j \cdot c$, respectively.
4. To reduce complexity, instead of the least squares solution suggested in [51], the estimated amplitude of the strongest path $\hat{\alpha}_j$ is defined as the amplitude of \mathbf{y} at sample l_j .
5. The estimated neighbor response ($\hat{\alpha}_j s_{tx}(t - \tau_j)$) is subtracted from the observed signal \mathbf{r} (dashed orange line in Fig 5.2a). Fig. 5.2b shows the matched filter output of the remaining signal after subtracting the strongest peak corresponding to the first neighbor response.
6. Repeat steps 2 to 5 with the remaining signal until $N - 1$ strongest paths are detected.
7. Independently of their amplitude $\hat{\alpha}_j$, the responses defined by the parameters $\hat{\alpha}_j$ and τ_j , respectively, are sorted in ascending order starting with the one of the closest neighbor. Being τ_1 the path delay of responder 1, the distance of responder j is estimated as:

$$\hat{d}_j = \hat{d}_1 + \frac{c \cdot (\tau_j - \tau_1)}{2}. \quad (5.3)$$

This algorithm allows to reliably detect the responses of all neighbors in the CIR and hence to resolve the first open challenge of concurrent ranging. Additionally, in [66] we have shown that the algorithm reliably detects responses from devices at similar distances, i.e., in the case of overlapping responses. To this end, we compared the algorithm with a conventional *threshold-based algorithm* as proposed by Falsi et al. [51]. The threshold-based algorithm compares the




Figure 5.3: Pulse shape $s_{tx}^{(i)}$ for different values of the TC_PGDELAY register. The x-axis marks the sample, whilst the y-axis indicates the normalized amplitude. Adapted from [66].

CIR with a defined threshold. As soon as the first sample of the CIR crosses this threshold, the maximum of the following N_p samples, i.e., the pulse duration, is derived. This is repeated until $N - 1$ peaks are detected. Both the proposed algorithm and the *threshold-based* one obtain a reliable performance when the responses are well-separated. However, as soon as the responses are overlapping due to a similar distance (and hence time of flight) of several responders, the proposed algorithm outperforms threshold-based algorithms. Indeed, when placing two responders at the same distance $d_1 = d_2 = 4\text{ m}$, the threshold-based algorithm detects both responses in only 48% of the tests, while the proposed algorithm is successful in 92.6% of the trials [66].

5.1.3 Encoding responder ID in the CIR using pulse shaping

To associate distance estimates to responders in concurrent ranging, it is required to encode the ID of each responder in the channel impulse response. Indeed, the preamble used to estimate the CIR consists of a fixed sequence of pulses independent of the responder (see Sect. 2.4) and hence does not embed the identity of the sender. Thus, the responses of the neighbors derived from the CIR do not contain any information allowing to associate the responses to the corresponding neighbors. Consequently, distance estimations in a traditional concurrent ranging scheme are anonymous. To solve this issue, we propose the concept of *pulse shaping*.

Pulse shaping. Up to now, the transmitted pulse shape $s_{tx}(t)$ remained unchanged. Fig. 5.3a shows the normalized pulse over the number of samples. In pulse shaping, instead, the transmitted pulse is changed as a function of the responder ID. This allows to associate signal components in the CIR to a responder. To this end, UWB radios provide the ability to change the shape of the transmitted pulses. For instance, the 8-bit register TC_PGDELAY in the DW1000 allows to regulate the width of the transmitted pulse. Typically, this is used to tune the output stage of the transceiver to meet the spectrum mask (see Sect. 2.2) as changing the shape effectively alters the output bandwidth [39]. While making the pulse narrower (i.e., increasing the bandwidth) is not an option due to the regulatory spectral mask, making the pulse wider, instead, does not violate the regulations. Fig. 5.3 exemplarily shows the pulse shape $s_{tx}^{(j)}$ obtained when configuring TC_PGDELAY with $0x93$ ($s_{tx}^{(1)}$), $0xC8$ ($s_{tx}^{(2)}$), $0xE6$ ($s_{tx}^{(3)}$), and $0xF0$ ($s_{tx}^{(4)}$), respectively. Note that $0x93$ corresponds to the employed default configuration hence $s_{tx}^{(1)} = s_{tx}$.




Figure 5.4: Estimated CIR (a) and matched filter output (b) in the case of two responders at $d_1 = 4m$ and $d_2 = 10m$, respectively, replying with different pulse shapes. Adapted from [66].

Identifying pulse shapes. Fig. 5.4a depicts an estimated CIR when one neighbor at a distance $d_1 = 4m$ responds using the default pulse $s_{tx}^{(1)}(t)$ (see Fig. 5.3a) and when a second neighbor at a distance of $d_2 = 10m$ responds with a wider pulse $s_{tx}^{(3)}(t)$ (see Fig. 5.3c). The different transmitted pulses are evident in the CIR. Executing the algorithm described in Sect. 5.1.2 for $N_{PS} = 3$ different pulse templates $s_{tx}^{(i)}(t)$ (with $i \in \{1, \dots, N_{PS}\}$) results in the matched filter outputs $y_i(t)$ illustrated in Fig. 5.4b. To detect the transmitted pulse shape of the responders and hence their ID, the estimated amplitudes of the neighbor responses $\hat{\alpha}_{j,i}$ (with j denoting the number of the response) of all N_{PS} matched filter outputs $y_i(t)$ are compared. The pulse shape i at which the amplitude $\hat{\alpha}_{j,i}$ is maximized, determines the pulse shape used by the responder [66]. Hence, in Fig. 5.4b, the first response corresponds to a transmitter employing the pulse template $s_{tx}^{(1)}(t)$ (blue, solid) and the second to a responder using the pulse template $s_{tx}^{(3)}(t)$ (purple, dotted). In [66] we have shown that in 99.2% of the cases the responder was successfully detected with $N_{PS} = 3$. Although in principle up to 108 concurrent responders are possible, for a higher N_{PS} the possibility of wrong classifications increases.

5.1.4 Mitigating the impact of multipath using response position modulation

As already discussed intensively throughout this thesis, in an indoor and multipath-rich environment it is common to observe several strong multipath components. In case of blocked and attenuated line-of-sight, the energy level of a received MPC might even exceed the energy of the LOS component. Thus, this challenges the concurrent ranging scheme as it is delicate to differentiate between a response from a neighbor and a dominant MPC from another responder.

Response position modulation. Using different pulse shapes as described in the previous section possibly mitigates the impact of strong multipath components, however, the most effective solution is to generally avoid that MPCs are overlapping with responses from other transmitters. To this end, we propose a technique called *response position modulation (RPM)*, which separates the responses of the neighbors in the time domain (in addition to the separation due to a different ToF), i.e., it modulates the position of the responses and hence the response delay Δ_R (see Fig. 5.1). Thus, we introduce a modulated response delay as $\Delta'_R = \Delta_R + \delta_j$, where δ_j sets an additional individual delay for each responder j . Doing so reduces the probability

of overlapping responses, as the latter are spread across the CIR estimate and are eventually assigned to a slot separated by the delay δ_j . Due to limited size of the CIR register in UWB transceivers, the separation of the responses and hence the number of slots spanned by the RPM mechanism is finite. For instance, the CIR estimate provided by the Decawave DW1000 has a length of 1016 samples with a sampling period of $T_s = 1.0016$ ns. Thus, the maximum offset is $\delta_{max} \approx 1017$ ns, which relates to a maximum distance offset of $\delta_{max} \cdot c \approx 307$ m. Consequently, knowing the maximum communication range and an estimate of the delay spread allows to define the number of non-overlapping responses fitting in the CIR register. In this thesis – since the exploitation of the individual delay δ_j affects the concurrency of the neighbor responses – we coin the term *quasi-simultaneous* responses instead of concurrent or simultaneous responses.

In Sect. 5.2, we will employ response position modulation to assign an individual delay δ_j to static anchors in a scalable indoor positioning system. Furthermore, we have shown in [66] that a combination of pulse shaping (see Sect. 5.1.3) and response position modulation may be used to increase the number of supported devices in the concurrent ranging scheme.

5.2 SnapLoc: From Concurrent Ranging to Scalable Positioning

The concept of concurrent ranging including the presented enhancements to make it applicable in real-world environments allows an energy-efficient and fast distance estimation within a network. Likewise, time-based positioning systems require time and distance estimates to multiple devices. Hence, within this section, we will utilize quasi-simultaneous responses to design *SnapLoc*, a UWB-based, highly scalable, GPS-like positioning system allowing an unlimited number of tags to self-localize at a theoretical position update rate of 2.3 kHz. To this end, Sect. 5.2.1 discusses the limitations of existing time-based ranging techniques as well as concurrent ranging in terms of scalable positioning. Following, the design principles and inner working mechanisms of SnapLoc allowing to overcome the limitations of current approaches as well as the implementation details of SnapLoc on a low-cost UWB platform are described in Sect. 5.2.2. Furthermore, Sect. 5.2.3 discusses the limited transmit timestamp resolution of off-the-shelf transceivers and presents methods to overcome this constraint and to enable decimeter-accurate position estimates. In Sect. 5.2.4, SnapLoc is evaluated in challenging indoor environments, showing that SnapLoc exhibits a 90% and median error of 33.7 cm and 18.4 cm, respectively, while preserving a high scalability and a high responsiveness.

5.2.1 Limitations of existing time-based positioning approaches

As discussed in the introduction of this chapter, most of the existing UWB-based positioning solutions focus on accuracy, often disregarding properties such as multi-tag support and position update rate [158]. Indeed, solutions based on two-way ranging do not scale in terms of tag density due to a large message overhead and the use of scheduling techniques to avoid packet collisions (Sect. 5.2.1.1). Positioning systems employing one-way ToA or TDoA, instead, require a tight synchronization of the anchors (Sect. 5.2.1.2). Furthermore, even the concept of concurrent ranging as suggested in Sect. 5.1 still suffers from limitations hindering its application in an indoor positioning system that allows to localize an unlimited number of tags (Sect. 5.2.1.3).

5.2.1.1 Limitations of TWR-based approaches

In two-way ranging (TWR)-based positioning systems, mobile tags derive the distance to static anchors in a sequential fashion, where each distance estimate demands multiple packet exchanges (see Sect. 2.3.3). Thus, this approach requires to allocate specific timeslots to each tag in order to avoid packet collisions. Such scheduling and collision avoidance techniques – which are incidentally also needed in the SALMA positioning system (see Chapter 4) – limit the number of supported tags and hence the scalability. Furthermore, the multiple packet exchanges per distance estimate result in a large communication overhead effectively decreasing the position update rate [108] as well as increasing the tag’s radio-on time. The latter heavily affects the tag’s energy consumption and thus its possible uptime in case of battery-powered tags. Additionally, sequentially estimating the distance to each anchor leads to inconsistent measurements in mobile and highly-dynamic settings eventually limiting the achievable positioning accuracy.

5.2.1.2 Limitations of TDoA-based approaches

As in one-way ToA ranging only a single message is exchanged per distance estimate, a tight clock synchronization between participating nodes is required (see Sect. 2.3.1). Exploiting the time difference of arrival (TDoA) between multiple anchors rather than the absolute time of flight, instead, relaxes the synchronization requirement at the tags (see Sect. 2.3.2). This simplifies the system design and its applicability while still allowing one-way communication. In case the tag initiates the position estimation by sending a broadcast message, in fact, only one message per position estimate is required [161, 193]. The broadcast message is received by the synchronized anchors, which compute the TDoA and estimate the tags’ position at a central positioning engine [162]. This is especially advantageous in tracking and monitoring applications, where the tags do not necessarily need to know their own position (see Sect. 1.1.2). Whilst such an approach allows to minimize the number of transmissions carried out by a tag, one still needs to allocate specific timeslots to each tag in order to avoid collisions limiting the scalability of the system. Letting the anchors initiate the position determination by sequentially broadcasting a signal, instead, allows the tags to carry out passive self-localization without actively transmitting packets [117]. However, this still requires a tight nanosecond-level synchronization between anchors, which is challenging [223] and typically increases message overhead [193, 227]). Furthermore, the anchors still send messages sequentially, which requires also the tag to correct the tags’ clock skew due to the long reception phase of the sequential messages.

5.2.1.3 Limitations of concurrent ranging

Exploiting quasi-simultaneous responses in a concurrent ranging scheme as discussed in Sect. 5.1 could potentially counteract the limitations of TWR as well as TDoA-based approaches. This is because it does not require any time synchronization between the participating nodes and because it reduces the communication overhead to a minimum (as just a single transmit and receive operation is required at the tag to estimate the distance to the anchors). However, it is hardly applicable in a scalable positioning system as (i) the number of users is not unlimited due to technical constraints, (ii) a high amount of payloads is lost when responders are located at similar distances, and (iii) low-cost UWB transceivers limit the possible timestamp resolution.

Concurrent ranging not fully scalable. The concept of pulse shaping presented in Sect. 5.1.3 to encode the responder ID in the estimated CIR allows to associate a distance estimate with the corresponding transmitter. However, the concept is restricted to 108 pulse shapes and hence concurrent responders. In practice, the limitation is even lower at around ten responders as the pulse shape is highly influenced by distortions due to an imperfect communication channel eventually leading to wrong classification of responders. On the other hand, the concept of response position modulation presented in Sect. 5.1.4 is highly effective to mitigate the impact of strong multipath components. However, the individual delay δ_j is assigned to mobile tags and hence the number of supported users is also limited due to the finite size of the CIR register in off-the-shelf UWB transceivers. Although combining RPM with pulse shaping as suggested in [66] allows to support a reasonable number of tags, however, their number is still limited. In SnapLoc, instead, one of the main design goals is to support an unlimited number of tags.

Corrupted timestamps due to packet collisions in SS-TWR. As illustrated in Fig. 5.1, concurrent ranging requires the estimation of the distance to the closest neighbor \hat{d}_1 using single-sided two-way ranging. That is because the remaining distances derived from the CIR are related to \hat{d}_1 (see Eq. 5.3). Thus, concurrent ranging requires to reliably detect the timestamps included in the payload of the closest neighbor's *RESP* message ($t_{rx,1}$ and $t_{tx,1}$ in Fig. 5.1). However, the quasi-simultaneous responses might corrupt the payload of the closest neighbor, particularly when one or more responders are located at similar distances, as we have demonstrated in [72]. Indeed, the measurements showed that even with just two responders placed at similar distance to the initiator up to 60% of all packets and hence concurrent ranging attempts are corrupted.

Limited transmit timestamp resolution. In concurrent ranging all responders transmit a *RESP* message after a predetermined delay $\Delta'_R = \Delta_R + \delta_j$. To this end, we use the *delayed transmission* feature of UWB transceivers such as the Decawave DW1000. It enables to set a future timestamp at which the transceiver actually sends a message and hence allows to align a pre-set timestamp with the real transmit timestamp and embed it in the message being transmitted ($t_{tx,i}$ in Fig. 5.1). Unfortunately, low-cost transceivers hardly achieve the required transmit timestamp resolution when using the delayed transmission feature to achieve decimeter-level distance estimations in concurrent ranging. Indeed, the DW1000 ignores the least significant 9 bits of the timestamp, effectively limiting the transmission resolution to approximately 8 ns [39, p. 26]. This is not an issue in the classical single-sided two-way ranging scheme, as the real transmit timestamp is anyway embedded in the message. However, this is not the case in concurrent ranging since there is no possibility to derive the real transmission timestamp for the responses acquired from the CIR estimate. Hence, this has a severe impact on the precision of concurrent ranging as it negatively affects the concurrency of the *RESP* messages from the neighbors.

5.2.2 Design principles

SnapLoc counteracts the aforementioned limitations of concurrent ranging and applies a modification of the latter to a TDoA-based and anchor-initiated positioning system. This enables that the system scales regardless of the tag density and without negatively impacting the system performance (Sect. 5.2.2.1). By employing quasi-simultaneous responses from the anchors, SnapLoc further rectifies the disadvantages of current TWR or TDoA based approaches. Indeed,




Figure 5.5: Concurrent ranging applied in an anchor-based positioning system. The tag (T) initiates a position estimation by broadcasting an *INIT* message to the simultaneously responding static anchors ($A_1 \dots A_4$). Taken from [72].

SnapLoc reduces the radio-on time at the tag to a single receive operation allowing the tags to passively localize themselves requiring neither a tight synchronization of the anchors nor a clock correction at the tag. To avoid misclassification of strong multipath components as dedicated anchor responses in the estimated CIR, we employ the response position modulation introduced in Sect. 5.1.4 to assign an individual delay in the nanosecond range to each anchor (Sect. 5.2.2.2). Following, Sect. 5.2.2.3 discusses the detection of the anchor responses within the CIR as well as the estimation of the TDoA and, consequently, the position of the tags $\mathbf{p}^{(i)}$. Finally, Sect. 5.2.2.4 describes the implementation details of SnapLoc on a low-cost UWB platform.

5.2.2.1 Applying concurrent ranging in an anchor-based positioning system

Instead of estimating the distance between multiple mobile tags, in SnapLoc quasi-simultaneous responses are employed to estimate the TDoA between multiple anchors by solely reading and analyzing a single CIR estimate. Fig. 5.5 illustrates this concept with $N = 4$ anchors A_j ($j = 1 \dots N$) and one tag T . The latter initiates a position estimation by broadcasting an *INIT* message to all surrounding anchors (solid arrows), which simultaneously respond with a *RESP* message after a constant delay Δ_R (dashed arrows). The time offsets of the anchor responses $\Delta\tau_{i,j}$ in the estimated channel impulse response contain information related to the position of the tag, namely the TDoA between the anchors A_i and A_j ($i \neq j$):

$$\Delta\tau_{i,j} = 2 \cdot (t_{j,T} - t_{i,T}). \quad (5.4)$$

Consequently, the time difference of arrival $\Delta t_{i,j}$ follows as:

$$\Delta t_{i,j} = t_{j,T} - t_{i,T} = \frac{\Delta\tau_{i,j}}{2}. \quad (5.5)$$




Figure 5.6: In SnapLoc a reference anchor A_R sends an *INIT* message, to which all surrounding anchors ($A_1 \dots A_4$) reply quasi-simultaneously with a *RESP* message. Adapted from [72].

Due to the exploitation of the TDoA instead of raw distance estimations, this approach removes the need to carry out a SS-TWR exchange and hence to correctly receive the timestamps embedded in the *RESP* message, which is one of the key limitations of concurrent ranging, as outlined in Sect. 5.2.1.3. Thus, the tags can estimate the TDoA between multiple anchors using only information contained in the CIR estimated from a single read operation.

Enabling the system to scale. The aforementioned approach allows the tag to initiate a position estimation individually and aperiodically. However, this requires to actively transmit an *INIT* message. Hence, in order to avoid collisions between sending tags, one would still need to allocate specific timeslots to each tag. This decreases the scalability of the system, as described in Sect. 5.2.1.2. Therefore, in SnapLoc we use an *anchor-initiated* approach in which one of the anchors is selected to act as the initiator broadcasting the *INIT* message, as illustrated in Fig. 5.6. This initiating anchor is referred to as reference anchor A_R in the remainder of this thesis. The key advantage of such an anchor-initiated TDoA approach is that the tag is not actively involved in the communication and, thus, no scheduling of the transmission between multiple tags is required. Hence, similar to GPS, this approach allows *passive self-localization*. This also enables tags to remain anonymous and maximize their privacy, as well as to achieve a constant system performance independent of the number of users exploiting the service.

5.2.2.2 Reusing response position modulation

The SnapLoc system consists of N anchors placed at known positions $\mathbf{a}^{(j)} \in \mathbb{R}^3$ (with $j = 1, \dots, N$) to localize N_p tags located at an unknown position $\mathbf{p}^{(i)} \in \mathbb{R}^3$ (with $i = 1, \dots, N_p$). One of the anchors is selected as reference anchor A_R to broadcast the *INIT* message, to which the remaining $N - 1$ anchors responds quasi-simultaneously¹ (see Fig. 5.6). To avoid overlaps

¹To reduce the number of anchors, in fact, A_R could also respond with a *RESP* message to the *INIT* message.




Figure 5.7: Using response position modulation by assigning an individual delay δ_j to each anchor, the responses are well separated in the estimated CIR. Taken from [72].

of anchor responses and MPCs in the estimated CIR, we employ response position modulation, as described in Sect. 5.1.4. To this end, an individual delay δ_j at each anchor is used to separate the responses in time. Fig. 5.7 shows an estimated CIR when four anchors are responding quasi-simultaneously. Although the first response exhibits strong MPCs, it is possible to distinguish them from the remaining anchor responses due to the additional individual delay δ_j . Note that in the concurrent ranging scheme presented in Sect 5.2.1.3 the individual delay δ_j is assigned to mobile tags: this limits the number of users that can be supported and hence the scalability of the system (see Sect. 5.2.1.3). In SnapLoc, instead, the individual delay is assigned to the anchors, whose number is limited and typically remains constant over the service life. Still, due to the limited length of the CIR register in common UWB transceivers, there is a trade-off between how much the anchor responses can be separated and the number of supported anchors. In SnapLoc, we assign each anchor the individual delay $\delta_j = (j - 1) \cdot \beta$, where β represents the size of the assigned slot and is set to $\beta = 128$ ns. The latter relates to a distance offset of $\beta \cdot c \approx 38.4$ m. Due to path- and reflection losses, this separation makes it unlikely that a strong MPC of an earlier response interferes with the current response and allows to use up to eight anchors when using the DW1000 transceiver (due to the maximum offset of $\delta_{max} \approx 1017$ ns, see Sect. 5.1.4). In case the anchor density is insufficient, one can reduce β to increase the number of supported anchors, however, this is not advised in multipath-rich environments. Instead, similar to mobile networks such as GSM, the area of operation can be divided in cells, where neighboring cells are assigned to different channels, preamble codes, and/or time slots (see Chapter 6) [72].

5.2.2.3 Estimating response time, time difference of arrival, and position

As discussed in Sect. 5.2.2.1, SnapLoc employs an *anchor-initiated* approach, where a reference anchor A_R broadcasts the initialization message. The remaining anchors (marked as $A_1 \dots A_4$ in Fig. 5.6) respond quasi-simultaneously with a *RESP* message after a delay $\Delta_R + \delta_j$ (with $j = 1 \dots 4$). A tag T within the communication range of the anchors listens to the response signals sent by the anchors and has to reliably detect the responses in the estimated channel impulse response. To this end, the algorithm described in Sect. 5.1.2 might be exploited.

Estimating response time. However, the modified *search and subtract* algorithm searches for K strong MPCs in the entire CIR estimate. Instead, in SnapLoc we are interested in the first path component of each slot defined by the individual delays δ_i (see Fig. 5.7). Thus, the threshold-based algorithm is sufficient for this scenario. To this end, after upsampling the estimated CIR and employing a matched filter, the first sample n_j of each slot of the matched filter output exceeding a given threshold S_{th} indicates the first path of each anchor response. The threshold S_{th} is set as the 10-fold power of the noise floor. Consequently, the estimated time difference of the responses $\Delta\hat{\tau}_{i,j}$ is determined by $\Delta\hat{\tau}_{i,j} = (n_j - n_i) \cdot (T_s/L)$ with $(i \neq j)$ and T_s denoting the sampling period and L the upsampling factor [72].

Estimating time difference of arrival. Similarly to the tag-initiated approach discussed in Sect. 5.2.2.1, the responses in the estimated CIR encode information related to the time difference of arrival between the responding anchors. However, due to the anchor-initiated nature in SnapLoc, the time difference $\Delta\tau_{i,j}$ of the responses corresponding to A_i and A_j follow as [72]:

$$\Delta\tau_{i,j} = (\delta_j - \delta_i) + (t_{R,j} + t_{j,T}) - (t_{R,i} + t_{i,T}). \quad (5.6)$$

Given that the individual delays δ_i and δ_j as well as $t_{R,i}$ and $t_{R,j}$ (due to the static nature of the anchors) are known, the TDoA $\Delta t_{i,j}$ between the anchors A_i and A_j follows as:

$$\Delta t_{i,j} = t_{j,T} - t_{i,T} = \Delta\tau_{i,j} - (\delta_j - \delta_i) - t_{R,j} + t_{R,i}. \quad (5.7)$$

Estimating position. After deriving the time difference of arrival $\Delta t_{i,j}$ between the anchors from the channel impulse response, the tags have acquired all the necessary information to estimate their unknown position $\mathbf{p}^{(i)}$ using TDoA multilateration [235]. To find the position $\hat{\mathbf{p}}$ minimizing the squared error of the observed TDoA estimates, we use an iterative quasi-Newton method [175] with an initial position estimate chosen at the center of the room [72].

5.2.2.4 Implementation

We implement SnapLoc on the Decawave DWM1001 development board [40] and a self-made low-cost UWB platform². The latter is based on the DW1000 transceiver and the STM32 Nucleo-64 board, which employs a low-power ARM Cortex-M3 based STM32L152RE microcontroller [14, 72]. Every node uses an off-the-shelf omnidirectional UWB dipole antenna. Note that the employed hardware highly influences the minimum response delay $\Delta_{R,min}$ at the anchors, i.e., the minimum time necessary to switch from receiving the *INIT* message to responding with the *RESP* message. It is mainly defined by the SPI communication overhead and delays introduced by the processing of a packet reception and transmission. Additionally, since the first symbol of the physical layer header (PHR) determines the transmit timestamp [38], the minimum applicable response delay $\Delta_{R,min}$ is influenced by the duration of PHR and payload of the *INIT* message, as well as the duration of preamble and SFD of the *RESP* message. Overall, this deterministic contribution to the minimum response delay $\Delta_{R,min}$ corresponds to a delay of 178.5 μs . In [72], we have evaluated the non-deterministic part, namely, the minimum switching time experimentally. It showed that the switching time is approximately 100 μs , however,

²<https://github.com/BGWH/NetLoc>

when employing a constrained microcontroller with lower CPU and SPI speed, e.g., the self-made UWB platform, the resulting minimum response delay increases to $\Delta_{R,min} \approx 850 \mu s$. This delay affects the position update rate achievable by SnapLoc, as shown in Sect. 5.2.4.2. Furthermore, the longer response time also causes clock drifts at the anchors. Hence, to support constrained devices and low-cost oscillators, we have presented in [72] an optional clock correction scheme for the anchors. We configure the DW1000 to use channel 4, i.e., a bandwidth of 900 MHz and a carrier frequency of 3.9936 GHz (see Table 3.1), maximum data rate (6.8 Mbps), a pulse repetition frequency of 64 MHz, as well as a preamble symbol repetition of 128.

5.2.3 Improving timestamp resolution

To respond quasi-simultaneously at the anchors to the *INIT* message, we employ the delayed transmission feature provided by the Decawave DW1000 transceiver. According to the common response delay Δ_R and the individual delay δ_j , it allows to set the timestamp at which the *RESP* message is transmitted. Although the DW1000 transceiver represents receive (RX) and transmit (TX) timestamps as 40-bit values with a resolution of 15.65 ps [38], it ignores the least significant 9-bits when performing delayed transmissions, as discussed in Sect. 5.2.1.3. This lowers the effective transmission resolution from (theoretical) 15.65 ps to $4/(499.2 \cdot 10^6) \approx 8$ ns. Without correction, this transmission uncertainty results in a uniformly distributed and memoryless error $e^{TX} \sim \mathcal{U}(-8 ns \cdot c, 0)$ of the timestamps derived from the estimated CIR. Considering that an error of 1 ns in the time domain results approximately in an error of 30 cm in the spatial domain, it is evident that e^{TX} highly affects the positioning performance, as we show experimentally in Sect. 5.2.4. Thus, to sustain a decimeter-level accuracy in SnapLoc, we propose two techniques to improve the transmit timestamp resolution [72].

Wired correction. Since the DW1000 stores the timestamps as 40-bit values, the lost 9-bits are known in the transceiver. Hence, the first correction method records the lost 9-bits at each anchor and sends these correction values back via a wired backbone to the reference anchor A_R . This method is referred to as wired correction or optimal correction as all missed bits are reproduced. Alternatively, the correction values may also be transferred to the reference anchor A_R via a different wireless technology [185]. Following, A_R broadcasts the missing transmit timestamp information embedded in the next *INIT* message to all nearby tags to avoid additional messages. The tags then correct the timestamps of the anchor responses derived from the previous channel impulse response. Hence, this causes a delay by one initialization interval T_{init} . Due to the high update rate of SnapLoc, this trade-off is tolerable, as discussed in Sect. 5.2.4.2.

Wireless correction. Since the wired correction method recovers all the missing 9-bits of the transmit timestamp at each anchor, it is the optimal correction. However, it requires a backbone network to send the correction values back to the reference anchor. Thus, to facilitate a fast and simple setup of SnapLoc, we propose another technique to increase the timestamp resolution that does not require a wired connection between the anchors and reference anchor. So far, the reference anchor A_R was used to trigger a position estimation by sending an *INIT* message and afterwards may act as a regular anchor by responding to its own initialization message. In the *wireless correction* scheme, instead, A_R listens to the responses of the anchors and derives the estimated CIR, similar to a regular tag. Consequently, from the perspective of the reference

anchor, this behaves as a *tag-initiated* approach (see Sect. 5.2.2.1 and Fig. 5.5). Thus, according to Eq. 5.5, the time difference of arrival between the anchors can be derived by analyzing the estimated CIR. As the anchors are static and their positions are known beforehand, the estimated distance information can be compared with the true values. Deviations of the estimated values from the true values are treated as errors due to ignoring the least significant 9-bits of the transmit timestamp. To recover the lost precision, we differentiate between the correction at anchor A_1 and the remaining anchors. This is due to the fact that the response of anchor A_1 corresponds to the first peak in the CIR and is hence detected by the internal leading edge or first path detector of the transceiver [39]. For the correction at anchor A_1 , we define the transmit error due to the limited timestamp resolution $e_{A_1}^{TX}$ as the difference between the true round trip time $t_{R,1}^{RT}$ between A_1 and A_R and the estimated one $\hat{t}_{R,1}^{RT}$:

$$e_{A_1}^{TX} = t_{R,1}^{RT} - \hat{t}_{R,1}^{RT}. \quad (5.8)$$

The true round trip time $t_{R,1}^{RT}$ is defined by

$$t_{R,1}^{RT} = 2 \cdot t_{R,1} + \Delta_R + \delta_1 + 2 \cdot \Theta_a, \quad (5.9)$$

where $t_{R,1}$ is the time of flight between A_R and A_1 , Δ_R the common response delay for all anchor nodes, δ_1 the individual delay of A_1 , and Θ_a the antenna delay. The latter is required to correct for delays introduced by the antenna, PCB, and internal and external components [39, p.205 ff.]. We have determined Θ_a experimentally for each type of used UWB platform [72]. The estimated round trip time $\hat{t}_{R,1}^{RT}$ is defined as the difference of the receive timestamp $t_{resp,1}^{RX}$ of A_1 's *RESP* message and the transmit timestamp t_{init}^{TX} of the *INIT* message at the reference anchor. Therefore, the TX timestamp error of the first anchor $e_{A_1}^{TX}$ follows as [72]:

$$e_{A_1}^{TX} = (2 \cdot t_{R,1} + \Delta_R + \delta_1 + 2 \cdot \Theta_a) - (t_{resp,1}^{RX} - t_{init}^{TX}). \quad (5.10)$$

The transmit timestamp error of the remaining anchors $e_{A_j}^{TX}$ ($j = 2, \dots, N$) is defined by the true TDoA $\Delta t_{j,1}$ between the j^{th} anchor A_j and A_1 and the one estimated from the CIR $\hat{\Delta t}_{j,1}$:

$$e_{A_j}^{TX} = \Delta t_{j,1} - \hat{\Delta t}_{j,1}. \quad (5.11)$$

The true time difference of arrival $\Delta t_{j,1}$ is derived from the known positions of the reference anchor A_R and the j^{th} anchor A_j , which follows as:

$$\Delta t_{j,1} = t_{R,j} - t_{R,1} \quad (5.12)$$

where $t_{R,j}$ is the time of flight between A_R and A_j . The estimated TDoA $\hat{\Delta t}_{j,1}$ is derived from the CIR according to Eq. 5.5 and has to be corrected by the previously acquired transmit error of the first anchor $e_{A_1}^{TX}$. Thus, the resulting error of the j^{th} anchor A_j is:

$$e_{A_j}^{TX} = (t_{R,j} - t_{R,1}) - (\hat{\Delta t}_{j,1} + e_{A_1}^{TX}). \quad (5.13)$$

The resolution of the error value $e_{A_j}^{TX}$ is restricted by the sampling period of the estimated CIR $T_s = 1.0016 \text{ ns}$. Thus, 3-bits in the *INIT* message broadcasted by the reference anchor are enough to represent the error correction value. Therefore, the overhead due to a longer packet size is slightly shorter in the *wireless correction* method compared to the *wired correction*.

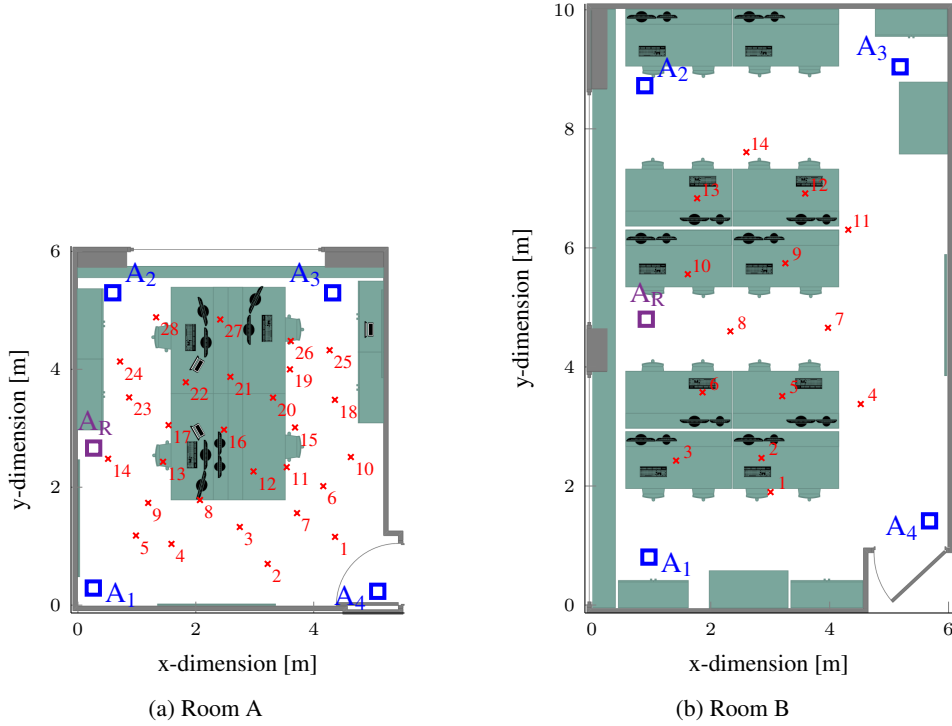


Figure 5.8: Evaluation setup: Two environments with 28 and 14 evaluation points, respectively, four anchors (blue square) and a reference anchor (purple square). Adapted from [72].

5.2.4 Evaluation

We evaluate SnapLoc experimentally in the same challenging office environment as SALMA (Room A, see Fig. 5.8a) as well as in a larger laboratory classroom (Room B, see Fig. 5.8b). Sect. 5.2.4.1 describes the experimental setup in more detail. This is followed by an analysis of the achievable position update rate and the energy efficiency in terms of over-the-air time in Sect. 5.2.4.2. We then evaluate the performance of SnapLoc in Sect. 5.2.4.3, showing that it can achieve decimeter-level positioning accuracy when employing quasi-simultaneous responses and the proposed methods to improve the transmit timestamp resolution.

5.2.4.1 Experimental setup

The evaluation is performed in an office with a size of $5.2 \times 6.03 \text{ m} \approx 31.36 \text{ m}^2$ (Room A, see Fig. 5.8a) and a larger laboratory classroom with $6.05 \times 10 \text{ m} = 60.5 \text{ m}^2$ (Room B, see Fig. 5.8b). Both rooms contain several scattering and reflecting objects such as monitors, desks, and chairs. The reference anchor (purple square) and the other anchors (blue squares) are placed on tripods at known static positions and at a height of 1.60 m, i.e., in the same 2D plane. All evaluations are performed with only $N = 4$ anchors in order to examine SnapLoc using minimal infrastructure. The evaluation points ($N_{\text{EP}} = 28$ in Room A and $N_{\text{EP}} = 14$ in Room B) are

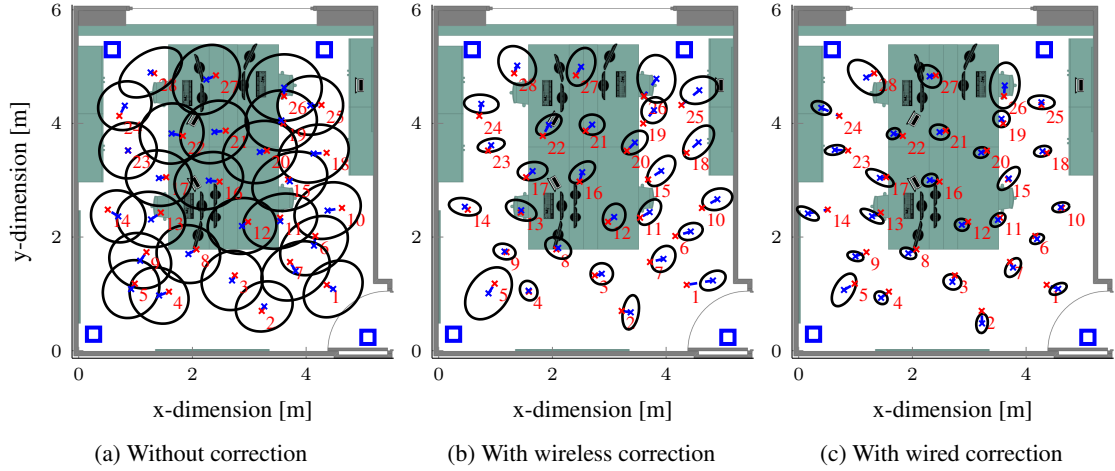


Figure 5.9: Mean position estimate (blue crosses) and standard deviation (black error ellipses) for each evaluation point in Room A without correction of the transmit timestamp (a), with the wireless correction (b), and with the wired correction (c). Adapted from [72].

randomly distributed in the rooms. At each evaluation point, 500 position estimates are observed. The absolute error of each positioning trial is calculated as the Euclidean distance between the position of the evaluation point \mathbf{p}_{EP} and the i^{th} position estimate $\hat{\mathbf{p}}_i$.

5.2.4.2 Position update rate and efficiency

The simultaneous acquisition of the anchor signals within a single read operation and without transmitting a packet allows tags in SnapLoc to minimize the radio-on time and hence the energy consumption. Indeed, acquiring the necessary information to self-localize a tag requires only approximately $82.4 \mu\text{J}$ [72]. Besides a low energy consumption, quasi-simultaneously responding anchors also highly affect the achievable position update rate, as the latter relates to the total time needed to provide the tag with the required information to estimate its position. In SnapLoc, this total time consists of the duration of *INIT* and *RESP* messages, as well as the time to switch between receive and transmit mode at the anchors. According to [72] and Sect. 5.2.2.4, the switching time is approximately $100 \mu\text{s}$ when using the DW1000 evaluation modules and the duration of the *INIT* and *RESP* message with the settings listed in Sect. 5.2.2.4 is roughly $334 \mu\text{s}$. Thus, deriving the information required to estimate the tag’s position takes only $434 \mu\text{s}$ overall. Theoretically, this enables an update rate of more than 2.3 kHz for SnapLoc, without any restrictions on the number of tags. Even when using our self-made hardware with the constrained microcontroller presented in [72], we still achieve an update rate of about 996 Hz . This should be considered as an upper bound, as it ignores the execution time of the positioning algorithm and the impact of the methods to improve the timestamp resolution.




Figure 5.10: Performance of SnapLoc depending on the method used to correct the limited transmit timestamp resolution in the two rooms used in our evaluation. Adapted from [72].

5.2.4.3 Positioning performance and effectiveness of TX uncertainty correction

We evaluate next the performance of SnapLoc in terms of positioning accuracy and precision, as well as the effectiveness of the methods to overcome the limited transmit timestamp resolution proposed in Sect. 5.2.3 in two different environments.

Effectiveness of TX uncertainty correction. The first evaluation of SnapLoc is performed in a smaller office (see Fig. 5.8a) using $N_{EP} = 28$ evaluation points. Fig. 5.9 shows the mean (blue crosses) and the standard deviation (black error ellipses) for the 500 position estimates at each evaluation point individually for the transmit timestamp correction methods presented in Sect. 5.2.3. The mean indicates the estimation bias (distance to true position marked with red crosses); the error ellipses, instead, give insights in the radial (facing the LOS) as well as the tangential (perpendicular to the LOS) error. Fig. 5.9a shows the performance of SnapLoc *without transmit timestamp correction*. Fig. 5.9b shows the performance of SnapLoc with the *wireless correction*, whilst Fig. 5.9c illustrates it with the *wired correction*. As clearly indicated, the latter performs best, as it recovers all the ignored 9-bits of the transmit timestamp at all anchors. The *wireless correction*, instead, restores theoretically a time resolution of 15.56 ps for anchor A_1 and a resolution of 1 ns for the remaining anchors, thus, its performance is slightly worse compared to the *wired correction* method. Without any correction, each of the transmit timestamps at the anchors has a resolution of just 8 ns, which reduces the positioning precision, as illustrated by the larger ellipses in Fig. 5.9a. Moreover, it is noticeable in Fig. 5.9b and Fig. 5.9c that the evaluation points within a distance of 1.5 m to an anchor ($EP \in \{1, 5, 24, 25, 26, 28\}$) perform worse than those located further away from the anchors. This is due to the high signal strength of the close anchor, which causes the CIR accumulation and register to saturate [72]. As the amplitude of the other anchor responses remains relatively low, a correct response detection is impaired. Thus, when deploying SnapLoc, a distance of at least 1.5 m between the tag and the anchors should be maintained. This is anyway typically the case in indoor positioning systems, as anchors are usually mounted close to the ceiling.

Overall positioning accuracy and precision. To derive statistically meaningful observations of the overall performance of SnapLoc, we derive its accuracy and precision using the cumulative distribution function (CDF) over the error Err_i of all position estimates. Due to the saturation

effects at tag positions close to the anchors, we have ignored the corresponding evaluation points $EP \in \{1, 5, 24, 25, 26, 28\}$ for this analysis. Fig. 5.10a shows the CDFs depending on the used method to correct the limited TX timestamp resolution. *Without correction* (dotted purple line), a 90% error of 1.15 m and a median error of 0.68 m was achieved. Instead, the integration of *wireless correction* allows to reduce the 90% error to 55.8 cm and the median error to 25.4 cm (dashed orange line) and the *wired correction* even reaches a 90% error of just 33.7 cm and a median error of 18.4 cm (solid blue line). Thus, by using the proposed correction methods, SnapLoc achieves decimeter-level accuracy despite the limited transmit timestamp resolution of 8 ns and the CIR sampling period of about 1 ns.

Performance in larger rooms. To confirm the plausibility of the results in Room A and to show that the performance of SnapLoc is not degrading in larger areas, we carry out an evaluation in a laboratory classroom (see Fig. 5.8b) that is significantly larger than the previously employed office room ($31.36 m^2$ vs. $60.5 m^2$). Fig. 5.10b shows the CDF of all position estimates in the $N_{EP} = 14$ evaluation points shown in Fig. 5.8b. Without using a transmit timestamp correction, the 90% error is 1.30 m and the median error 0.73 m. The *wireless correction* reduces the 90% error to 74 cm and the median error to 22.3 cm. With the *wired correction* or optimal correction, SnapLoc achieves a median error of 17 cm and a 90% error of 35.2 cm. The slight differences compared to the evaluation in Room A are mainly due to the presence of a few more outliers with a position error above 0.5 m, as indicated in the longer tail in Fig. 5.10b. Still, the results are consistent to the evaluation in Room A despite the use of a larger area. We can hence conclude that SnapLoc is easily applicable to larger areas of operation. Additionally we have performed in [72] an evaluation while moving the tag around freely: this shows that SnapLoc achieves reasonable results also in dynamic conditions even without using any (tracking) filter on the position estimates, such as a Kalman or particle filter.

Conclusions and Outlook

The fact that we spend up to 90% of the day in enclosed buildings and the highly increasing number of connected devices emphasize the demand for an indoor equivalent to satellite-based positioning systems to provide humans, robots, machines, and objects with seamless information about their current position and status. Towards this goal, this thesis provides solutions and techniques to provide robust, efficient, and scalable positioning information. We implement each of these contributions on low-cost UWB platforms and evaluate them in real-world environments.

We first present a scheme to adapt UWB physical layer settings at runtime to provide *robust* and *energy-efficient* UWB communication even in dynamic and harsh environmental conditions. To this end, we characterize experimentally the impact of PHY settings on the performance of UWB communication showing that one can gain up to 8 dB additional link margin. Combined with measurements of the energy consumption, this allows the proposed adaptation scheme to privilege PHY configurations in order to optimize link reliability at minimal energy costs. To trigger parameter adaptations, we propose a UWB link state indicator providing the adaptation algorithm with real-time information about the link quality. The link state indicator exploits the estimated channel impulse response to derive the received signal strength and to extract relevant information about the characteristics of the surrounding environment. Evaluations using the Decawave DW1000 transceiver show the effectiveness of the adaptation scheme to maintain reliable and efficient UWB communication laying the foundation for a robust indoor positioning system, which requires to reliably share packets and timestamps.

Next, we exploit multipath information to cut down the required infrastructure of a UWB-based indoor positioning system to a single physical anchor. Therefore, differently to state of the art systems, not only the first path component of the CIR is exploited but also the position-related information provided in the multipath components. The system called SALMA does not need a time-consuming setup phase, prior calibration or training phase, instead, it solely requires a floor plan and the position of the anchor to model the multipath propagation. Additionally, by using directional UWB antennas, we exploit the angular domain of the MPCs and manage to increase the robustness of SALMA against overlapping MPCs. Indeed, the evaluations show that SALMA achieves a median position error below 8 cm and 90% of all position estimates

exhibit an error below 20.17 cm. Even under obstructed LOS conditions, SALMA attains a high positioning performance with a 90% error still below 30.7 cm. Thus, the exploitation of multipath components allows to decrease the deployment effort and hence to increase the *time-* and *cost efficiency*, while still allowing to sustain a high accuracy and precision.

Finally, the last contribution tackles the scalability property by presenting an energy-efficient and fast ranging method followed by a *scalable* indoor positioning system. The concept of both is similar to the presented multipath-assisted positioning, however, the significant components in the estimated CIR do not originate from reflecting objects but rather from multiple distinct transmitters, which intentionally inject signal components in the CIR estimate. Within this doctoral thesis we exploit this concept to provide a practical solution for concurrent ranging to enable an efficient distance estimation to multiple users in parallel. To this end, we first propose an algorithm to reliably detect responses in the estimated CIR at runtime. Second, we introduce novel techniques called pulse shaping and response position modulation to allow the association of a distance estimate to an individual responder and to mitigate the impact of strong multipath components on the response detection. Following, we apply the enhanced concept of concurrent ranging to develop a highly scalable and responsive indoor positioning system named SnapLoc. Instead of multiple mobile transmitters responding to an initialization message, the static anchor nodes are responding quasi-simultaneously. This enables tags in the communication range to acquire the necessary information to estimate their position within a single receive operation. In contrast to classical TDoA-based systems, this avoids the necessity of a tight synchronization of the anchors and correction of clock deviations at the tag. To overcome an intrinsic limitation of low-cost UWB transceivers regarding their achievable transmit timestamp resolution, we present correction techniques allowing to achieve decimeter-level positioning accuracy. Indeed, SnapLoc achieves a 90% error of 33.4 cm and a median error of 18.4 cm at theoretical update rates of up to 2.3 kHz and independent of the number of tags utilizing the positioning service.

Considering the initial thesis statement in Sect. 1.3, this doctoral thesis confirms that the high time resolution and the channel information provided by UWB transceivers enables robust, efficient, and scalable indoor positioning. Towards robustness, it is required to adapt to environmental changes at runtime by gathering and analyzing the channel information. The latter furthermore enables to employ virtual anchors instead of physical anchors by exploiting not just the first path component, but, additionally, the position-related information contained in multipath components. Finally, towards scalability and high responsiveness of the positioning system, we recommend to exploit quasi-simultaneous responses of multiple transmitters.

6.1 Future work

Although this doctoral thesis strives for robust, efficient, and scalable indoor positioning, still, the presented work has a number of limitations and open issues, which might be tackled in future work. We present these potential improvements of the thesis in this section.

Network-wide parameter adaptation. The PHY parameter adaptation presented in Chapter 3 is solely discussed in a single link setup and the current implementation considers only one global PHY configuration corresponding to the worst link state in the network. This ensures that all communication links achieve a minimum target link quality. However, in a large (multi-hop)

network it is likely that some nodes experience a higher link quality than necessary, hence, the system resources in terms of energy budget can be used more efficiently by employing local optimal PHY configurations. To this end, each node may maintain a neighbor table with individual PHY settings for each communication link or one may divide the network in clusters, where the latter use local PHY parameters satisfying the requirements of all members of the cluster. Furthermore, we did not integrate the adaption scheme in one of the presented positioning systems.

Multipath-assisted quasi-simultaneous positioning. In SnapLoc, multiple anchors inject signal components in the estimated CIR of the tag, which in fact corresponds to a simultaneous CIR estimation of multiple communication links, i.e., the ones between the tag and each responding anchor. This multipath information might be used to perform a NLOS or destructive interference detection of multiple anchors by analyzing the CIR estimate acquired from a single receive operation (see Sect. 3.2.2). In addition, one can combine the quasi-simultaneous positioning technique employed in SnapLoc (see Chapter 5) with the multipath assistance utilized in SALMA (see Chapter 4). To this end, SnapLoc could provide coarse position estimates and instead of employing a timestamp correction method (see Sect. 5.2.3), the refinement can be done by SALMA with candidate points placed around the coarse position estimate and using the redundant multipath information of multiple anchors. The use of SnapLoc for the coarse position estimation allows to eliminate the need for a DS-TWR exchange in SALMA and, hence, to allow passive self-localization with multipath assistance. Combining this with a more efficient selection of candidate points for which the log-likelihood in Eq. 4.10 is evaluated, could further speed up the execution time of the algorithm used in SALMA. Indeed, estimating the likelihood of one candidate point in SALMA-full takes almost 1 ms in MATLAB [70]. Thus, reducing the number of candidate points for which the likelihood function is evaluated, increases the position update rate. This can be achieved by applying a gradient ascent algorithm to iteratively find the candidate point with the maximum likelihood. Alternatively, dividing the area of interest in grids and executing the likelihood function iteratively for one candidate point per grid likewise enables to eliminate the need for a DS-TWR exchange in SALMA. However, both options still require a floor plan, thus, in future we plan to design an efficient simultaneous localization and mapping (SLAM) algorithm employing multipath assistance [60, 121] and to employ low-cost radar systems to derive the floor plan automatically [64, 75].

Scalability in area. The focus of this doctoral thesis in terms of scalability is to support an unlimited number of tags without performance degradation and, hence, to cope with the highly increasing number of connected devices. However, the proposed concepts were solely evaluated in single room setups, thus, neglecting their scalability in terms of area, i.e., to support multi-room, multi-hop setups. To this end, similar to mobile communication standards such as GSM, we recommend to divide the area of operation in cells and assign individual channels and preamble codes to neighboring cells. In case the devices do not support multiple channels, we suggest to assign cells to individual time slots. Mobile tags can listen to multiple cells and select the one with the highest link quality. Periodically performing this quality check enables a handover between different cells and, to configure and synchronize the cells, we recommend to perform Glossy-like flooding [55]. Thus, we plan to extend SALMA and SnapLoc with multi-cell support and install and evaluate them in multi-room multi-level environments.

Improve accuracy and precision of positioning systems. In this thesis, we aim to demonstrate the raw capabilities of the presented techniques to improve robustness, efficiency, and scalability of UWB indoor positioning by utilizing solely single-shot measurements. Combining these concepts with other known techniques will improve the positioning performance even further. For instance, one can extend the presented positioning systems with tracking filters such as Kalman or particle filter to benefit from previous position estimates. Furthermore, fusing the positioning algorithms with low-cost inertial measurement units would significantly improve the accuracy and precision. This might also be achieved by enhancing the employed hardware. Indeed, improving the directionality of the UWB antenna system and a higher number of antennas results in a lower positioning error, as indicated in Sect. 4.6.2.2. Furthermore, we plan to evaluate the positioning system in a more dynamic environment using drones and robots.

Combine concepts with AoA measurements. As indicated in the related work section of Chapter 4, combining time-based positioning systems with AoA estimates is a common technique to reduce the required infrastructure. Hence, we plan to work on UWB-based AoA estimation and to strive after a multipath-assisted AoA positioning system using quasi-simultaneous responses.

IEEE 802.15.4f and IEEE 802.15.4z

A.1 LRP UWB PHY specification (based on IEEE 802.15.4f)

Besides the HRP UWB PHY described in Sect. 2.2.2.1, the IEEE 802.15.4-2015 standard encompasses a second physical layer definition related to ultra-wideband communications, namely, low-rate pulse repetition frequency (LRP) PHY [95]. The latter was originally published in 2012 as part of the IEEE 802.15.4f standard to support autonomous active RFID systems [94]. Similar to HRP, LRP UWB uses an impulse radio signal scheme, but there are significant differences in the band allocation and modulation. First, the standard specifies only three channels covering the frequency range from 6.2896 GHz to 9.1856 GHz. Second, depending on the mode, the modulation used is on-off keying (OOK), i.e., the presence/absence of pulses, or pulse position modulation (PPM), i.e., the information is encoded in the position of the pulses. Three modes are supported, which differ in terms of robustness and data rate. The highest data rate of 1 Mbps is provided in the *base mode*, where each symbol carries one bit of information and the PRF is 1 MHz. As suggested by the name, the significantly lower PRF depicts one of the main differences between LRP and HRP UWB. In the *extended mode* the data rate is reduced to 250 kbps and a 1/4 convolutional code is implemented to increase the robustness. However, the highest robustness and sensitivity is achieved in the *long-range mode* by employing 64 pulses per symbol and, in contrast to the other modes, this mode employs pulse position modulation instead of on-off keying. Similar to HRP, the synchronization header is divided into a preamble and a SFD consisting of a continuous stream of pulses and a fixed sequence of 16 bits, respectively. Optionally, all modes can append a location-enhancing information postamble (LEIP) to the transmitted packet consisting of a sequence of pulses. The purpose of the LEIP is to enhance ranging performance [18, 94]. In terms of security, proponents of the LRP UWB PHY claim that, in contrast to HRP, secure ranging with LRP UWB is provable [1, 122].

A.2 IEEE 802.15.4z draft standard

Although fundamental research on UWB was performed in the 1990s and early 2000s, it lasted more than a decade until the final breakthrough of UWB technology. To this end, one of the main driving applications is in the automotive domain to avoid relay attacks in car access systems [192]. The latter exploit the vulnerability of passive keyless entry systems, which allow the driver to enter the car and start the engine while keeping the keys in the pocket. One criminal stays close to the key fob, e.g., outside of the driver's home, and relays the signal from the key fob to another criminal positioned close to the car. This misleads the car into thinking that the driver and key are present, unlocks the car and allows to start the engine [15]. To combat this relay theft, UWB technology is utilized, as it allows to precisely estimate the distance of the vehicle and key fob. In the event of a relay attack, the car will note that the key fob is not in the vicinity and, hence, will not allow to open the car and start the engine [44]. Supporting afore-said use cases requires the adaptation of the UWB standards to improve the integrity, efficiency, and accuracy of existing range measurement methods as well as enhancing the MAC support for ranging procedures. For this purpose, a new task group was formed working on the IEEE 802.15.4z amendment [96]. At the time of writing, the standard is still in a draft status [91]. However, parts of the final standard are already becoming apparent such as the scrambled timestamp sequence (STS). The latter is a cryptographically generated sequence of pulses inserted into the PHY frame of HRP UWB devices. A receiver knowing the correct seed will then generate its own sequence and cross correlate it with the received one. This enables the estimation of a CIR used to detect malicious attacks and thus to increase the integrity of the message [171]. For LRP-based devices, authenticated ranging is achieved with a technique called distance commitment, which instructs the receiver to decode the data only during short periods related to the earliest detected path [96]. If an attacker would pretend a shorter distance by sending the preamble earlier, the receiver would detect that due to the random demodulated data [195].

Impact of clock variations in TWR

This appendix continues the discussion about two-way ToA ranging in Sect. 2.3.3 by analyzing the uncertainty of distance estimates in SS-TWR and DS-TWR due to clock variations.

B.1 Clock variations in SS-TWR

The individual clock drift of the initiator e_i and responder e_r , respectively, causes a ToF estimation error $\epsilon_{c,ss}$. The latter is defined in SS-TWR as following (if other error sources than clock drift are ignored) and with $\hat{\tau}_{f,ss}$ being the estimated time of flight according to Eq. 2.6 [167]:

$$\begin{aligned}\epsilon_{c,ss} = \hat{\tau}_{f,ss} - \tau_{f,ss} &= \frac{1}{2}(\hat{T}_{round} - \hat{T}_{reply}) - \frac{1}{2}(T_{round} - T_{reply}) \\ &= \frac{1}{2}[T_{round}(1 + e_i) - T_{reply}(1 + e_r) - T_{round} + T_{reply}] \quad (\text{B.1}) \\ &= \frac{1}{2}(e_i T_{round} - e_r T_{reply}).\end{aligned}$$

Replacing $T_{round} = 2\tau_{f,ss} + T_{reply}$ according to Eq. 2.6 yields

$$\begin{aligned}\epsilon_{c,ss} &= \frac{1}{2}(e_i \cdot (2 \cdot \tau_{f,ss} + T_{reply}) - e_r \cdot T_{reply}) \\ &= e_i \tau_{f,ss} + \frac{1}{2}[T_{reply}(e_i - e_r)].\end{aligned} \quad (\text{B.2})$$

Thus, the reply time T_{reply} is the dominant factor of the distance estimation error caused by clock drifts. T_{reply} is determined by the packet length as well as the processing speed of the responder and can be in the order of milliseconds, causing estimation errors of several meters [144].

B.2 Clock variations in DS-TWR

In double-sided two-way ranging, a second round-trip time measurement is employed, which affects the ranging uncertainty due to clock variations. Again, assuming that the frequency

clock drift is the dominant error source and given the definition of the true time of flight $\hat{\tau}_{f,ds}$ in Eq. 2.8, the estimated time of flight $\hat{\tau}_{f,ds}$ follows as [99]:

$$\begin{aligned}\hat{\tau}_{f,ds} &= \frac{1}{4}[(T_{round,i} - T_{reply,i})(1 + e_i) + (T_{round,r} - T_{reply,r})(1 + e_r)] \\ &= \tau_{f,ds} + \frac{1}{4}[(T_{round,i} - T_{reply,i})e_i + (T_{round,r} - T_{reply,r})e_r].\end{aligned}\quad (\text{B.3})$$

From Eq. 2.8 and Eq. B.3, the ToF estimation error $\epsilon_{c,ds}$ follows as:

$$\epsilon_{c,ds} = \hat{\tau}_{f,ds} - \tau_{f,ds} = \frac{1}{4}[(T_{round,i} - T_{reply,i})e_i + (T_{round,r} - T_{reply,r})e_r]. \quad (\text{B.4})$$

According to Eq. 2.7, $T_{round,i} = 2\tau_{f,ds} + T_{reply,r}$ and $T_{round,r} = 2\tau_{f,ds} + T_{reply,i}$, which yields

$$\epsilon_{c,ds} = \frac{1}{2}\tau_{f,ds}(e_i + e_r) + \frac{1}{4}[(T_{reply,r} - T_{reply,i})(e_i - e_r)]. \quad (\text{B.5})$$

Therefore, the estimation error due to clock drifts $\epsilon_{c,ds}$ is minimal when the reply times at initiator and responder, respectively, are equal. Hence, the impact of clock drifts on the ToF estimation error is significantly reduced when using an additional message compared to SS-TWR. As discussed in Sect. 2.3.3, Neiryneck et al. [144] proposed an alternative approach to estimate the time of flight $\tau_{f,ds}$ (see Eq. 2.10). In the approach named asymmetric DS-TWR, the estimated ToF $\hat{\tau}_{f,ds}$ under the assumption of clock drift as the main error source follows as:

$$\begin{aligned}\hat{\tau}_{f,ds} &= \frac{k_i T_{round,i} k_r T_{round,r} - k_i T_{reply,i} k_r T_{reply,r}}{2k_i(T_{round,i} + T_{reply,i})} \\ &= k_i \tau_{f,ds},\end{aligned}\quad (\text{B.6})$$

where $k_i = (1 + e_i)$ and $k_r = (1 + e_r)$. Thus, the resulting estimation error $\epsilon_{c,ds} = \hat{\tau}_{f,ds} - \tau_{f,ds} = k_i \tau_{f,ds} - \tau_{f,ds} = e_i \tau_{f,ds}$ solely depends on the ToF and not on the reply times.

Part II

Included Publications

This doctoral thesis is mainly based on the publications listed in the following table. The right column indicates the chapter, where contents of the according publication are discussed.

Publication	Chapter
<p>Paper A: Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications. <u>B. Großwindhager</u>, C.A. Boano, M. Rath, and K. Römer. 19th IEEE Symp. on a World of Wireless, Mobile, and Multimedia Netw. (WoWMoM). June 2018.</p>	Ch. 3
<p>Paper B: SALMA: UWB-based Single-Anchor Localization System Using Multipath Assistance. <u>B. Großwindhager</u>, M. Rath, J. Kulmer, M.S. Bakr, C.A. Boano, K. Witrissal, and K. Römer. 16th ACM Conference on Embedded Networked Sensor Systems (SenSys). November 2018.</p>	Ch. 4
<p>Paper C: Concurrent Ranging with Ultra-Wideband Radios: From Experimental Evidence to a Practical Solution. <u>B. Großwindhager</u>, C.A. Boano, M. Rath, and K. Römer. 38th IEEE International Conference on Distributed Computing System (ICDCS). July 2018.</p>	Ch. 5
<p>Paper D: SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags. <u>B. Großwindhager</u>, M. Stocker, M. Rath, C.A. Boano, and K. Römer. 18th ACM/IEEE Intern. Conf. on Information Processing in Sensor Networks (IPSN). April 2019. <i>Best Paper Award.</i></p>	Ch. 5

In addition, the following table lists publications, which are not included in this doctoral thesis.

Publication
<p>Switchable Directional Antenna System for UWB-based Internet of Things Applications. B. Großwindhager, M.S. Bakr, M. Rath, F. Gentili, W. Bösch, K. Witrisal, C.A. Boano, and K. Römer. International Conference on Embedded Wireless Systems and Networks (EWSN). Feb. 2017. <i>Best Poster Award.</i></p>
<p>Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning. J. Kulmer, S. Hinteregger, B. Großwindhager, M. Rath, M.S. Bakr, E. Leitinger, and K. Witrisal. IEEE International Conference on Communications Workshops (ICC Workshops). May 2017.</p>
<p>Multipath-assisted Indoor Positioning Enabled by Directional UWB Sector Antennas. M. Rath, J. Kulmer, M.S. Bakr, B. Großwindhager, and K. Witrisal. 18th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC). July 2017.</p>
<p>Demo: UWB-based Single-anchor Low-cost Indoor Localization System. B. Großwindhager, M. Rath, J. Kulmer, S. Hinteregger, M.S. Bakr, C.A. Boano, K. Witrisal, and K. Römer. 15th ACM Conference on Embedded Network Sensor Systems (SenSys). November 2017. <i>Best Demo Award.</i></p>
<p>Dependable Internet of Things for Networked Cars. B. Großwindhager, A. Rupp, M. Tappler, M. Tranninger, S. Weiser, B. Aichernig, C.A. Boano, M. Horn, G. Kubin, S. Mangard, M. Steinberger, and K. Römer. International Journal of Computing. 16(4):226-237. December 2017.</p>
<p>Poster: Runtime Adaptation of PHY Settings for Dependable UWB Communications. B. Großwindhager, C.A. Boano, M. Rath, and K. Römer. 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN). April 2018.</p>
<p>Dataset: Single-anchor Indoor Localization with Decawave DW1000 and Directional Antennas. B. Großwindhager, M. Rath, J. Kulmer, M.S. Bakr, C.A. Boano, K. Witrisal, and K. Römer. 1st Workshop on Data Acquisition To Analysis (DATA). November 2018.</p>
<p>Dependable Wireless Communication and Localization in the Internet of Things. B. Großwindhager, M. Rath, M.S. Bakr, P. Greiner, C.A. Boano, K. Witrisal, F. Gentili, J. Grosinger, W. Bösch, and K. Römer. The Philosophy of Mission-Oriented Sensor Networks and Systems - Volume 2: Advances. 2019</p>
<p>A Compact Broadband Frequency Selective Microstrip Antenna and its Application to Indoor Positioning Systems for Wireless Networks. M. S. Bakr, B. Großwindhager, M. Rath, J. Kulmer, I. C. Hunter, R. A. Abd-Alhameed, K. Witrisal, C. A. Boano, K. Römer, and W. Bösch. IET Microwaves, Antennas and Propagation. March 2019.</p>
<p>Demo: SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags. M. Stocker, B. Großwindhager, C.A. Boano, and K. Römer. 18th International Conference on Information Processing in Sensor Networks (IPSN). April 2019.</p>

Paper A

B. Großwindhager, C.A. Boano, M. Rath, and K. Römer. **Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications.** In *Proceedings of the 19th IEEE International Symposium on a World of Wireless, Mobile, and Multimedia Networks (WoWMoM'18)*, pages 01–11, Crete, Greece. June 2018.

©2018 IEEE

ISBN: 978-1-5386-4725-7

DOI: 10.1109/WoWMoM.2018.8449776

Link: <https://ieeexplore.ieee.org/document/8449776>

Abstract. Ultra-wideband (UWB) technology is increasingly used to build location-aware IoT applications because of its outstanding positioning accuracy. Its communication performance, however, is unexplored and strongly affected by the chosen physical layer settings as well as by the surrounding environment. Finding an effective way to increase the dependability of UWB communications is yet an open problem. In this paper, we study the performance of different UWB physical layer settings and use them as tuning knobs to increase the energy efficiency and robustness of communications. Towards this goal, we first experimentally quantify the reliability and energy cost of each setting, in order to understand which physical layer configuration to privilege depending on the application requirements. We then use the estimated channel impulse response – a unique feature of UWB transceivers – to accurately measure the link quality and to extract relevant information about the characteristics of the surrounding environment, such as the presence of destructive interference. Capitalizing on this information, we design a scheme that adapts the UWB physical layer settings at runtime. An experimental evaluation using the Decawave DW1000 radio shows the effectiveness of the proposed adaptive scheme, highlighting the increased communication robustness and energy efficiency.

My contribution. I am the main author of this publication and developed the ideas to exploit the channel impulse response to optimize the robustness and energy efficiency of UWB communication at runtime. I have implemented all the software and algorithms and, furthermore, I have carried out all the experiments including the PHY parameter characterization and the evaluation. Carlo Boano helped me significantly in writing and structuring the paper and with Michael Rath I had several fruitful technical discussions. I presented the paper at WoWMoM'18.

Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications

Bernhard Großwindhager, Carlo Alberto Boano, Michael Rath, and Kay Römer
Faculty of Electrical and Information Engineering, Graz University of Technology, Austria
E-mail: {grosswindhager, cboano, mrath, roemer}@tugraz.at

Abstract—Ultra-wideband (UWB) technology is increasingly used to build location-aware IoT applications because of its outstanding positioning accuracy. Its communication performance, however, is unexplored and strongly affected by the chosen physical layer settings as well as by the surrounding environment. Finding an effective way to increase the dependability of UWB communications is yet an open problem. In this paper, we study the performance of different UWB physical layer settings and use them as tuning knobs to increase the energy efficiency and robustness of communications. Towards this goal, we first experimentally quantify the reliability and energy cost of each setting, in order to understand which physical layer configuration to privilege depending on the application requirements. We then use the estimated channel impulse response – a unique feature of UWB transceivers – to accurately measure the link quality and to extract relevant information about the characteristics of the surrounding environment, such as the presence of destructive interference. Capitalizing on this information, we design a scheme that adapts the UWB physical layer settings at runtime. An experimental evaluation using the Decawave DW1000 radio shows the effectiveness of the proposed adaptive scheme, highlighting the increased communication robustness and energy efficiency.

I. INTRODUCTION

Ultra-wideband (UWB) radio transceivers spread the signal power over a much wider bandwidth than traditional narrowband IoT technologies. This results in beneficial properties such as a high immunity to multipath fading [1], [2], a high data throughput [3], as well as a very good time-domain resolution allowing for accurate localization and tracking [4], [5].

Especially the possibility to achieve a localization accuracy at the centimeter level in scenarios with limited global navigation satellite system reception has attracted the attention of the research community [3], [6]. A large body of work has indeed produced complex algorithms to maximize the ranging accuracy indoors [7] for example, by using channel information and exploiting multipath reflections [8], or by combining inertial sensors and time of arrival measurements [9].

UWB communication performance still unexplored. The community, however, has not yet investigated in depth the communication performance of UWB, nor tried to build robust and highly-available IoT applications using UWB. Studying and maximizing the energy efficiency and reliability of UWB communications is still a necessity and a long-due step for two main reasons. First, to fully exploit the outstanding positioning accuracy of UWB and build location-aware applications such as assisted living [10], medical monitoring [11], and drone or robot navigation [12], [13], *robust communication links*

are necessary in order to reliably acquire and share the timestamps needed to carry out time-based location estimation. Second, UWB radios are more energy-hungry than common narrowband IoT transceivers such as BLE [14]. As a result, UWB-based localization systems are typically more efficient when using a *separate transceiver* for communication [15].

The role of UWB physical settings. Ultra-wideband is fundamentally different in nature from other IoT technologies. Besides the use of channels with a much higher bandwidth, in UWB one can configure several physical layer (PHY) settings that drastically affect the radio sensitivity and power consumption. This is different from narrowband IEEE 802.15.4 systems, where most of the knobs to tune the energy efficiency and robustness of communications are at the MAC layer, e.g., duty cycle [16], clear channel assessment threshold [17], and backoff times [18]. Moreover, whilst in common narrowband transceivers one can tune the transmission power to reduce energy consumption or to increase communication range, this is not easily possible with UWB radios, due to the restrictive transmit power regulations. Unfortunately, the research community has not yet characterized how PHY settings affect the performance of UWB communications – as opposed to the large body of work studying the role of PHY settings for narrowband IoT technologies such as LoRa [19], [20].

The need for runtime adaptation. Shedding light on the selection of PHY settings in low-cost UWB radios is hence a first necessary step to expose tuning knobs that can be used to control and to ultimately increase the energy efficiency and robustness of UWB communications. Such tuning knobs would be especially useful to dynamically react to the fluctuations in communication performance triggered by the mobility of UWB nodes. To date UWB systems make use of *static* PHY settings, i.e., hardcoded physical layer parameters that remain constant over time, which makes them incapable to sustain an enduring high packet delivery rate [13] and unable to cope with changes in the surrounding environment.

The role of the environment. Even though UWB has a high immunity to multipath fading and is more robust to interference than most narrowband technologies [3], its performance is still strongly affected by varying environmental conditions, such as the existence of obstacles limiting the line-of-sight and the presence of destructive interference. A detailed understanding of the characteristics of the surrounding environment and the ability to assess the quality of a wireless link, would be valuable to maximize the reliability of UWB transmissions.

Estimating the link quality. The body of work on link quality estimation for UWB-based systems is, however, quite limited, and the information that can be retrieved from UWB radios is significantly different from the one that has been largely studied in narrowband IoT technologies [21]. Among others, conventional energy detection methods are not available on UWB transceivers due to the low power spectral density [22]. Furthermore, how to make use of physical layer information that is specific to UWB transceivers (such as the *estimated channel impulse response*) to estimate the link quality or to characterize the surrounding environment has not yet been thoroughly investigated by the research community.

Contributions. In this paper we first carry out a comprehensive study of the performance of low-cost UWB transceivers as a function of different PHY settings. Our experimental evaluation reveals that by tuning, among others, pulse repetition frequency, data rate, bandwidth, and carrier frequency, one can increase the radio sensitivity by up to 8 dB – a value sufficient to transform a useless link into one able to sustain a high delivery rate. We then quantify the energy consumption of different PHY configurations and highlight which change of settings to privilege in order to obtain an increase in communication robustness at minimal energy costs.

Second, we analyze the PHY information that can be derived from off-the-shelf UWB transceivers and study how it can be used to (i) accurately estimate the link quality, as well as to (ii) extract information about the characteristics of the surrounding environment, in order to guide a better selection of PHY settings. In this regard, our investigation has shown that the estimation of channel impulse response (CIR) measured by UWB radios can be used, among others, to detect the presence of destructive interference. We use this information, combined with the received signal power derived from the estimated CIR, to build a robust link state indicator for UWB systems.

Third, we show that such link state indication can be obtained even in *absence of a complete packet reception*. The estimated CIR is indeed extracted only from the preamble of a UWB packet, and we experimentally show that a preamble can still be fully decoded even when the SNR is 12 dB lower than the one necessary to decode the payload. This can be explained with the different modulation scheme used to encode preamble and payload – another key feature of UWB radios.

Finally, we make use of all aforementioned results to design an adaptive scheme that derives an optimal set of PHY settings at runtime based on the estimated link quality, the characteristics of the surrounding environment, and the application requirements. Among others, our adaptation logic exploits the fact that some of the changes in UWB PHY settings *do not require prior agreement* among sender and receiver, and makes use of acknowledgement (ACK) messages consisting of just a preamble, hence exploiting the more robust modulation scheme of the preamble to eliminate the problem of *asymmetric links*.

An experimental evaluation using the Decawave DW1000 radio shows the increased communication robustness and energy efficiency when using our adaptation scheme.

In summary, this paper makes the following contributions:

- We characterize the performance of UWB transceivers as a function of different physical layer settings and experimentally quantify their impact on the reliability and energy efficiency of communications (Sect. III);
- We use physical layer information to estimate the link quality and to characterize the surrounding environment also in absence of a complete packet reception (Sect. IV);
- We design a scheme to adapt the PHY settings of UWB radios at runtime (Sect. V), and evaluate experimentally the increased robustness and energy efficiency (Sect. VI).

II. DEMISTIFYING UWB TECHNOLOGY

Ultra-wideband radios spread the signal power over a much wider bandwidth (≥ 500 MHz) than traditional narrowband IoT technologies such as BLE and LoRa, yielding an extremely low power spectral density and, as a consequence, reducing interference to other wireless systems.

Besides the higher bandwidth, the UWB PHY for low-rate wireless personal area networks formalized in 2007 by the IEEE 802.15.4 working group is fundamentally different from its narrowband counterpart. We describe next the main differences, namely: the employed modulation schemes (Sect. II-A), the configurable physical settings (Sect. II-B), the international regulations on the radiated power (Sect. II-C), as well as the information that can be extracted from the channel (Sect. II-D). We finally show that, despite its robustness and higher immunity to multipath fading, UWB communications may experience overlapping multipath components and suffer from destructive interference (Sect. II-E).

A. Different Modulation Schemes

Devices communicating using the IEEE 802.15.4 UWB standard use two different modulation schemes. A UWB frame consists indeed of two main blocks: a synchronization header (SHR), and a data portion consisting of a physical layer header (PHR) and a payload, as illustrated in Fig. 1. The SHR is sent using single pulse modulation, i.e., it consists of single pulses of fixed duration $t_p = 1/f_p$, where f_p corresponds to the highest frequency at which standard-compliant UWB transceivers are allowed to transmit pulses (499.2 MHz). The data portion is sent as bursts using BPM/BPSK modulation, i.e., the information is encoded in the position of the burst (burst position modulation or BPM) and in the phase of the burst (binary-phase shift keying or BPSK). We will show in Sect. III-B how the SHR and the data portion exhibit a different robustness due to the diverse modulation schemes.

B. Configurable PHY Settings

Synchronization header settings. The SHR is composed of a preamble (used for signal detection as well as frame synchronization) and a start-of-frame delimiter (SFD). The SFD marks the end of the preamble and indicates that the receiver has to switch to BPM/BPSK modulation to receive the data portion. The length of the SFD depends on the data rate of the payload: it contains 64 preamble symbols if the latter is sent at 110 kbps (the lowest data rate available), whilst it contains 8 preamble symbols for any other data rate used.




Fig. 1. UWB PHY frame structure according to the IEEE 802.15.4 standard. Settings marked with * are also available in the Decawave DW1000 transceiver, whilst the ones marked with || are not supported [23].

The duration of an SHR depends on two tunable settings: the *pulse repetition frequency (PRF)* and the number of *preamble symbol repetitions (PSR)*. One preamble symbol consists of N time intervals in which either a positive, negative, or no pulse is sent. N depends on the PRF, which defines the number of transmitted pulses within a certain period: N is 496 or 508 for a PRF of 16 and 64 MHz, respectively. The resulting duration of a preamble symbol is $t_{ps} = N/f_p$. The number of preamble symbols sent is defined by the PSR, which directly affects the duration of a frame (due to the higher preamble duration t_{psr}) and hence the radio's energy expenditure (see Sect. III-C).

Data portion settings. The robustness with which a data portion can be received depends, among others, on the *data rate*. Lowering the latter increases the likelihood to sustain a more reliable link, as we show experimentally in Sect. III-C.

The physical header is 19 bits long and contains the length of the payload and the data rate used to transmit it. It is sent at 110 kbps if the payload is also sent at 110 kbps (the lowest data rate available), whilst it is sent at 850 kbps for any other data rate used to transmit the payload. This implies that a transmitter can change the data rate of the payload (for example from 850 kbps to 6.8 Mbps) without the necessity of a prior agreement with a receiver node. We will use this feature in our runtime adaptation scheme described in Sect. V.

C. UWB Frequencies and Emission Limits

The IEEE 802.15.4 standard dictates which frequencies should be used for communications and specifies that the maximum transmit power shall conform with local regulations.

Diversification of channels. The IEEE 802.15.4 UWB PHY allocates frequencies in three ranges: below 1 GHz, between 3 and 5 GHz, and between 6 and 10 GHz. Within these three bands, the standard suggests 16 different channels for UWB, out of which twelve support a bandwidth of 500 MHz and four make use of an increased bandwidth of up to 1331.2 MHz. Channels with a higher bandwidth can be used to increase the robustness of wireless links, as we will illustrate in Sect. III.




Fig. 2. Estimated channel impulse response obtained from a Decawave DW1000 in an indoor environment. The LOS component is marked with a dashed green line and significant multipath components with dotted red lines.

Restrictive power regulations. The transmission power of UWB communications is severely constrained by the Federal Communications Commission (FCC) and International Telecommunication Union Radiocommunication Sector regulations¹. As a result, the bandwidth and the mean power of a UWB transmitter need to be individually calibrated and tuned so that the spectrum fits the regulatory mask. This requires expensive measurement tools and time-consuming calibration. Furthermore, tuning these settings affects important ranging parameters such as the transmit and receive antenna delays [26], [27]. This hinders the adjustment of transmission power at runtime as often carried out with narrowband IEEE 802.15.4 devices [28].

D. Available Channel Information

Narrowband radios typically provide hardware indicators (capturing the energy of the received signal or the chip error rate) that are often used to estimate the quality of a wireless link [21]. UWB transceivers additionally provide a *channel impulse response (CIR)* estimation, i.e., information about the multipath propagation consisting of reflections from walls and scattering from other objects. Fig. 2 depicts an estimated CIR obtained from a Decawave DW1000 radio. It shows the line-of-sight (LOS) component (marked with dashed green line) and significant multipath reflections (marked with dotted red lines). Due to the high bandwidth (and hence the very high time resolution of UWB), one can extract these multipath reflections: in narrowband radios, the latter would overlap with the LOS component, leading to severe multipath fading.

CIR information is used in UWB radios to precisely estimate the arrival time of a packet by detecting the first path or leading edge of the CIR. One can also make use of the multipath information in the CIR for high-accuracy indoor positioning [8]. We propose, instead, to use this CIR information to *characterize the surrounding environment*, e.g., to detect the presence of destructive interference, and guide an optimal selection of UWB PHY settings, so to increase the reliability of communications, as described in Sect. IV and V.

E. Susceptibility to the Environment

As for every wireless technology, also the performance of Ultra-wideband is strongly affected by the surrounding environment. Communication performance may indeed vary when moving from multipath-rich indoor settings to outdoor areas, or in the presence of obstacles limiting the LOS.

¹The maximal equivalent isotropically radiated power in any direction should not exceed -41.3 dBm/MHz [24], [25].




Fig. 3. Decrease in packet reception rate (PRR) caused by destructive interference in a multipath-rich indoor environment.

What is not obvious, is that one can perceive destructive interference also in UWB links, despite the high immunity to multipath fading. In fact, since the bandwidth is not infinite, one can still suffer from overlapping multipath components, thus there is fading also in UWB [29]. We have measured this situation experimentally by using two UWB nodes embedding a Decawave DW1000 radio configured with its default settings (see Table I). The nodes were mounted at 1m height and exchanged 1000 packets at specific distances in a hallway about 60 meters long. Fig. 3 shows the packet reception rate (PRR) at different distances. The fade at a distance of about 20 meters is clearly visible: the link degrades and the PRR drops close to zero. The link sustains again a perfect reception rate at 28 meters and beyond. In Sect. IV we show that this fade is due to destructive interference, and that the cause of packet loss can be identified using the CIR. We will integrate this environment classification technique into our adaptive scheme to increase the reliability of communications (Sect. V).

III. CHARACTERIZING UWB PERFORMANCE

We study next the impact of the IEEE 802.15.4 UWB modulation schemes and the different PHY settings on the robustness and energy efficiency of communications. After describing our experimental setup in Sect. III-A, we show that the detection of a preamble is significantly more robust than decoding a payload (Sect. III-B). We then quantify the impact of each PHY setting on the reliability and efficiency of UWB communications (Sect. III-C) and on the precision of ranging (Sect. III-D). Based upon these findings, we shed light on the change of PHY settings that should be privileged to increase UWB communication performance in Sect. III-E.

A. Experimental Setup

To characterize the performance of UWB, we make use of the Decawave DW1000 UWB radio [23]. In particular, we employ EVB1000 boards consisting of a DW1000 transceiver, an ARM Cortex M3 processor, and an omnidirectional antenna.

Cable connection. In order to maximize reproducibility and avoid environmental influences such as weather, temperature, or obstacles, we connect two EVB1000 boards using SMA cables and place a Mini-Circuits RCDAT-8000-90 programmable attenuator in between. The latter allows us to control the level

TABLE I
DEFAULT CONFIGURATION OF THE DW1000 RADIO ON POWER-UP.

PHY Setting	Value
Channel	5
Pulse repetition frequency	16 MHz
Preamble symbol repetitions	128
Data rate	6.8 Mbps
Payload size (including MAC header)	12 Bytes

of attenuation in 0.25 dB steps and hence to finely explore the differences in communication robustness for each setting.

Two-way ranging. For each attenuation step, we let the two EVB1000 boards carry out 1000 ranging operations. The used two-way ranging scheme consists of three messages to estimate the distance without the need of a tight synchronization [30]. The exchanged packets consist of an 11-byte MAC header that embeds source and destination addresses, as well as a 16-bit checksum. The payload of the first two messages is 1-byte long, whilst the last message includes three timestamps and is 16-bytes long. For each message, both sender and receiver record if the preamble was detected and the payload received, as well as status information provided by the radio.

Testing settings individually. We let the two UWB nodes communicate using the DW1000 default configuration on power-up (see Table I) and vary each PHY setting individually, so to clearly quantify the impact of each setting on communication performance. We repeat our measurements using six different boards to make sure that our results are not hardware-specific.

Measuring reliability. For each 0.25 dB step of the programmable attenuator, we compute the *packet reception rate* (PRR) as the ratio between the number of packets for which both SHR and payload are correctly received, and the number of packets sent. We also compute the *header reception rate* (HRR) as the ratio between the number of packets for which a valid SHR is received, and the number of packets sent.

Measuring energy efficiency. To quantify the energy efficiency of each setting, we measure the average current consumption of each transmission and reception using a Keysight MSO-S 254A oscilloscope, distinguishing between the SHR and the data portion. We then derive the energy consumption using the packet's over-the-air time and the supply voltage (3.3 V).

B. Impact of different Modulation Schemes

Using the setup described in Sect. III-A, we experimentally investigate the robustness of the different modulation schemes used in the SHR and the data portion. Our results show that an SHR can still be successfully detected even at an attenuation 7 dB higher than the one at which the first packets are lost. Fig. 4 shows the PRR and HRR over the attenuation when using the default configuration (see Table I). To eliminate measurement noise, the curves were smoothed with a five-sample moving average filter. One can clearly identify in Fig. 4 the three distinct reception regions: connected, transitional, and disconnected [31]. At an attenuation of approximately 70 dB, when the radio can no longer decode the data portion, the SHR is still successfully detected with a probability of 100%. The higher robustness of the SHR is due to the different modulation scheme and leads to two important observations.

CIR info available in absence of a complete reception. Even when a receiver cannot decode the data portion of a packet, it may still correctly decode its SHR and extract its information. Among others, one can extract the estimated CIR and use it to characterize the surrounding environment and to anticipate the degradation of a wireless link by detecting the presence of destructive interference, as described in Sect. IV.




Fig. 4. Packet reception rate (PRR) and header reception rate (HRR) for different attenuation levels when using the default configuration of the DW1000 radio on power-up. The SHR can be detected even at an attenuation 7 dB higher than the one at which the first packets are lost.

Preamble as a binary ACK. Several systems make use of acknowledgement (ACK) messages to confirm the reception of packets and agree upon specific information, e.g., a new time slot or frequency channel. These ACK messages should be received reliably, so to avoid disagreements between nodes [32]. When using UWB radios, one could make use of the higher robustness of the SHR (and hence of the preamble) to increase the reliability of these ACK messages. The reception of a preamble can, for example, be used as a reliable binary ACK to signal the correct reception of a full packet. The probability that a data portion was received in one direction, and that the preamble (i.e., an ACK) was lost in the other direction is, indeed, low, given the results presented in Fig. 4.




Fig. 5. Packet reception rate (PRR) and header reception rate (HRR) for different attenuation levels as a function of preamble symbol repetitions (PSR).




Fig. 6. Packet reception rate (PRR) and header reception rate (HRR) for different attenuation levels as a function of pulse repetition frequency (PRF).

TABLE II
DECAWAVE DW1000 SUPPORTED CHANNELS.

Ch.	Carrier Freq. [MHz]	Bandwidth [MHz]	TX power [dBm]
1	3494.4	499.2	-14.32
2	3993.6	499.2	-14.32
3	4492.8	499.2	-14.32
4	3993.6	1331.2 (900)	-11.76
5	6489.6	499.2	-14.32
7	6489.6	1081.6 (900)	-11.76

C. Impact of different PHY Settings

We characterize next the performance of different UWB PHY settings using the setup described in Sect. III-A.

Preamble symbol repetitions (PSR). The IEEE 802.15.4 UWB standard defines four possible PSR: 16, 64, 1024, and 4096. The DW1000 does not allow selecting PSR of 16, but allows to choose other values (see Fig. 1). We hence run experiments using PSR of 64, 128, 1024, and 4096. Fig. 5 shows the PRR and HRR over the different steps of the programmable attenuator for different PSR values. First, we can observe that PSR does not affect the PRR: this is to be expected, as the PSR tunes the SHR and not the data portion. Second, longer preambles result in a correct SHR reception also at much lower signal-to-noise ratios: increasing the PSR from 64 to 4096 allows to lower the sensitivity level of the SHR by 6 dB. Third, by tuning the PSR, one can further increase the margin between the successful detection of the SHR and the correct decoding of the data portion up to 12 dB. Table III provides a comparison of the energy required to send packets using different PSR at the transmitter and at the receiver. Using the longest preamble (PSR = 4096) results in an energy consumption 5, 30, and 57 times higher compared to the one measured with PSR of 1024, 128, and 64, respectively. Note that Table III just reflects the energy consumption when actively transmitting and receiving a packet: listening for a preamble also accounts for a current draw of about 130 mA.

Pulse repetition frequency (PRF). The PRF determines the number of pulses sent within a certain period and influences both SHR and data portion of a packet. Although the standard specifies possible PRF values of 4, 16, and 64 MHz, the former is not supported by the DW1000, and we hence only compare the use of a PRF of 16 and 64 MHz. A single preamble symbol is slightly shorter when using a lower PRF. In addition to the lower amount of pulses sent, this also results in an energy consumption up to 24% lower when transmitting a preamble. Since also less pulses are sent within a burst, the use of a smaller PRF results in an energy consumption 6% lower also when transmitting a data portion (see Table III). The higher amount of pulses per burst improves the robustness of the data portion and increases the reliability of communications. Fig. 6 shows that an increase in PRF from 16 to 64 MHz gives up to 1 dB gain when the PRR is 90% (red-dotted line). We can also observe that the benefit of the PRF tuning depends on the selected carrier frequency, and that an increase in PRF results in a stronger reception of an SHR (i.e., a higher HRR).

TABLE III
ENERGY CONSUMPTION PER TRANSMITTED AND RECEIVED PACKET DEPENDING ON THE PHY SETTINGS.

PHY Setting	Transmission				Reception				Transmission + Reception	
	SHR [μ J]	Δ [%]	Payload [μ J]	Δ [%]	SHR [μ J]	Δ [%]	Payload [μ J]	Δ [%]	Total [μ J]	Δ [%]
Default	31.23	-	6.198	-	66.34	-	18.03	-	121.8	-
PSR = 4096	942.5	2918	6.198	-	2001.9	2917.6	18.03	-	2969	2337
PSR = 1024	237.0	658.8	6.198	-	503.4	658.8	18.03	-	764.6	527.8
PSR = 64	16.54	-47.06	6.198	-	35.12	-47.06	18.03	-	75.9	-37.7
Ch. 1 (500 MHz)	28.39	-9.118	5.177	-16.48	59.91	-9.686	16.03	-11.10	109.5	-10.10
Ch. 2 (500 MHz)	28.96	-7.278	5.329	-14.01	63.35	-4.50	17.09	-5.201	114.7	-5.80
Ch. 3 (500 MHz)	28.37	-9.163	5.323	-14.12	62.79	-5.353	17.16	-4.848	113.6	-6.701
Ch. 4 (900 MHz)	33.35	6.78	5.710	-7.859	62.16	-6.30	17.04	-5.464	118.3	-2.901
Ch. 7 (900 MHz)	30.83	-1.30	6.233	0.563	67.72	2.084	19.06	5.711	123.8	1.676
PRF = 64 MHz	38.65	23.75	6.567	5.96	69.11	4.179	18.03	-	132.3	8.648
DR = 850 kbps	31.45	0.72	22.71	266.4	67.78	2.173	66.29	267.7	188.2	54.55
DR = 110 kbps	44.46	42.33	176.2	2742	92.48	39.41	554.7	2977	867.8	612.5

Data rate. Although the IEEE 802.15.4 UWB standard supports data rates up to 27 Mbps, the DW1000 operations are limited to 6.8 Mbps. We hence investigate the use of a data rate of 110 kbps, 850 kbps, and 6.8 Mbps. Table III shows that transmitting and receiving a header at 110 kbps consumes about 39-42% more energy because of the longer SFD (the preamble itself is independent from the data rate). The energy spent to transmit and receive a payload, instead, is severely more affected by the employed data rate and increases linearly with the payload size. Transmitting and receiving the payload of 12 bytes with a data rate of 850 and 110 kbps consumes respectively up to 3.7 and 31 times more energy than when using the default settings. Fig. 7 (top) shows how increasing the data rate increases the robustness of the data portion. Using a data rate of 850 and 110 kbps raises the link margin by 2.9 and 5.5 dB, respectively. The HRR, instead, is independent from the data rate, as shown in Fig. 7 (bottom).




Fig. 7. Packet reception rate (PRR) and header reception rate (HRR) for different attenuation levels as a function of data rate.




Fig. 8. Packet reception rate (PRR) and header reception rate (HRR) for different attenuation levels as a function of carrier frequency and bandwidth.

Channel and bandwidth. The IEEE 802.15.4 standard defines 16 different channels for UWB, out of which the DW1000 transceiver supports six (see Table II). Fig. 8 shows that using a lower carrier frequency (e.g., channel 1) increases the robustness of both SHR and data portion. Similarly, the use of channels with the same carrier frequency, but a higher bandwidth (channel 4 and 7, see Table II) results in a more reliable communication. In terms of energy consumption, the channels below 5 GHz are slightly more efficient, as shown in Table III, whilst the higher bandwidth comes at the cost of a slightly higher energy expenditure.

Please note that the performance of different channels highly depends on the transceiver calibration parameters, as discussed in Sect. II-C. Furthermore, as our setup is cable-based, the Friis equation and more realistic UWB channel models [1], [33] have to be considered on top of our observations. In fact, the Friis equation strengthens our measurements even more as the free space path loss is reduced at lower carrier frequencies.

D. Impact of PHY Settings on Ranging Precision

Sect. III-C has shown that tuning the PHY settings of UWB radios can increase or decrease the chances of successfully receiving a packet. Whilst changes in PHY settings can maximize the probability to acquire and share the timestamps (a necessary step to estimate the distance between nodes), they may also affect ranging performance. Indeed, as described in Sect. II-D, the ranging performance depends on the signal-to-noise ratio (SNR) of the CIR estimation. The latter, for example, is improved by a longer preamble and a higher PRF. We hence investigated the impact of different PHY settings on the precision of ranging, and observed that such impact is minimal on low-cost transceivers such as the DW1000 (as also confirmed by [8]). Although no quantitative comparison is possible, we determined an upper bound on the ranging error for all possible settings to make sure that adapting the PHY settings at runtime does not affect the ranging performance.




Fig. 9. Cumulative distribution function (CDF) of the ranging error. Solid blue line for reliable links and green dotted line for unreliable links.

TABLE IV
TRADEOFF BETWEEN PACKET/PREAMBLE RELIABILITY AND ENERGY EFFICIENCY.

Change in PHY setting	Reliability of data portion	Reliability of SHR	Energy efficiency	Suggestions
Higher PSR	-	↑↑↑↑	↓↓↓↓	Use shortest PSR possible, except in cases it should be ensured that the preamble is reliably detected.
Higher PRF	↑	↑	↓↓	Limited impact on robustness, prefer a lower PRF (16 MHz) for highest energy efficiency
Lower data rate	↑↑↑↑	-	↓↓↓↓	Try to use the highest data rate sustaining a reliable link. Lower the data rate if link degrades heavily.
Lower carrier frequency	↑↑↑	↑↑↑	↑	Use lowest channel available to increase robustness at minimal energy costs.
Higher bandwidth	↑↑	↑↑	↓	Use channel with a higher bandwidth, except when energy is by far the most critical resource.

We use the setup described in Sect. III-A, and perform 11,000 two-way ranging trials for each of the PHY configurations investigated in Sect. III-C. For each trial, we compute the distance and then calculate the cumulative distribution function of the ranging error for all PHY settings which is depicted in Fig. 9. The precision of the range estimate lies within a bound of ± 10 cm for 99.7% of all range estimations. Fig. 9 further shows the influence of a highly reliable ($\text{PRR} \geq 90\%$) and a less reliable ($\text{PRR} < 90\%$) link on the precision: the error for 90% of the measurements is below 3.77 and 4.22 cm, respectively. Hence, increasing the reliability of UWB communications even slightly improves ranging performance.

E. Optimal Selection of PHY Settings

The experiments described in Sect. III-C allow us to assess which PHY settings are the most effective in increasing the robustness of communications, and to quantify the energy expenditure they incur. Table IV summarizes the influence of each PHY setting on the reliability of SHR and payload, as well as the energy cost they entail. One can derive from this table which change of settings to privilege in order to obtain an increase in communication robustness depending on the application requirements. If one aims to quickly increase the PRR without any constraint on the energy expenditure, the first option should be to decrease the data rate, as the latter has, by far, the highest impact on the reliability of a link. On the contrary, if one aims to increase the PRR with the least energy costs, one should first decrease the carrier frequency and switch to a channel with higher bandwidth, then try to increase the PRF, and, only as a last resort, lower the data rate. We will make use of these observations in Sect. V to design an adaptive scheme that derives at runtime the most energy-efficient PHY settings sustaining a reliable communication.

IV. ESTIMATING THE LINK STATE FROM THE CIR

To increase the reliability of communications by adapting PHY settings at runtime, one needs the ability to accurately assess the quality (or predict the degradation) of a UWB link. Whilst several solutions have investigated how to estimate the link quality using narrowband technologies [21], the link quality information that can be retrieved from UWB radios is still under-explored. Furthermore, as UWB communication performance is also affected by the surrounding environment, it is desirable to obtain valuable information about its characteristics and to make use of the latter to react accordingly.

In this section we present a novel *link state* indicator that makes use of CIR information to (i) estimate the *link quality*, and to (ii) obtain a detailed understanding of the *environmental state* in order to detect the cause for a degrading channel. The principle of our link state indicator is shown in Fig. 12: the link quality is obtained from the received signal power (RSP) derived from the CIR estimation. Binary environmental state information such as the presence of destructive interference or a clear LOS is also derived from the CIR estimation.

A. Deriving the Environmental State

As for every wireless technology, also the performance of UWB is strongly affected by the surrounding environment.

LOS/NLOS. The presence of obstacles blocking the line-of-sight between nodes significantly decreases the received signal strength and, therefore, the quality of a link. In time-based localization systems, such as the ones using UWB radios, this leads to positively biased range estimations. Several research groups have investigated concepts to mitigate the impact of NLOS scenarios on UWB-based localization systems. Maranò et al. [34] apply machine learning techniques on the CIR to assess whether a signal was transmitted in LOS or NLOS conditions. This technique can be exploited to derive a binary information about the presence of LOS/NLOS conditions to be included in the link state indicator: this info can be of great value also to increase the reliability of UWB communications. One can, for example, temporarily re-route packets via another link providing LOS conditions, or make use of directional antennas to focus the radiated energy in the direction of a multipath component that is not blocked by an obstacle [35].

Destructive interference. The presence of several walls and obstacles in proximity of a wireless node results in a high amount of multipath reflections. The overlapping of several multipath reflections may lead to a deep fade due to destructive interference: as experimentally shown in Sect. II-E, this can heavily degrade a UWB communication link. To escape such destructive interference, the only viable solution is a switch to another carrier frequency, as merely changing the data rate or varying the bandwidth would not be of significant help. We show next in Sect. IV-B how to derive the presence of destructive interference using the estimated CIR and embed this info as an environmental state bit in the link state indicator sketched in Fig. 12. To the best of our knowledge, we are the first to present an efficient technique detecting the presence of destructive interference in UWB communication links.




Fig. 10. CIR estimation in the presence and absence of destructive interference, as well as at the end of the SHR detection range.

B. Detecting Destructive Interference using CIR Information

We record the CIR estimation in the setup described in Sect. II-E at a distance of 19 and 21 m from the sender, i.e., shortly before, and inside the fade, as well as at a distance of 60 m, i.e., at the end of the SHR detection range. Fig. 10 shows the three CIR curves. At a distance of 19 m, the link is still highly reliable (see Fig. 3), thus its CIR (solid blue line) does not exhibit presence of destructive interference. At a distance of 21 m, the PRR of the link drops to zero: the CIR shown in Fig. 10 (dashed green line) exhibits a reduced amplitude in the LOS component. In contrast, the amplitude of the multipath components is not affected as much as the one of the LOS. Instead, at 25 ns, the amplitude of a multipath component was even stronger than the one measured at a distance of 19 m. This is a clear indication that there is destructive interference between the LOS component and another reflection. Indeed, in the CIR measured at a distance of 60 m (red dotted line) the LOS path is strongly attenuated, but so are all other multipath components, whose amplitudes are close to the noise level.

Based on these observations, we propose an efficient mechanism to detect the presence of destructive interference. The algorithm keeps track of the ratio PR between the power of the LOS component and the power of the multipath components which can be derived from the CIR $\hat{h}[n]$.

$$PR = \frac{\sum_{n=1}^{N_{LOS}} |\hat{h}[n]|^2}{\sum_{n=N_{LOS}}^{N_{MP}} |\hat{h}[n]|^2} \quad (1)$$

We derive suitable values for N_{LOS} and N_{MP} empirically, and propose the use of a threshold $PR_{TH} = 1.5$ to assess the presence of destructive interference. When using $N_{LOS} = 5$ and $N_{MP} = 30$, we obtain a PR of 8.28, 1.26, and 3.61 at 19, 21, and 60 meters, respectively. Thus, in absence of destructive interference and at the end of the communication range, PR is respectively 6.57 and 2.87 times larger than the one computed inside the fade. We show the effectiveness of this algorithm in Sect. VI. Please note that the presence of destructive interference can also be derived without the reception of a complete packet, as the CIR is solely estimated from the preamble. Hence, due to the higher robustness of the SHR modulation (as shown in Sect. III-B), one can also derive environmental state information with highly unreliable links.

C. Deriving the Received Signal Power

We have shown how to make use of the CIR to estimate the power of the LOS and multipath components already in the previous section. The *integral of the entire CIR* can be used to also provide an estimate of the total received signal power (RSP) and, consequently, to instantaneously estimate the *quality* of a link. The DW1000 provides a register indicating the integral of the CIR that we use to derive the RSP as a




Fig. 11. PRR as a function of the received signal power (RSP) measured using the DW1000 default settings (violet) and a data rate of 110 kbps (blue).

dBm value [23]. We then characterize the relationship between RSP and PRR experimentally, by using the data collected in the hallway experiments described in Sect. II-E.

Fig. 11 illustrates the distribution of PRR as a function of the RSP (mean \pm standard deviation) when using the DW1000 default settings (violet) and when using a data rate of 110 kbps (blue). Each point is computed over 1000 packets, and the interpolated dotted lines illustrate the trend for these two PHY configurations as well as the relation between PRR and RSP.

Fig. 11 also shows that the latter *varies for different PHY settings*: the curve obtained with a data rate of 110 kbps is indeed shifted by 5-6 dB to the left compared to the one obtained with the default settings. This is consistent with the results captured with the cable-based setup shown in Sect. III-C, which have shown that the PRR depends on the employed PHY settings. Indeed, also Fig. 7 shows a difference of 5-6 dB between the use of the default settings and a data rate of 110 kbps. This ability of the RSP to inherently capture differences in the performance of PHY settings without the need of a correction is clearly an asset.

The RSP value can be used to estimate the link quality and to monitor if the sustained PRR becomes insufficient (e.g., below a PRR of 90%, as indicated by the dotted red line in Fig. 11). To this end, one can define a threshold RSP_{TH} , whose value can either be selected depending on the current PHY configuration, or by considering the worst-case: in Fig. 11, this corresponds to a RSP_{TH} of -93.3 dBm.

V. ADAPTING PHY SETTINGS AT RUNTIME

We design an adaptive scheme that determines at runtime a configuration of PHY settings increasing the dependability of UWB communications. To this end, our scheme makes use of the RSP to detect a degrading link. It then exploits a ranking of PHY settings (derived from the application requirements and from our experimental observations in Sect. III-E) to choose a new PHY configuration for the UWB radio. The latter is also driven by environmental state info computed from the CIR, such as the presence of destructive interference (see Fig. 12).

Detection of a degrading link. The link state indicator is used to trigger a change of the PHY setting in case of a degraded channel. For this purpose, the estimated received signal power is compared with the RSP_{TH} threshold derived in Sect. IV-C. If the link quality is below this threshold, the adaptation logic initiates a setting change to increase link reliability. We further define a second threshold RSP_{EE} defining the received signal power at which the quality of the channel allows to select a more efficient PHY setting. In the evaluation in Sect. VI we have set $RSP_{EE} = -88.8$ dBm, which ensures a good tradeoff between energy savings and a high PRR.




Fig. 12. Sketch of the inputs of our adaptation logic deriving a new PHY settings configuration for the Decawave DW1000 UWB transceiver.

Ranking UWB PHY settings. To keep the proposed solution generic, the adaptation logic optimizes its operation depending on the application requirements. For example, based on the requirements of the application at hand, one can define an *energy-conservative* adaptation policy maximizing reliability while minimizing the energy consumption. Similarly, for applications demanding a high reliability, one can define an *aggressive* adaptation policy maximizing reliability regardless of the energy expenditure. The ranking of the PHY settings is, besides the application requirements, also dependent on the impact of PHY settings on the robustness and efficiency. We derive this information from our experimental characterization presented in Sect. III-C. According to Table III, an energy-conservative ranking of UWB PHY settings is: (1) lower carrier frequency, (2) higher bandwidth, (3) higher PRF, (4) lower data rate. An aggressive ranking of PHY settings gives a lower data rate the highest priority. Discrete choices within a PHY setting (e.g., carrier frequency) are also ranked accordingly.

Influence of environmental state on adaptation logic. In case of an abrupt change in the surrounding environment (such as a blocked LOS or destructive interference), the ranking of the UWB PHY settings has to be adapted accordingly. For example, in the presence of a destructive interference, the change of the carrier frequency should always get highest priority, independently from the ranking of the settings.

Coordination of configuration changes. In order to correctly communicate, it has to be ensured that the PHY settings of all involved parties are synchronized. Thanks to the properties of UWB highlighted in Sect. II-B, changes in data rate and PSR *do not require a prior agreement* between transmitter and receiver, as this info can be derived at runtime from the SHR and PHR. Changes in other PHY settings (e.g., carrier frequency), instead, require coordination by means of a packet exchange. If this is the case, two aspects need to be considered. First, the information about the configuration change embedded in these packets should be acknowledged reliably, so to avoid disagreements between nodes [32]. Thanks to the higher robustness of the SHR, one can make use of a preamble as a binary ACK, as highlighted in Sect. III-B: this also allows to mitigate the problem of asymmetric links [21]. Second, one may need to adapt the local PHY settings of a node based on the configuration of all its neighbors. In our current implementation, we use the settings of the node with the worst link state as a global PHY configuration. However, as in a large (multi-hop) network it is unlikely that the same PHY configuration can meet the requirements of all nodes, one may

investigate the use of different time-slots for each neighbor or define multiple clusters in the network accordingly [36]. Such networking-wide issues are beyond the scope of the paper and subject of future work.

Implementation. We implement the proposed adaptation scheme by piggybacking 2 bytes in each exchanged packet. The first byte embeds the next PHY configuration to be used. The second byte contains link state information and consists of 7 bits RSP and 1 bit indicating the presence of destructive interference. If the current link state shows the necessity for a new configuration, the adaptation logic follows the devised ranking. For PHY settings with more than two discrete choices (e.g., carrier frequency), we first use up all the possible values, before moving on with the next PHY setting in the ranking. We provide an implementation of both the aggressive and energy-conservative policy, and evaluate their performance in Sect. VI.

VI. EVALUATION

We evaluate next the performance of the proposed adaptive scheme. First, we compare the reliability and energy efficiency of UWB communications in highly-dynamic environments when using static and adaptive PHY settings (Sect. VI-A). We then show the effectiveness of our approach in detecting and escaping destructive interference (Sect. VI-B).

A. Increasing the Dependability of UWB Communications

We design a scenario where two EVB1000 nodes (a transmitter and a receiver) are mounted on tripods at a distance of 5 m. The transmitter periodically sends packets to the receiver at a rate of 50 Hz. To emulate the mobility of nodes in a reproducible manner, we insert a programmable attenuator between the SMA connector and the transmitter's antenna. This enables us to simulate repeatable changes of the received signal strength due to varying environmental conditions. The attenuation sequence over time is shown in Fig. 13 (top). One can recognize two phases: a triangle-shaped sequence with 2-5s step size (0 to 130 s) that we use to push the devices to the edge of their communication range and a random sequence (130 to 360 s) emulating irregular fluctuations.

Starting with the DW1000 default settings, we compare the reliability and energy efficiency of static PHY settings, with an adaptive PHY configuration following an energy-conservative and an aggressive adaptation policy. Fig. 13 (second from top) shows the PRR over time sustained by the three schemes. The use of static PHY settings (blue, circles) causes the link to lose a large number of packets. The use of the proposed adaptive scheme, instead, allows to sustain a high packet delivery rate, with an average PRR always higher than 95 and 98% when using an energy-conservative (orange, squares) and an aggressive (yellow, diamonds) adaptation policy, respectively.

Fig. 13 (third from top) shows the measured RSP over time. The red solid line marks the RSP_{TH} threshold (-93.3 dBm) used to trigger the change to a more robust PHY configuration. The red dotted line marks the RSP_{EE} threshold (-88.8 dBm) used to trigger a more energy-efficient PHY configuration. Please note that, in this scenario, we do not make use of environmental state information to detect destructive interference.




Fig. 13. Comparison of the reliability (PRR) and energy expenditure of a static PHY configuration with the one sustained by an adaptive PHY configuration following an energy-conservative or an aggressive adaptation policy.

Fig. 13 (bottom) shows the energy cost of the three schemes over time, computed as the energy per packet in μJ for both transmission and reception. The adaptive logic allows to sustain a higher reliability at a cost of a 5 and 29% higher energy expenditure for the energy-conservative and aggressive policy, respectively. We can distinguish four main phases: in the first and third (Ph. 1 and 3), the adaptation logic triggers a more energy-efficient PHY setting, since the reliability of the link allows it. This saves significant energy compared to the default setting (blue, solid). In a second phase, the high attenuation requires, among others, the use of a lower data rate to increase the reliability of the link: this comes at an increased energy cost (that is higher when using the aggressive adaptation policy). The difference between the energy-conservative (orange, dashed) and the aggressive (yellow, dotted) adaptation policy is even more distinctive in the fourth phase.

B. Recognizing and Escaping Destructive Interference

We evaluate next the ability of the proposed adaptation algorithm to recognize and escape destructive interference at runtime. Towards this goal, we reuse the experimental setup described in Sect. II-E and keep the transmitter at a static position, while moving the receiver across the hallway. In particular, the receiver is moved from a distance of 18 m (reliable link) to a distance of 20.5 m, hence entering the area where destructive interference is present (see Fig. 3). After 30 seconds, the receiver is moved back to a distance of 18 m. We then compare the PRR when using default static PHY settings and when using an aggressive adaptation policy that makes use or does not make use of the environmental state information capturing the presence of destructive interference.

Fig. 14 (top) shows the PRR for these three schemes. When using static PHY settings (blue line with circles), the movement of a few meters degrades the link up to a point in which packets are no longer received. The same occurs when using an aggressive adaptive scheme that does not make use of environmental state information (red line with squares), as lowering the data rate does not help significantly in the presence of destructive interference. The aggressive adaptive




Fig. 14. The proposed adaptation scheme can correctly recognize and escape destructive interference, sustaining a high reception rate over time.

scheme (yellow line with diamonds), instead, recognizes and escapes destructive interference by directly changing the carrier frequency without a switch to a lower data rate.

Fig. 14 (bottom) shows the calculation of the ratio between LOS and multipath power PR (blue dotted line). At time 30s, as soon as PR is smaller than the PR_{TH} threshold of 1.5 (red solid line), the destructive interference bit of the environmental state indicator is set to high and a change of PHY configuration (carrier frequency) is triggered. After this change, the measured PR is again well above the PR_{TH} threshold: this is expected as, after the change of carrier frequency, the node no longer suffers destructive interference.

VII. RELATED WORK

UWB technology has attracted a large body of research, especially after the formal addition of the UWB physical layer to the IEEE 802.15.4 standard in 2007, and the availability of the first compliant low-cost transceivers.

UWB-based localization. Undoubtedly, the interest of the community was mostly attracted by the possibility of using UWB to carry out precise positioning. Researchers have produced a large number of algorithms to maximize the positioning accuracy indoors [3], for example by exploiting multipath reflections [8], or by making use of inertial sensors [9]. Other works have instead investigated how to make use of antenna diversity [26], as well as of directional antennas [35] to increase localization accuracy. Given the large number of techniques that have been produced, the community started to quantitatively compare the accuracy of different localization systems in competitions [6]. The latter showed the benefits of UWB and spread the word about its potential. Differently from these works, our study does not focus on accuracy, but instead on increasing the dependability of UWB communications.

UWB measurements. Researchers have experimentally quantified the effect of LOS and NLOS [37], the impact of clock offset on different ranging schemes [38], as well as the UWB operating range as a function of different preamble parameters [39]. Unlike these works, we systematically study the performance of *both* preamble and payload settings and experimentally quantify their sensitivity and energy cost, in order to understand which configuration to privilege. A work in spirit close to ours is the one by Mikhaylov et al. [40], who have experimentally measured the impact of various settings on the accuracy of indoor localization. The focus of their work, however, is not on communication performance.

UWB link quality estimation. A plethora of works has proposed different strategies to estimate the quality of a wireless link [21]. The proposed solutions have combined the knowledge of the expected number of transmissions [41] with information from the network, link, and physical layers [42], often exploiting the link quality indicators returned by IEEE 802.15.4 narrowband hardware such as RSSI and LQI [43]. The link quality information that can be retrieved from UWB radios is, however, unexplored, and radically different from the ones that have been studied in narrowband IEEE 802.15.4 transceivers by the aforementioned studies. To the best of our knowledge, the work presented in this paper is the first one studying link state estimation with UWB radios and proposing an efficient technique that makes use of the estimated CIR information to detect the presence of destructive interference in UWB communication links.

Runtime adaptation. Increasing the performance of a wireless system by adapting protocol parameters at runtime is a well-known problem. The research community has explored how to adapt the transmission power of the radio [28] or MAC protocol parameters such as duty cycle [16] and clear channel assessment threshold [17], but always focused on narrowband IEEE 802.15.4 transceivers only. To the best of our knowledge, we are the first studying an automatic selection of UWB parameters at runtime to improve communication performance.

VIII. CONCLUSIONS

In this paper, we have exploited UWB PHY settings as tuning knobs to increase the dependability of communications. Towards this goal, we have first experimentally quantified the reliability and energy cost of each setting and understood which physical layer configuration to privilege depending on the specific application requirements. We have then used the estimated CIR to accurately measure the link quality and to extract relevant information about the characteristics of the surrounding environment, such as the presence of destructive interference. Building upon this information, we have designed an adaptation scheme that tunes the PHY settings of an UWB transceiver at runtime, and demonstrated experimentally that it effectively increases the communication performance.

Future work will investigate the design of more complex algorithms that make use of the estimated CIR information to accurately characterize the environment and the use of directional antennas as an additional tuning knob.

ACKNOWLEDGMENTS

This work was supported by the TU Graz LEAD project “Dependable Internet of Things in Adverse Environments”.

REFERENCES

- [1] A. Molisch, “Ultra-Wide-Band Propagation Channels,” *Proceedings of the IEEE*, vol. 97, no. 2, Feb. 2009.
- [2] Z. Irahautent et al., “UWB Channel Measurements and Results for Office and Industrial Environments,” in *Proc. of the ICUWB Conf.*, 2006.
- [3] A. Alarifi et al., “Ultra Wideband Indoor Positioning Technologies: Analysis and Recent Advances,” *Sensors*, vol. 16, no. 5, May 2016.
- [4] M. Win, D. Dardari, A. Molisch, W. Wiesbeck, and J. Zhang, “History and Applications of UWB,” *Proc. of the IEEE*, vol. 97, no. 2, 2009.
- [5] J. Zhang et al., “UWB Systems for Wireless Sensor Networks,” *Proc. of the IEEE*, vol. 97, no. 2, 2009.
- [6] D. Lymberopoulos et al., “The Microsoft Indoor Localization Competition,” *IEEE Sig. Proc. Mag.*, vol. 34, no. 5, pp. 125–140, 2017.
- [7] D. Dardari, P. Closas, and P. Djuric, “Indoor Tracking: Theory, Methods, and Technologies,” *IEEE Trans. on Vehic. Techn.*, vol. 64, no. 4, 2015.
- [8] J. Kulmer et al., “Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning,” in *IEEE ANLN Workshop*, 2017.
- [9] M. Kok et al., “Indoor Positioning Using Ultrawideband and Inertial Measurements,” *IEEE Trans. on Vehicular Techn.*, vol. 64, no. 4, 2015.
- [10] D. Espes et al., “Ultra-Wideband Positioning for Assistance Robots for Elderly,” in *Proc. of the 10th SPPRA Conference*, Feb. 2013.
- [11] R. Chávez-Santiago et al., “Ultrawideband Technology in Medicine: A Survey,” *Journal of Electrical and Computer Engineering*, Feb. 2012.
- [12] K. Guo et al., “Ultra-Wideband-Based Localization for Quadcopter Navigation,” *Unmanned Systems Journal*, vol. 4, no. 1, Jan. 2016.
- [13] B. Kempke et al., “PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization,” in *Proc. of HotWireless Workshop*, 2015.
- [14] J. Lee et al., “A Comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee, and Wi-Fi,” in *Proc. of the IECON Conf.*, 2007.
- [15] P. Alcock, U. Roedig, and M. Hazas, “Combining Positioning and Communication Using UWB Transceivers,” in *Proc. of the 5th IEEE DCOSS Conference*, Jun. 2009.
- [16] M. Zimmerling et al., “pTunes: Runtime Parameter Adaptation for Low-power MAC Protocols,” in *Proc. of IPSN Conference*, Apr. 2012.
- [17] M. Sha et al., “Energy-efficient Low Power Listening for Wireless Sensor Networks in Noisy Environments,” in *IPSN Conf.*, Apr. 2013.
- [18] S. Brienza et al., “Just-in-Time Adaptive Algorithm for Optimal Parameter Setting in 802.15.4 WSNs,” *ACM TAAAS*, vol. 10, no. 4, 2016.
- [19] A. Augustin et al., “A Study of LoRa: Long Range & Low Power Networks for the Internet of Things,” *Sensors*, vol. 16, no. 9, Sep. 2016.
- [20] M. Bor and U. Roedig, “LoRa Transmission Parameter Selection,” in *Proc. of the 13th IEEE DCOSS Conference*, Jun. 2017.
- [21] N. Baccour et al., “Radio Link Quality Estimation in Wireless Sensor Networks: a Survey,” *ACM Trans. Sens. Netw.*, vol. 8, no. 4, Nov. 2012.
- [22] Y. Qi et al., “Clear Channel Assessment with Multiplexed Preamble Symbols for Impulse UWB Communications,” in *ICUWB Conf.*, 2006.
- [23] Decawave Ltd., *DW1000 User Manual. Version 2.10*, 2016.
- [24] Federal Communications Commission (FCC), *Revision of Part 15 of the Commission’s Rule Regarding UWB Transmission Systems*, Std., 2002.
- [25] ITU-R, *Recommendation ITU-R SM. 1755, Characteristics of UWB technology*, Std., 2006.
- [26] B. Kempke et al., “SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization,” in *Proc. of the 14th ACM SenSys Conference*, Nov. 2016.
- [27] Decawave Ltd., *Application Note APS023 (Part 1+2), version 1.0*, 2016.
- [28] S. Lin et al., “ATPC: Adaptive Transmission Power Control for Wireless Sensor Networks,” in *Proc. of the 4th ACM SenSys Conference*, 2006.
- [29] A. Molisch, *Wireless communications*. John Wiley & Sons, 2012.
- [30] Decawave Ltd., *Application Note APS013, version 2.0*, 2015.
- [31] M. Zúñiga and B. Krishnamachari, “Analyzing the Transitional Region in Low-Power Wireless Links,” in *Proc. of IEEE SECON Conf.*, 2004.
- [32] C. Boano et al., “JAG: Reliable and Predictable Wireless Agreement under External Radio Interference,” in *Proc. of IEEE RTSS Symp.*, 2012.
- [33] A.F. Molisch et al., “IEEE 802.15.4a Channel Model,” Tech. Rep., 2004.
- [34] S. Marano et al., “NLOS Identification and Mitigation for Localization Based on UWB Experimental Data,” *IEEE J-SAC*, vol. 28, no. 7, 2010.
- [35] M. Rath et al., “Multipath-assisted Indoor Positioning Enabled by Directional UWB Sector Antennas,” in *IEEE SPAWC Workshop*, 2017.
- [36] C. Frank and K. Römer, “Algorithms for Generic Role Assignment in Wireless Sensor Networks,” in *Proc. of the 3rd SenSys Conf.*, 2005.
- [37] T. Ye et al., “An Experimental Evaluation of IEEE 802.15.4a UWB Technology for Precision Indoor Ranging,” *IJACI*, vol. 4, no. 2, 2012.
- [38] K. Mikhaylov et al., “On the Selection of Protocol and Parameters for UWB-based Wireless Indoor Localization,” in *ISMICT Symp.*, 2016.
- [39] T. Gigl et al., “Ranging Performance of the IEEE 802.15.4a UWB Standard Under FCC/CEPT Regulations,” *JECE*, Jan. 2012.
- [40] K. Mikhaylov et al., “Impact of IEEE 802.15.4 Communication Settings on Performance in Asynchronous Two Way UWB Ranging,” *Intl. Journal of Wireless Information Networks*, vol. 24, no. 2, Jun. 2017.
- [41] D. De Couto et al., “A High-throughput Path Metric for Multi-hop Wireless Routing,” in *Proc. of the 9th MobiCom Conference*, Sep. 2003.
- [42] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis, “Four-Bit Wireless Link Estimation,” in *Proc. of the 6th HotNets Workshop*, Nov. 2007.
- [43] K. Srinivasan et al., “An Empirical Study of Low-power Wireless,” *ACM Transactions on Sensor Networks*, vol. 6, no. 2, Mar. 2010.

Paper B

B. Grosswindhager, M. Rath, J. Kulmer, M.S. Bakr, C.A. Boano, K. Witrisal, and K. Römer. **SALMA: UWB-based Single-Anchor Localization System Using Multipath Assistance**. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems (SenSys'18)*, pages 132–144, Shenzhen, China. November 2018.

©2018 ACM

ISBN: 9781450359528

DOI: 10.1145/3274783.3274844

Link: <https://dl.acm.org/doi/10.1145/3274783.3274844>

Abstract. Setting up indoor localization systems is often excessively time-consuming and labor-intensive, because of the high amount of anchors to be carefully deployed or the burdensome collection of fingerprints. In this paper, we present SALMA, a novel low-cost UWB-based indoor localization system that makes use of only one anchor and that does neither require prior calibration nor training. By using only a crude floor plan and by exploiting multipath reflections, SALMA can accurately determine the position of a mobile tag using a single anchor, hence minimizing the infrastructure costs, as well as the setup time. We implement SALMA on off-the-shelf UWB devices based on the Decawave DW1000 transceiver and show that, by making use of multiple directional antennas, SALMA can also resolve ambiguities due to overlapping multipath components. An experimental evaluation in an office environment with clear line-of-sight has shown that 90% of the position estimates obtained using SALMA exhibit less than 20 cm error, with a median below 8 cm. We further study the performance of SALMA in the presence of obstructed line-of-sight conditions, moving objects and furniture, as well as in highly dynamic environments with several people moving around, showing that the system can sustain decimeter-level accuracy with a worst-case average error below 34 cm.

My contribution. Together with Michael Rath, I am the main author of this publication. We have both contributed equally to this work. My contributions were mainly focused on the hardware and implementation, whereas Michael Rath's main part was the design of the positioning algorithms. We have carried out the evaluations together. Especially related to the algorithms, this work would not have been possible without the tremendous help of Josef Kulmer. Carlo Boano helped us in writing and structuring the paper and Mustafa Bakr contributed the antennas used in SALMA-full. I have presented the paper at SenSys'18.

SALMA: UWB-based Single-Anchor Localization System using Multipath Assistance

Bernhard Großwindhager[†], Michael Rath[†], Josef Kulmer, Mustafa S. Bakr, Carlo Alberto Boano, Klaus Witrisal, and Kay Römer

Faculty of Electrical and Information Engineering, Graz University of Technology, Austria
{grosswindhager, mrath, kulmer, mustafa.bakr, cboano, witrisal, roemer}@tugraz.at

[†] Authors contributed equally to this work

ABSTRACT

Setting up indoor localization systems is often excessively time-consuming and labor-intensive, because of the high amount of anchors to be carefully deployed or the burdensome collection of fingerprints. In this paper, we present SALMA, a novel low-cost UWB-based indoor localization system that makes use of only one anchor and that does neither require prior calibration nor training. By using only a crude floor plan and by exploiting multipath reflections, SALMA can accurately determine the position of a mobile tag using a single anchor, hence minimizing the infrastructure costs, as well as the setup time. We implement SALMA on off-the-shelf UWB devices based on the Decawave DW1000 transceiver and show that, by making use of multiple directional antennas, SALMA can also resolve ambiguities due to overlapping multipath components. An experimental evaluation in an office environment with clear line-of-sight has shown that 90% of the position estimates obtained using SALMA exhibit less than 20 cm error, with a median below 8 cm. We further study the performance of SALMA in the presence of obstructed line-of-sight conditions, moving objects and furniture, as well as in highly dynamic environments with several people moving around, showing that the system can sustain decimeter-level accuracy with a worst-case average error below 34 cm.

CCS CONCEPTS

• **Computer systems organization** → *Embedded and cyber-physical systems*; • **Networks** → *Location based services*;

KEYWORDS

Indoor localization, ultra-wideband, multipath, single-anchor.

ACM Reference Format:

Bernhard Großwindhager[†], Michael Rath[†], Josef Kulmer, Mustafa S. Bakr, Carlo Alberto Boano, Klaus Witrisal, and Kay Römer. 2018. SALMA: UWB-based Single-Anchor Localization System using Multipath Assistance. In *The 16th ACM Conference on Embedded Networked Sensor Systems (SenSys '18)*, November 4–7, 2018, Shenzhen, China. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3274783.3274844>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SenSys '18, November 4–7, 2018, Shenzhen, China

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5952-8/18/11...\$15.00

<https://doi.org/10.1145/3274783.3274844>

1 INTRODUCTION

Localizing people and objects in a precise and accurate way is a key requirement for future location-aware Internet of Things (IoT) applications such as assisted living [54], health care [20], and robot navigation [14, 24]. As of today, achieving an accurate position estimation is still a grand challenge especially *indoors*, where global navigation satellite systems such as GPS, Galileo, GLONASS, and Beidou are not applicable due to the limited signal reception [4].

Challenges of indoor positioning. When estimating the position of a device indoors, one needs to deal with severe *attenuation, multipath, and scattering* of signals due to walls, furniture, or other surrounding objects. This is, for example, a major challenge for localization systems making use of narrowband RF technologies: solutions based on IEEE 802.15.4 [18, 29], Bluetooth [1, 3], and Wi-Fi [9, 16] are indeed highly susceptible to multipath fading, and can hardly achieve a sub-meter accuracy in these settings [38].

Indoor environments are also *highly dynamic*: moving people and objects may obstruct the line-of-sight (LOS) path between a source and a receiver. This is, for example, particularly challenging for localization systems based on optical technologies. Furthermore, the unpredictable presence of interference sources (e.g., co-located wireless devices using the same frequency band) can cause loss of information and fluctuations in the received signal strength that drastically affect the accuracy of positioning algorithms.

A practical localization system needs to achieve a high positioning accuracy despite these inherent properties of indoor environments. This task is further complicated by the fact that an ideal indoor localization system should maximize the accuracy, efficiency, and responsiveness of position estimation, while *minimizing deployment efforts and costs*. Although a plethora of distinct approaches has been proposed in the literature, none of them can yet achieve a high accuracy at minimal costs and is thus widely accepted [35].

Deployment overhead still too high. After comparing the performance of more than 100 state-of-the-art indoor localization systems under the same settings, Lymberopoulos and Liu [35] have concluded that the set-up procedure of existing solutions is excessively *time-consuming* and *labor-intensive*. As a consequence, the use of most systems is still impractical in real-world deployments.

This state of affairs represents a serious problem, because recent solutions based on Ultra-wideband (UWB) could easily achieve accuracies in the order of decimeters [24, 31]. However, one cannot fully exploit this outstanding positioning accuracy, because of the high overhead in deploying the required infrastructure. These systems employ indeed multiple anchors (e.g., at least 8 [52], 9 [23],

or 15 [24]), each of which needs to be carefully placed [19] in order to maximize the system's performance – a burden that is not sustainable on a large scale. Similarly, localization systems based on RSS profiling [16, 43] require a laborious offline data collection process before deployment to acquire the radio maps (fingerprints). Even worse, fingerprinting needs to be frequently repeated to cope with environmental changes, such as furniture setup and human motion: this makes the installation costs prohibitively high.

Minimizing the deployment effort of localization systems while still allowing to sustain a high positioning accuracy is hence a long-due fundamental step towards the creation of solutions that are viable for real-world IoT applications.

Contributions. In this paper we present SALMA, a novel UWB-based indoor localization system that can sustain a decimeter-level accuracy despite the use of only a *single* anchor. SALMA removes the need of multiple anchors by exploiting multipath propagation, i.e., specular reflections originating from static objects. The system works out of the box without any time-consuming setup phase, as it does not require any prior calibration, training, or position estimates (i.e., SALMA is based on neither fingerprinting nor other learning algorithms). All that is needed is a crude floor plan showing the geometry of the building in which the system is installed. The map includes static objects such as walls and windows only, to avoid adaptations if furniture or other objects are moved.

Starting from this floor plan and the known location and orientation of the anchor, SALMA models the theoretical multipath propagation and compares it with the estimated channel impulse response (CIR) derived by the anchor node, as shown in Sect. 3. Exploiting the position-related information embedded in the CIR allows to unambiguously determine the position of a tag using a single anchor with an accuracy comparable to the one achieved by common multi-anchor UWB systems. This way, SALMA reduces the infrastructure costs and setup time, hence addressing the omnipresent trade-off between accuracy and deployment costs.

We implement SALMA on off-the-shelf UWB devices based on the popular Decawave DW1000 transceiver, building – to the best of our knowledge – the *first* low-cost single-anchor UWB-based indoor localization system. In particular, as shown in Sect. 4, we support multiple tags simultaneously and shift the burden of position estimation to the anchor node. This allows to keep the design of the mobile tag simple, so to preserve its limited battery capacity.

We specifically implement SALMA for two-dimensional settings in order to support map-based navigation and tracking applications such as locating patients in hospitals [21], assistance for visually impaired, disabled, and elderly people [7, 40, 47, 54], as well as monitoring sport events [33, 45]. The applicability to three-dimensional settings and the resulting challenges are discussed in Sect. 8.

We further show how overlapping multipath components (MPCs) may limit the performance of SALMA when using a single omnidirectional antenna. To alleviate this problem, we illustrate in Sect. 5 how to improve the robustness of SALMA using *multiple directional antennas*. The latter enable the exploration of the angular information of MPCs to enhance the system's performance remarkably.

In Sect. 6, we carry out a thorough experimental evaluation¹ of the performance of SALMA in an office environment with clear LOS

conditions. Among others, our results reveal that 90% of position estimates obtained with SALMA exhibit less than 20 cm error, with a median below 8 cm. This performance was obtained with a single measurement snapshot from four directional antennas. We further simulate how the accuracy of SALMA can be improved with a higher number of antennas with narrower bandwidth.

In Sect. 7, we study the performance of SALMA in the presence of *obstructed LOS*, showing that 90% of position estimates exhibit less than 30 cm error, with a median below 15 cm. Furthermore, we evaluate SALMA in a challenging setting (stockroom) reaching a 90% error of 44.5 cm, and show that moving objects and furniture has a limited effect on the performance. We also deploy SALMA in an office crowded with tens of people moving in/out across 24-hours, and show that – even in such a highly dynamic environment – SALMA sustains a worst-case average error below 34 cm. Therefore, by exploiting the redundancy offered by multipath reflections, SALMA achieves a high accuracy even with obstructed LOS, addressing an inherent vulnerability of traditional systems.

In summary, this paper makes the following contributions:

- We present SALMA, a UWB-based indoor localization system making use of only a single anchor and requiring neither prior profiling nor calibration (Sect. 3);
- We implement SALMA on off-the-shelf UWB devices and support multiple tags simultaneously (Sect. 4);
- We increase the robustness of the system to overlapping MPCs by using multiple directional antennas (Sect. 5);
- We evaluate the performance of SALMA experimentally in different scenarios with clear LOS and show that 90% of position estimates exhibit less than 20 cm error (Sect. 6);
- We show that SALMA is resilient to obstructed LOS situations and that it sustains a high accuracy even in dynamic environments with objects and people moving around (Sect. 7).

2 SALMA: OVERVIEW

Fig. 1a shows a sketch of SALMA's design. The system consists of a single anchor (fixed infrastructure) and multiple battery-powered mobile tags (devices to be localized). The anchor is connected to and powered by a central notebook running a *localization engine* that computes the position of each tag.

Every tag initiates a double-sided two-way ranging (DS-TWR) with the anchor node, following a time division multiple access scheme. The two-way ranging process allows the anchor to estimate the distance $d_0 = \|\mathbf{p} - \mathbf{a}\|$, with \mathbf{p} and \mathbf{a} being the tag and the anchor location, respectively (Fig. 1b). Upon completion of the DS-TWR process, the anchor records the estimated distance \hat{d}_0 , as well as an estimate of the channel impulse response (CIR) provided by the UWB transceiver, and forwards this info to the localization engine.

Exploiting multipath propagation. The CIR embeds information about the multipath propagation consisting of reflections from walls. Traditional UWB localization systems employ the CIR to estimate the distance d_0 , which is related to the path delay τ_0 as follows: $d_0 = \tau_0 \cdot c_0$, with c_0 being the speed of light. Therefore, these systems only use the path delay τ_0 , and forgo remaining multipath components (MPCs). SALMA, instead, additionally uses delays of reflected multipath components, which contain additional geometric information (cf. τ_k and d_k in Fig. 1b for $k = 1, \dots, 4$).

¹All datasets are publicly accessible under <http://www.iti.tugraz.at/SALMA> [12].

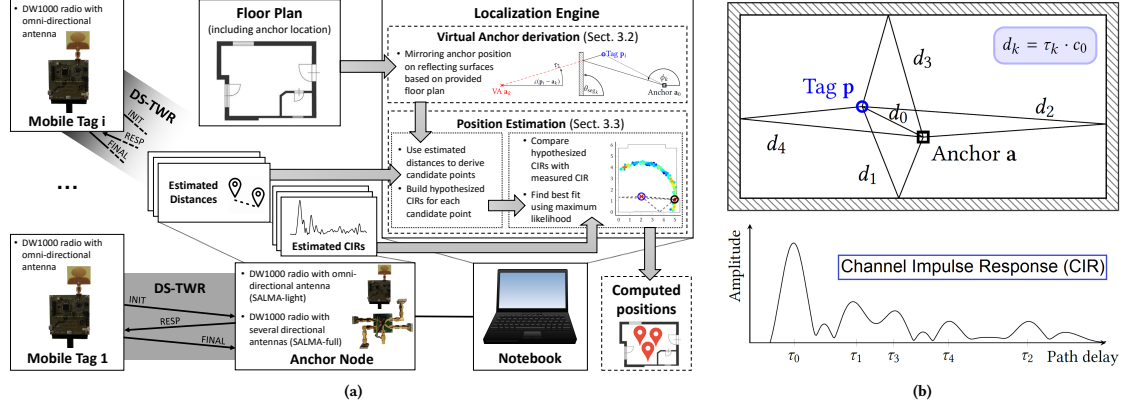


Figure 1: Overview of SALMA’s design (a): the system makes use of the multipath propagation between a single anchor and a tag i . The multipath propagation is characterized by the estimated CIR containing position-related information (b).

Localization engine. Starting from a floor plan showing the geometry of the building in which the system is installed², and the known location of the anchor, SALMA models the theoretical multipath propagation by employing the concept of virtual anchors [39] and by building a hypothesized CIR for several candidate positions. The latter are selected on a circle of radius \hat{d}_0 centered in \mathbf{a} , with \hat{d}_0 being the estimated distance derived from the DS-TWR. The localization engine then compares the hypothesized CIR of each candidate point with the one measured through the DS-TWR process, and returns the best fit using maximum likelihood estimation.

As we will show in the next sections, exploiting the position-related information encoded in the MPCs allows to unambiguously and accurately determine the position of a tag using a single anchor.

3 SALMA: DESIGN PRINCIPLES

We describe next the mathematical principles behind the functional stages of SALMA, showing how the system can leverage the information that is contained in the observed CIR to accurately narrow down the tag position. First, we present a model of the observed CIR including the multipath component (MPC) parameters in Sect. 3.1. We then explain in Sect. 3.2 how to use the known anchor position and floor plan to determine *virtual anchors* that can relate the tag position to parameters embedded in the CIR. Third, we describe in Sect. 3.3 how these parameters are used in combination with the observed CIR to obtain a position estimate. Sect. 4 then outlines how these methods are implemented on off-the-shelf hardware.

3.1 Signal model

Taking advantage of multipath propagation requires its proper modeling. In the following, we introduce the signal model relating the effective system response (i.e., the observed CIR) and the parameters of multipath components. We assume that a tag is equipped with a single omni-directional antenna, while the anchor can carry

M antennas. Each antenna with index $m = 1, \dots, M$ is characterized with its beampattern $b_m(\phi)$. The observed CIR $r_m(t)$ between a single tag and the anchor’s m^{th} antenna can be modeled as:

$$r_m(t) = \sum_{k=0}^K \alpha_k b_m(\phi_k) s_{\text{DW}}(t - \tau_k) + w_m(t). \quad (1)$$

The first term on the right-hand-side describes K specular MPCs, i.e., dominant reflections, of the transmitted signal $s_{\text{DW}}(t)$. The latter includes de-spreading and filtering at the receiver. Each MPC is characterized by its complex-valued amplitude α_k , angle of departure ϕ_k and delay τ_k . These MPCs are resulting from reflections at flat surfaces such as walls, windows or doors and will be further discussed in Sect. 3.2. The last term $w_m(t)$ denotes zero-mean white Gaussian measurement noise with variance σ_w^2 . Note that the proposed signal model in (1) can model single omni-directional as well as multiple directional antenna measurements.

The signal $r_m(t)$ is sampled with frequency $f_s = 1/T_s$ and N_s samples are acquired. Hence, we use vector notation [26, 27] to compactly describe the signal model in (1) as:

$$\mathbf{r} = \mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi}) \boldsymbol{\alpha} + \mathbf{w} \quad (2)$$

with

$$\begin{bmatrix} r_1 \\ \vdots \\ r_M \end{bmatrix} = \begin{bmatrix} X_1(\boldsymbol{\tau}, \boldsymbol{\phi}) \\ \vdots \\ X_M(\boldsymbol{\tau}, \boldsymbol{\phi}) \end{bmatrix} \boldsymbol{\alpha} + \begin{bmatrix} w_1 \\ \vdots \\ w_M \end{bmatrix} \quad (3)$$

and

$$\begin{aligned} \mathbf{r}_m &= [r_m(0 \cdot T_s), \dots, r_m([N_s - 1] \cdot T_s)]^T \\ \mathbf{X}_m(\boldsymbol{\tau}, \boldsymbol{\phi}) &= [b_m(\phi_0) s(\tau_0) \dots b_m(\phi_K) s(\tau_K)] \\ s(\tau_k) &= [s_{\text{DW}}(0 \cdot T_s - \tau_k), \dots, s_{\text{DW}}([N_s - 1] \cdot T_s - \tau_k)]^T \\ \mathbf{w}_m &= [w_m(0 \cdot T_s), \dots, w_m([N_s - 1] \cdot T_s)]^T \\ \boldsymbol{\phi} &= [\phi_0, \dots, \phi_K]^T; \boldsymbol{\tau} = [\tau_0, \dots, \tau_K]^T; \boldsymbol{\alpha} = [\alpha_0, \dots, \alpha_K]^T. \end{aligned}$$

Thus, the proposed signal model connects the MPC parameters (α_k, ϕ_k, τ_k) with the expected CIR. In Sect. 3.2, we relate these parameters to the tag position.

²While furniture and other objects do affect the performance of the system, the impact stays in reasonable bounds as demonstrated in Sect. 6 and 7. Thus, there is no need to keep track of whether tables, shelves, or other furniture have moved.




Figure 2: The concept of virtual anchors (VA) and its use in calculating the angle of departure ϕ_k and delay τ_k .

3.2 Geometric model and virtual anchors

The MPC parameters of the CIR contain position-related information regarding the tag location as well as the environment [28]. Following Euclidean geometry, simple relations can be obtained for ϕ_k and τ_k . In particular, we employ the concept of virtual anchors (VAs) [39] in order to relate MPC parameters to the tag positions (see Fig. 2). To obtain the positions of the virtual anchors \mathbf{a}_k ($k > 0$), the position of the physical anchor $\mathbf{a}_0 \triangleq \mathbf{a}$ is mirrored at each reflective flat surface. Fig. 2 illustrates the top-view of a single reflection. A specular MPC (black solid) originates at the wall segment. Assignment of the specular MPC to a virtual anchor (red cross) enables an efficient calculation of the MPC parameters, delay, and angle. The delay τ_k follows as geometric distance between tag and VA, divided by the speed of light c_0 , according to

$$\tau_k = \frac{1}{c_0} \|\mathbf{p} - \mathbf{a}_k\|. \quad (4)$$

We describe the angle of departure ϕ_k via the azimuth angle between tag and VA $\angle(\mathbf{p} - \mathbf{a}_k)$ according to

$$\phi_k = 2\theta_{\text{seg}_k} - \angle(\mathbf{p} - \mathbf{a}_k). \quad (5)$$

Here, θ_{seg_k} denotes the angle of the involved reflective surface that was used to generate the VA \mathbf{a}_k (see Fig. 2).

Note that, in this work, we limit the multipath propagation to *single-bounce* reflections, i.e., only a single reflective object is bounced during the path's propagation. Hence, the number of considered surfaces also determines the number of used MPCs K and VAs, e.g., for the floorplan shown in Fig. 1b we set $K = 4$ resulting in four VAs. In principle, the virtual anchor model can be extended to cover higher-order reflections as well. However, higher-order reflections are attenuated strongly, due to their increased path length and additional reflection losses. It should be also noted that, for each tag position \mathbf{p} , the visibility of the VAs has to be taken into account. This means that we have to check the direct path from \mathbf{p} to the VA position \mathbf{a}_k for intersections with any obstacles or wall segments. Only if there is a single involved intersection with the correct wall segment, we can use the k -th MPC in the signal model.

While the parameters τ_k and ϕ_k can be directly derived from the geometric model using the known VAs, a proper model for the MPC amplitudes α_k is difficult to obtain [27]. Hence, we propose to treat α_k as nuisance parameter, estimated directly from the observation \mathbf{r} .

3.3 Position estimation

In the following, we present a position estimator based on the CIR measurements. We aim for a maximum likelihood (ML) estimator, derived from the signal model in (2). To allow efficient computations, we assume complex-valued white Gaussian measurement noise \mathbf{w} .

The likelihood $p(\mathbf{r}|\mathbf{p})$ of observation \mathbf{r} conditioned on tag position \mathbf{p} follows as:

$$p(\mathbf{r}|\mathbf{p}) = \left(\frac{1}{\pi\sigma_w^2}\right)^{MN_s} \exp\left\{-\frac{1}{\sigma_w^2}\|\mathbf{r} - X(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha}\|^2\right\} \quad (6)$$

where $\boldsymbol{\tau}$ and $\boldsymbol{\phi}$ are related to the tag position via (4) and (5). Taking the log of (6) results in the log-likelihood function

$$\log p(\mathbf{r}|\mathbf{p}) = -MN_s \log(\pi\sigma_w^2) - \frac{1}{\sigma_w^2}\|\mathbf{r} - X(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha}\|^2. \quad (7)$$

This function depends on MPC amplitudes $\boldsymbol{\alpha}$. We propose to estimate $\boldsymbol{\alpha}$ as least squares solution [27] according to

$$\hat{\boldsymbol{\alpha}} = (X^H(\boldsymbol{\tau}, \boldsymbol{\phi})X(\boldsymbol{\tau}, \boldsymbol{\phi}))^{-1}X^H(\boldsymbol{\tau}, \boldsymbol{\phi})\mathbf{r} \quad (8)$$

with $(\cdot)^H$ denoting the conjugate and transposed. The position estimate $\hat{\mathbf{p}}$ maximizing the log-likelihood function can be formulated as a non-linear optimization problem:

$$\hat{\mathbf{p}} = \arg \max_{\mathbf{p} \in \mathcal{P}} \log p(\mathbf{r}|\mathbf{p}) = \arg \min_{\mathbf{p} \in \mathcal{P}} \|\mathbf{r} - X(\boldsymbol{\tau}, \boldsymbol{\phi})\hat{\boldsymbol{\alpha}}\|^2. \quad (9)$$

The parameters $\boldsymbol{\tau}$ and $\boldsymbol{\phi}$ are determined via the geometry, and these parameters in turn build the hypothesized CIR $X(\boldsymbol{\tau}, \boldsymbol{\phi})\hat{\boldsymbol{\alpha}}$, which is compared to the observed CIR \mathbf{r} . The position for which the hypothesis comes closest to the observation (and thus maximizes the likelihood) is chosen as the position estimate $\hat{\mathbf{p}}$. Searching for a global maximum requires to evaluate (9) at each feasible tag position \mathcal{P} , i.e., all positions within the communication range to the anchor. As shown in [26], this exhaustive search can be limited to potential candidate points that are located along a circle around \mathbf{a} with radius \hat{d}_0 . We consider N_C candidate points $\mathcal{P} = \{\mathbf{p}^{(j)}\}_{j=1}^{N_C}$ where each point is drawn independently with Gaussian distributed radius $d^{(j)} \sim \mathcal{N}(\hat{d}_0, \sigma_{\text{pw}}^2)$ and uniformly distributed angle $\phi^{(j)} \sim \mathcal{U}(0, 2\pi)$ [26]. Candidate points lying outside of the room are discarded. These can be determined with simple line equation tests using the given floor plan. The number of candidate points N_C has a direct impact on the accuracy of the found estimate (9) and will be studied in Sect. 6.3.

4 IMPLEMENTATION ON OFF-THE-SHELF DEVICES

We implement SALMA on off-the-shelf UWB devices. After introducing the hardware in Sect. 4.1, we sketch the scheme used to derive the distance between the tag and anchor as well as the CIR \mathbf{r}_m in Sect. 4.2. We then illustrate how the system can support multiple tags in Sect. 4.3 and describe the implementation of the position estimation in Sect. 4.4.

4.1 Hardware

The system consists of Decawave EVB1000 platforms used for both anchor and tags (Fig. 3a). These platforms employ the low-cost IEEE 802.15.4-compliant UWB transceiver DW1000 [5]. The tags are battery-powered and can be moved around freely. The anchor, instead, is located at a fixed position \mathbf{a} and is connected to a notebook running MATLAB. The antenna at the tag is a self-made linearly polarized omni-directional dipole antenna (Fig. 3b), but any off-the-shelf omni-directional UWB antenna is suitable. At the anchor, instead, we employ either a single omni-directional antenna (Sect. 4.4) or multiple directional antennas (Sect. 5.2).




Figure 3: Decawave EVB1000 node (a) with self-made omnidirectional dipole antenna (b) and switchable directional antenna system employed in Sect. 5.2 (c).

Transmitted pulse shape. The proposed signal model in (1) requires a known transmitted pulse shape $s_{\text{DW}}(t)$. The IEEE 802.15.4-2015 standard allows the generation of an arbitrary pulse shape, as long as it fulfills certain requirements on its cross-correlation with a standard reference signal, a root raised cosine pulse with a roll-off factor of $\beta = 0.5$ [50]. Decawave follows the IEEE 802.15.4-2015 standard, but does not provide information regarding the transmitted signal of the DW1000. Therefore, we identify $s_{\text{DW}}(t)$ in a measurement campaign. We place a transmitter and receiver 1 m apart from each other in clear LOS conditions. The receiver logs 1000 CIRs. In a post-processing step, we separate the LOS from the CIR and calculate an average over these signals, which defines the transmitted pulse shape $s_{\text{DW}}(t)$ of the DW1000.

4.2 Acquiring CIR and ranging

As shown in Fig. 1a, the localization engine of SALMA requires to estimate the distance \hat{d}_0 between the tag and the anchor, and to derive information about the multipath propagation by acquiring the CIR provided by the DW1000.

Two-way ranging. Due to the missing synchronization between anchor and tags, we employ a double-sided two-way ranging scheme (DS-TWR) to estimate the distance $\hat{d}_0 = \|\mathbf{p} - \mathbf{a}\|$ between each tag and the anchor (see Fig. 4). The DS-TWR scheme consists of three messages, each of which contains an 11-byte MAC header embedding source and destination address, as well as a 16-bit checksum. The payload of the first message (*INIT*) and the second message (*RESP*) is 1 byte long (MSG_ID). The last message (*FINAL*) is 16 bytes long and contains the message ID as well as three 5-byte timestamps [13]. The uncalibrated distance \hat{d}_{TWR} is calculated in the DS-TWR scheme with [6, pp. 213]:

$$\hat{d}_{\text{TWR}} = \frac{T_{\text{round1}} \cdot T_{\text{round2}} - T_{\text{reply1}} \cdot T_{\text{reply2}}}{T_{\text{round1}} + T_{\text{round2}} + T_{\text{reply1}} + T_{\text{reply2}}} \quad (10)$$

To calibrate the distance estimate, we perform 5000 DS-TWR trials between the anchor and a tag placed 2 m apart from each other. The derived variance and mean of the difference between the reported distance \hat{d}_{TWR} and the true distance $d_0 = 2$ m is $\sigma_{\text{DW}}^2 = (0.054 \text{ m})^2$ and $\mu = 0.26$ m, respectively. Hence, the calibrated distance estimate follows as $\hat{d}_0 = \hat{d}_{\text{TWR}} - \mu$. The distance \hat{d}_0 and the variance σ_{DW}^2 define the distribution of the candidate points around the anchor, as shown in Sect. 3.3.

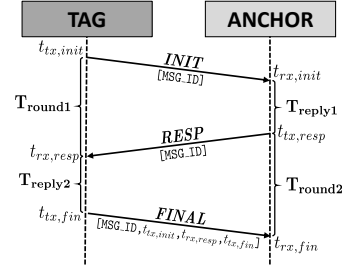


Figure 4: Double-sided two-way ranging scheme.

Acquisition of CIR. Besides deriving the distance \hat{d}_0 between anchor and tag, the anchor acquires the CIR r_m from the *FINAL* message received from the tag. Fig. 1b illustrates an exemplary r_m . The sampling period is set to $T_s = 1/f_s = 1/(2 \cdot 499.2 \text{ MHz}) = 1.0016$ ns. Each sample consists of a 16-bit real integer and a 16-bit imaginary integer resulting in a total size of 4048 Bytes. To reduce the amount of data read via SPI from the DW1000, we limit the length of the CIR to $N_s = 100$ samples.

Having the tags initiate the DS-TWR lets the anchor receive the required information to run the localization algorithm (i.e., the *INIT* and *FINAL* message). At the same time, it also allows to shift the burden on the anchor, which is typically static and much more powerful than tags, as it is line-powered and connected to a backbone localization engine that performs the CPU-intensive calculations. This is advantageous in real-world deployments, as tags are able to control the position update rate based on their energy budget. For example, by equipping a tag with an accelerometer, one can initiate a position update only in case of a movement, and remain in low-power mode otherwise.

4.3 Supporting multiple tags

We have so far considered only a single tag placed at an unknown position \mathbf{p} . SALMA can support up to N_t tags placed at positions \mathbf{p}_i (with $i = 1 \dots N_t$) by employing a slotted ALOHA scheme. The duration of a time-slot is related to the computation time necessary to obtain a position estimate (evaluated experimentally in Sect. 6.4) plus a guard interval of 1 ms at the start/end of each time-slot to overcome mis-alignments due to clock drifts.

In our current implementation, the anchor periodically broadcasts beacon messages embedding information about the time-slots' occupancy every 30 seconds³. Tags are not assigned to specific time-slots, but have instead the freedom to use any of the unoccupied ones: this enables a tag to use several time-slots in case it requires a higher update rate. In principle, this scheme may lead to collisions if two tags pick the same time-slot. This is, however, a well-known issue that has been largely studied in RFID systems where the reader sends a request and tags pick a random slot to answer [8, 42]. Existing anti-collision schemes can be readily applied also in SALMA. For example, the anchor can monitor the number of occupied slots and adjust their number accordingly, or adapt the slot duration by changing the number of candidate points.

³Due to the stable clock of the EVB1000 board (10 ppm), an even higher interval between beacon messages can be safely selected.

4.4 SALMA-light: Position estimation using omni-directional antennas

After the anchor has acquired the estimated distance \hat{d}_0 and the CIR \mathbf{r}_m , SALMA needs to carry out the position estimation as described in Sect. 3.3. We provide a first implementation of such a position estimation by equipping the anchor node with a single omni-directional antenna: we call this implementation *SALMA-light*. When using a single antenna, only a single CIR observation is available, which greatly simplifies the signal model from (3) with $M = 1$.

Obtaining a position estimate. We use \hat{d}_0 to obtain candidate points, as described in Sect. 3.3. At each candidate point, only the MPC delays τ_k are calculated using (4), since the beampattern has no effect on the estimate. The amplitude estimate from (8) requires a computationally demanding matrix inversion, and, in the case of overlapping MPCs, the matrix might not even be invertible. Hence, we approximate the log-likelihood value from (9) iteratively [26]:

$$\text{init : } \mathbf{r}^{(0)} = \mathbf{r} \quad (11)$$

$$\text{for } k = 1 \dots K : \alpha_k = \mathbf{s}^H(\tau_k) \mathbf{r}^{(k-1)} \quad (11)$$

$$\mathbf{r}^{(k)} = \mathbf{r}^{(k-1)} - \alpha_k \mathbf{s}(\tau_k) \quad (12)$$

Essentially, we take the observed CIR \mathbf{r} and sequentially subtract sub-hypotheses ($\alpha_k \mathbf{s}(\tau_k)$) by using pulses shifted to the respective τ_k and weighted by single amplitude estimates α_k . The resulting $\mathbf{r}^{(k)}$ is then the left-over 'residual' signal. The latter represents how similar the hypothesized and measured CIRs are, and is thus used as an approximation of the log-likelihood. This procedure is repeated for each candidate point and the one with highest log-likelihood value is chosen to be the tag position estimate $\hat{\mathbf{p}}$.

Limitation: multipath ambiguities. While this method is simple, the non-accessible beampattern restricts the algorithm to delay information only. This restriction makes the algorithm sensitive to overlapping MPCs, as well as to ambiguities in the delay times of MPCs, which may degrade the positioning performance significantly, as shown in Sect. 6.2.

5 TACKLING MULTIPATH AMBIGUITIES

As discussed in Sect. 4.4, SALMA-light uses measurements from a single antenna only, which makes the algorithm sensitive to overlapping MPCs and ambiguities. In this section, we introduce *SALMA-full*: an enhanced version of the system in which the anchor makes use of multiple switchable directional antennas⁴. Hence, we may now take advantage of the full signal model from (3), where each antenna m is characterized by its beampattern $b_m(\phi)$ covering one sector of the azimuth plane.

The combined observations of the antennas enable the system to separate closely-arriving MPCs in the spatial domain. However, the combination of several antenna measurements requires phase-coherency between the measurements, which is not given by low-cost transceivers. In the following, we tackle the phase-coherency issue (Sect. 5.1), describe how to carry out position estimation using directional antennas (Sect. 5.2), and highlight the key differences in the employed hardware compared to SALMA-light (Sect. 5.3).

5.1 Non-phase-coherent amplitude estimates

Phase-coherency demands accurate radio clocks, which are not provided by off-the-shelf UWB transceivers like the DW1000. In our case, this affects the implementation of the presented amplitude estimates in (8). Inaccurate clocks between consecutive measurements are perceived as a phase change in the baseband-equivalent CIR. Thus, amplitude estimates from consecutive measurements differ in their complex-valued phase $\alpha_{k,m} \approx e^{j\varphi} \alpha_{k,m'}$, where φ denotes the unknown phase offset. However, the unknown phase offset φ is required for the position estimate in (9).

To overcome the necessity of phase coherency, we follow the approach presented in [27]. Assuming non-overlapping MPCs ($\mathbf{s}(\tau_k)^H \mathbf{s}(\tau_k) \approx 0$), an MPC amplitude α_k can be estimated independently as projection of the normalized signal $\frac{\mathbf{s}^H(\tau_k)}{\mathbf{s}^H(\tau_k) \mathbf{s}(\tau_k)}$ onto the m -th measurement \mathbf{r}_m according to

$$\alpha_{k,m} = \frac{\mathbf{s}^H(\tau_k) \mathbf{r}_m}{\mathbf{s}^H(\tau_k) \mathbf{s}(\tau_k)}. \quad (13)$$

Furthermore, the amplitude estimate in (8) can be written as complex-valued average. Relaxing the complex-valued weighted average by an absolute-valued average [41] results in an estimate of the k -th MPC amplitude α_k^{avg} according to

$$\alpha_k^{\text{avg}} = \frac{\sum_{m=1}^M |\alpha_{k,m}| \cdot |b_m(\phi_k)|^2}{\sum_{m'=1}^M |b_{m'}(\phi_k)|^2}. \quad (14)$$

The remaining phase $\angle \alpha_{k,m}$ is extracted from the individual antenna measurements and the amplitude estimate $\hat{\alpha}_{k,m}^{\text{avg}}$ of the m -th antenna and k -th MPC results in

$$\hat{\alpha}_{k,m}^{\text{avg}} = \alpha_k^{\text{avg}} \exp(j \angle \alpha_{k,m}). \quad (15)$$

This approximation combines MPC amplitudes from non-phase-coherent measurements, taking into account the directivity of the M antennas.

5.2 SALMA-full: position estimation using directional antennas

In contrast to SALMA-light, SALMA-full collects observations from M directional antennas. The antennas are physically separated (see Fig. 3c), which results in different range estimates from tag to each antenna. Since this difference is smaller than the standard deviation of the DW1000 ranging, this error can be neglected. However, to create the candidate points as described in Sect. 3.3, we use the mean value of all ranges.

For each candidate point, the MPC parameters τ and ϕ are calculated using the VA positions in (4) and (5), respectively. For the amplitude estimates, the same iterative approach is followed as in Sect. 4.4, but it is adapted to use the stacked observed CIRs $\mathbf{r} = [\mathbf{r}_1^T, \dots, \mathbf{r}_M^T]^T$ and to take the non-coherent amplitude estimates from the previous section into account. For this, in the iteration step (11), we use α_k^{avg} from (14), and for step (12), we use

$$\mathbf{r}_m^{(k)} = \mathbf{r}_m^{(k-1)} - b_m(\phi_k) \hat{\alpha}_{k,m}^{\text{avg}} \mathbf{s}(\tau_k).$$

⁴This system was showcased at SenSys'17 [11].

This gives us a (stacked) residual $\mathbf{r}^{(K)} = [r_1^{(K)T}, \dots, r_M^{(K)T}]^T$ representing the similarity between the hypothesized and measured CIR that is used as an approximation for the log-likelihood. The final estimate $\hat{\mathbf{p}}$ is obtained by evaluating the log-likelihood for each candidate point and by picking the one achieving the maximum.

Two aspects are worth of note regarding the beampatterns $b_m(\phi)$ of the antennas: first, we use 36 sampled values $b_m(i \cdot \Phi_s)$ with a spacing of $\Phi_s = 10^\circ$ obtained from a measurement campaign. Second, when the anchor is employed, it can be oriented with $\Phi_o = j \cdot 10^\circ$, where j might be chosen as desired. For the implementation, this value has to be known. SALMA then uses $b_m(\lfloor \frac{\phi_k + \Phi_o + 5}{10} \rfloor \bmod 36)$ to approximate $b_m(\phi_k)$, also taking the orientation into account.

5.3 Hardware differences

In contrast to SALMA-light, in SALMA-full we exploit four self-made and low-cost directional antennas with a half-power beamwidth of about 150° (see Fig. 3c). The antennas are mounted such that each one points in a different cardinal direction. The evaluation in Sect. 6.2 shows that, even with this wide beamwidth, SALMA achieves an error below 20 cm for 90% of the estimated positions. The higher number of antennas increases the acquisition time of CIRs and distance estimates. Hence, the acquisition duration is higher than that of SALMA-light, as discussed in Sect. 6.4.

6 EVALUATION

We evaluate the positioning capabilities of SALMA in challenging indoor environments: an office (Room A, see Fig. 5a and 6a), and a stockroom (Room B, see Fig. 5b and 6b). After describing the experimental setup in Sect. 6.1, we answer the following questions:

- What is the benefit of using SALMA-full over SALMA-light? (Sect. 6.2);
- What is the *accuracy* achieved by SALMA, and at which computational costs? (Sect. 6.3);
- How long does it take to estimate a position, and what are the implications on *scalability*? (Sect. 6.4);
- Do more (and better) antennas improve the performance of SALMA? (Sect. 6.5).

We answer all these questions in Room A under clear LOS conditions. In Sect. 7, we will then specifically evaluate how SALMA performs in more challenging environments with obstructed LOS (both rooms), and a dynamic environment due to moving objects and people (Room A).

6.1 Experimental setup

We carry out the evaluation in an office containing obstacles and scattering objects such as desks, chairs, shelves, and PC monitors, as shown in Fig. 5a and 6a (Room A). We place the tag in $N_{\text{EP}} = 35$ evenly distributed evaluation points, while fixing the anchor next to the table. We mount both anchor and tag on a tripod at a height of 1.50 m, i.e., well above the obstacles, so to have clear LOS conditions. The anchor is connected to a Lenovo ThinkPad T450s notebook running MATLAB. The tag, instead, is battery-powered and can move freely. The only pre-processing required by SALMA is to enter the anchor location and orientation as well as the coordinates of the surrounding four wall segments. The following settings are used

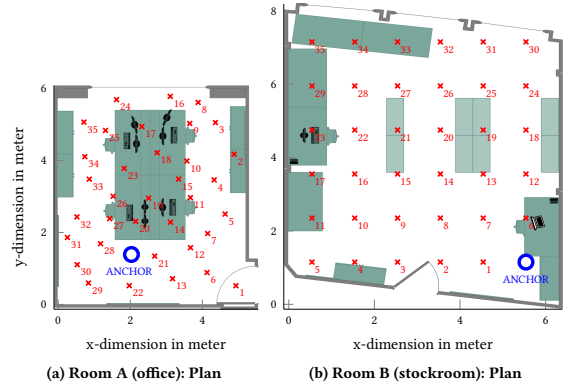


Figure 5: Evaluation setup (2D-plan): we consider 35 evaluation points (red crosses) in two different environments.



(a) Room A (office): Picture



(b) Room B (stockroom): Picture

Figure 6: Evaluation setup (Picture): the white dashed line marks the measurement height under obstructed LOS.

by the DW1000: maximum data rate (6.8 Mbps), pulse repetition frequency of 64 MHz, and a preamble symbol repetition of 1024. Channel 7 is used due to its high bandwidth (900 MHz) and since our self-made directional antennas are optimized for this band. At each evaluation point (denoted by \mathbf{p}_{EP}), we perform 100 position estimates, hence carrying out 3500 evaluations in total for both SALMA-light and SALMA-full. We denote the i -th position estimate by $\hat{\mathbf{p}}_i$, and obtain the absolute position error with

$$\text{Err}_i = \|\hat{\mathbf{p}}_i - \mathbf{p}_{\text{EP}}\|. \quad (16)$$

Statistically, we look at the cumulative distribution function (CDF) over the errors Err_i using all evaluation points (i.e., 3500 estimates) for different configurations, as illustrated in the upcoming sections.

6.2 SALMA-light vs. SALMA-full

In this section we examine the performance of both SALMA implementations comparatively.

Handling ambiguities. SALMA-light relies solely on the position information contained in the arrival times of MPCs. Hence, the

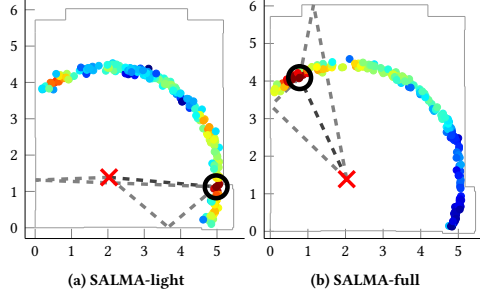


Figure 7: Evidence of multipath ambiguities.

resulting likelihood for the positions is highly multimodal, or in other words, there are multiple regions that seem to best fit the observed signal. This is demonstrated in Fig. 7, which shows the positioning result using SALMA-light (Fig. 7a) and SALMA-full (Fig. 7b) for one estimation run on position 34 (cf. Fig. 5a). The colored dots indicate the candidate point positions, where the color represents the likelihood values (red=high, blue=low). SALMA-light has three regions showing similarly high likelihood values (red and orange dots), caused by similarly long MPC paths, which results in a completely wrong estimate. In contrast, SALMA-full can narrow the estimate down to find the true position of the tag. This is possible due to the combined directional antenna observations, where wrong candidate points have low likelihood values because amplitude values do not fit to the antenna patterns.

Quantitative comparison. We show the improvement quantitatively by accounting for all 3500 estimates via the CDF of the absolute position error. Fig. 8 shows the CDF for SALMA-light (dashed blue line) and SALMA-full (solid orange line). With SALMA-light, 67.3% of all evaluations have a position error below 25 cm. On the other hand, 21.7% of the evaluations have an error above 1 m: these outliers are caused by the multipath ambiguities, as just explained. By using directional antennas, SALMA-full can mitigate these outliers: 90% of all evaluations are below 20.17 cm, whilst 99% of the evaluations are below 29.72 cm.

We can hence conclude that SALMA-full clearly outperforms SALMA-light thanks to the additional angular information. Hence, we focus the next evaluations on SALMA-full only.

6.3 Localization accuracy

With the promising results shown in Sect. 6.2, we investigate the accuracy of SALMA in more detail, and focus also on the computational costs.

Role of candidate points. We examine the impact of the number of candidate points used by SALMA-full. To this end, we perform 3500 estimates for different number of candidate points $N_C \in \{50, 100, 200, 500, 1000\}$. Fig. 9 shows the resulting CDF: even when using only 50 candidate points (blue, star), 90% of the estimates have an error below 30 cm. However, there are outliers for about 5% of the estimates. Increasing the number of candidate points removes the outliers and improves the performance to a “saturation point” at about 200 candidate points (i.e., a higher number of points gives negligible improvements). Hence, we make $N_C=200$




Figure 8: SALMA-light sustains an accuracy below 30 cm only in 70% of the cases due to multipath ambiguities. By exploiting the angular domain, SALMA-full exhibits an error below 30 cm in 99% of the cases (Room A).




Figure 9: Impact of the number of candidate points on SALMA’s accuracy: $N_C=200$ acts a good trade-off.

our preferred setup and use it for all further evaluations, unless stated otherwise. The number of candidate points increases the computation time linearly, so N_C can act as a trade-off between computational costs and accuracy, as illustrated in Sect. 6.4.

Individual evaluation points. A more detailed display of the accuracy is shown in Fig. 10a. For the 100 estimates at each position of Room A, the mean (blue circle) and the 3-fold standard deviation (black error ellipse) are shown. The former indicates that there is little estimation bias (distance to ground truth). As for the standard deviation, with the good ranging precision of UWB, the ranging deviation is small (facing the LOS), while info gained by the MPCs determines the angle deviation (perpendicular to the LOS). Overall, the good performance is reinforced, while there are certain positions (e.g., 9, 16, and 18) with a slightly higher bias. Of special note is that the accuracy of SALMA does not degrade at higher ranges: pos. 8, 16, and 24 are placed more than 4 meters apart from the anchor, but their estimates are as accurate as the ones obtained at positions much closer to the anchor. This is highlighted in Fig. 11 showing the average position error with respect to the real distance between anchor and tag. This stands in contrast to many other indoor positioning techniques, where the inaccuracy increases quickly with the range, e.g., visual systems [34].

6.4 Scalability

The number of supported tags by SALMA is limited by (i) the computation time of the position estimation, (ii) the duration of the DS-TWR, as well as (iii) the time needed to stream the CIR and additional info to the notebook via USB. The duration of a DS-TWR is mainly defined by the packet length of its three packets, which is 3.49 ms. Streaming one CIR to MATLAB takes 4.62 ms.

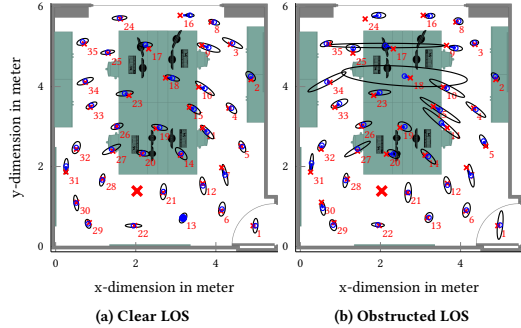


Figure 10: Error ellipses showing position bias and three-fold standard deviation (Room A).

The computation time of the position estimation depends on the notebook's performance and on the number of candidate points. In the evaluations, the algorithms are performed in MATLAB on a Lenovo ThinkPad T450s with 2.59 GHz clock and 8 GB RAM. An increase in the number of candidate points scales the computation time linearly. Thus, we evaluate the time needed per candidate point. SALMA-light takes $174.77 \pm 12.2 \mu\text{s}$ and SALMA-full requires $955.13 \pm 23.5 \mu\text{s}$ per candidate point. Thus, with $N_C = 200$ candidate points, the algorithms take 34.95 ms and 191.03 ms, respectively. The overall duration of a position estimation with SALMA-light when using $N_C = 200$ is hence 43.06 ms, resulting in an update rate of 23 Hz. Using $N_C = 50$, instead, gives an update rate of 60 Hz.

When using SALMA-full, for each antenna a DS-TWR trial is performed and four CIRs are acquired: this reduces the achievable update rate. In total, SALMA-full requires 223.5 ms for $N_C = 200$ and 79.8 ms for $N_C = 50$, resulting in an update rate of about 4.5 Hz and 12.5 Hz, respectively. Thus, SALMA can easily compete with comparable solutions and outdoor positioning systems like GPS.

6.5 The role of the antenna

We examine next how SALMA would perform when using more antennas with more directive beampatterns. As such antennas are not yet commercially available, we simulate artificial CIRs (r) and ranging (\hat{d}_0) for the same tag positions shown in Fig. 5a. For any tag position \mathbf{p} , we can create an artificial CIR in two steps:

- (1) Specular part: we shift and add pulses $s_{\text{DW}}(t - \tau_k)$ weighted by beampattern $b(\phi_k)$ and amplitude α_k using the known delays τ_k and angles ϕ_k . The amplitude exhibits free-space path-loss and each reflection halves the magnitude. We consider MPCs up to order two.
- (2) Scattering part: we simulate diffuse multipath by drawing realizations of a Gaussian random process whose variance is defined by a double exponential power delay profile according to equation (9) from [22]. Additionally we simulate AWGN measurement noise with an SNR of 29.5 dB at 1 m.

This simulation setup allows us to adjust the half-power beamwidth (HPBW) of the antennas and to recreate the effect of clutter by setting a signal-to-interference⁵ ratio (SIR). The latter is defined by the ratio between LOS and scattering energy. We determined empirically that an SIR of 3 dB properly describes the environment.

⁵Interference, in this case, refers to self-interference due the scattering part.




Figure 11: Average position error w.r.t. the real distance.




Figure 12: Simulated performance of SALMA when using multiple antennas with more directive beampatterns.




Figure 13: Performance of SALMA in clear LOS, obstructed MPC, and obstructed LOS situations in Room A.

To recreate the performance behavior from the SALMA-full measurement runs described in Sect. 6.3, we set the HPBW to 150° , matching the properties of the used antennas. Additionally, we carry out simulations using six antennas with a HPBW of 90° , reflecting a higher quality implementation. Fig. 12 shows the results. On the one hand, we can see that the SALMA-full simulation (red curve) fits the measured results (purple curve) closely. A slightly better performance is achieved in the simulation, because the impact of bias due to floor plan inaccuracies is not present. On the other hand, we can notice that, when simulating six antennas (blue curve), the performance of SALMA improves significantly: the 90% error decreases by about 10 cm to almost reach the sub-decimeter mark, whilst 99% of the estimates achieve an error below 20 cm.

7 ROBUSTNESS TO NON-LINE-OF-SIGHT AND DYNAMIC ENVIRONMENTS

Indoor environments are inherently highly dynamic due to moving humans and objects. Thus, the value of a localization system strongly depends on (i) its performance under obstructed LOS, (ii) its behavior in different environments and (iii) its robustness in crowded settings. In this section, we discuss the performance of SALMA under non-line-of-sight (NLOS) conditions (Sect. 7.1), when furniture is moved without updating the map (Sect. 7.2) and in the case of a highly-dynamic and crowded environment (Sect. 7.3).

7.1 Performance under NLOS conditions

In situations of a blocked LOS, range-based systems suffer from a positive bias [36]. This is either caused by the lower propagation speed in case the signal propagates through the obstacle, or, in case of a fully blocked LOS, due to the misinterpretation of a reflection as the direct path.

Distance bias. SALMA requires the distance estimate \hat{d}_0 between anchor and tag to distribute the candidate points on a circle around the anchor as described in Sect. 3.3. Thus, an obstructed LOS causes an increase in the radius of the circle. First, we analyze the impact of different objects on \hat{d}_0 . We place tag and anchor 2 m apart from each other and perform 1000 DS-TWR trials with different objects blocking the LOS. The objects included: a metal plate (800×450×3 mm), PC monitors, and humans. The threshold-based mechanism of the DW1000 was able to detect a leading edge corresponding to the LOS in each of the trials. But, indeed, the obstructed LOS leads to a positive bias in the range estimate. Metal plate and PC monitor caused a range bias of 11 cm and 13.1 cm, respectively. Even in the case of humans blocking the LOS, the leading edge was successfully detected, but two humans blocking the LOS already led to a range bias of 41.8 cm. Further evaluations will show that SALMA is robust even in the case of a range bias.

Accuracy evaluation. SALMA is not just making use of the LOS component, but also of specular MPCs. Thus, we evaluate SALMA also in situations of blocked MPCs. We repeat the evaluation in Room A described in Sect. 6.1: this time, however, we mount the tag and the anchor at a height of 1.20 m, corresponding to the height of monitors, shelves, and people in the room. Depending on the position of the evaluation points, this results in obstructed LOS for twelve of these points $EP_{OLOS} = \{3, 8, 9, 10, 15, 16, 18, 19, 25, 26, 33, 34\}$, which results in 1200 evaluations. Note that, for all these evaluation points, there were also specular MPCs blocked by objects. For twenty points, instead, the LOS was still clear but specular MPCs were blocked $EP_{OMPC} = \{1, 2, 4 - 7, 11 - 14, 17, 20 - 24, 27, 28, 32, 35\}$. In total, 2000 evaluations were acquired in these situations. The remaining 300 evaluations are still in clear LOS with no blocked MPCs, thus, they are ignored. Fig. 13 (magenta dash dotted line) shows the CDF of all evaluations under blocked LOS (EP_{OLOS}). The median is at 14.5 cm and the error for 90% of the estimates is still below 30.7 cm. The blue dashed line in Fig. 13 shows the CDF just considering evaluations where significant multipath components are blocked by obstacles or humans (EP_{OMPC}). The median is at 10.25 cm and the error for 90% of the estimates is below 30.52 cm. This shows that SALMA remains robust even in the case of blocked MPCs.

Qualitative evaluation. In Fig. 10b, we see again the accuracy for individual evaluation points, now for the obstructed LOS case. While the position bias (distance blue circles to red crosses) did not increase significantly, we can see that the variance in the angular direction increases for most of the evaluation points. The error ellipses shown in Fig. 10b indicate two evaluation points with significantly higher variances (no. 17 and 18) as the other points. The reason is the unfortunate position of the anchors in this case. The PC monitors and obstacles at the left and right wall block the respective MPCs, thus, the position information obtained at these positions comes only from the LOS and the reflection from the

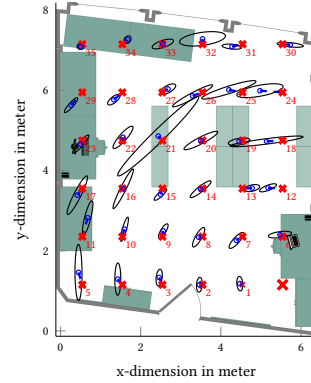


Figure 14: Error ellipses showing position bias and three-fold standard deviation with empty racks in Room B.




Figure 15: Performance of SALMA in Room B in the case of moving obstacles (storage racks are empty or filled).

window. Since these two reflections are arriving at the same angle, SALMA suffers from a poor geometric configuration. This results in ambiguities similar to the ones shown in Sect. 6.2. Due to significant MPCs from the left and right wall, this situation was not evident in the clear LOS case (see Fig. 10a).

7.2 Performance in stockroom with moving obstacles

In Sect. 6 and 7.1, we have performed all the measurements in Room A. To prove the capabilities of SALMA also in more challenging environments and in the presence of moving obstacles, we have evaluated its performance also in Room B (see Fig. 5b and 6b).

Performance in more challenging environments. To challenge SALMA, we chose a stockroom that is larger than Room A (46.7 m² vs. 31.6 m²) and cluttered with desks, storage racks (bright rectangles in Fig. 5b) and several other metal objects (see Fig. 6b). We have mounted anchor and tag at a height of 1.20 m. Fig. 15 (solid orange line) shows the CDF of all evaluations in Room B. The median is at 18.6 cm and 90% of all estimates obtain an error below 44.5 cm. Thus, compared to the evaluation in Room A, the performance of SALMA is slightly worse due to the larger room with more clutter and wall materials with unfavorable reflective properties (see Sect. 8). Fig. 14 shows the accuracy for individual evaluation points. Similar to Room A (see Sect. 7.1), some positions




Figure 16: Snippet of a 24-hours experiment in dynamic environments. The dashed blue line depicts the mean error of 50 position estimates over time, whilst the solid orange line shows the 90% error. Despite the people moving in/out of the room (green line), SALMA can sustain a decimeter-level position accuracy.

(e.g., 4, 5, 12, 18, 24) suffer from an unfavorable anchor placement as the LOS is arriving from the same angle as strong reflections.

Moving obstacles. The performance of localization systems based on RSS profiling and fingerprinting is highly affected by moving obstacles. Thus, changing the furniture in a room often requires to update or repeat measurements. To evaluate the performance of SALMA in the case of moving obstacles or furniture, we have stocked up the storage racks in Room B with full beer crates and other objects (see Fig. 6b). As the goal of SALMA is to minimize the setup effort, we do not model reflections from obstacles such as the full storage racks. Fig. 15 shows that the position error (dashed blue line), while higher due to the range bias introduced by obstructed LOS, still stays in reasonable bounds, relatively unaffected by the additional reflections. Thus, SALMA only slightly loses accuracy in favor of practicability and setup time.

7.3 Performance in a crowded environment

For a final stress test of SALMA and to evaluate its behavior in a dynamic environment including NLOS situations, we employed our system again in the office scenario (Room A) for a non-stop 24 hours run. The system was exposed to the usual ongoing work flow that involves multiple people passing by the system, thereby blocking the LOS or MPCs, hence creating a dynamic environment. During the 24 h experiment, SALMA localized three tags at representative positions (positions 3, 15, and 33) simultaneously. The positions were deliberately chosen to be under obstructed LOS. Every five seconds we estimated the tag positions resulting in 51840 position estimates. We evaluate the performance of the system on multiple levels. Fig. 16 shows the mean error (dashed blue line) and the 90% error (solid orange line) over 50 position updates from 07:00 - 21:00 o'clock. Additionally, we track the number of present people in the room during the experiment (green staircase graph). The figure focuses on daytime, since over night no one was in the room and the performance remained constant. It can be seen that the usual working environment (with the two designated working people present) does not impair the performance of the system providing an average error below 11.2 cm. In terms of present people, there are two events prominent in Fig. 16: at 11:00 o'clock there was a meeting with five people and at 14:00 o'clock we have presented SALMA to thirteen people making it in total fifteen people in the room simultaneously. We asked people to move around the room freely during the presentation, thus, the LOS and the MPCs were

obstructed in a dynamic fashion. Even though the error increases during these periods, still, when the room was completely filled with people, the average error was below 34 cm and the 90% error below 79 cm. The latter indicates that SALMA is robust also in a highly dynamic environment and under NLOS conditions.

Comparison to other multi-anchor systems. Comparing the accuracy of SALMA with other UWB-based systems is difficult, as they are either evaluated in mobile 2D [15, 31, 52] or static 3D [23, 24] scenarios. Silva et al. [49] report a 2D static LOS mean error of 16.6 cm. SALMA instead achieves an average error of just 9.85 cm in clear LOS. Kempke et al. [23, 24] report a 90% error of 77 cm and 50 cm in static 3D, respectively. In contrast, SALMA achieves a 90% error of 50 cm between 14:00-14:30 o'clock, thus, under obstructed LOS and when up to fifteen people were walking around. Therefore, it is fair to say that SALMA can compete and even outperform existing systems, despite using just a single anchor.

8 DISCUSSION

Our evaluation demonstrates the capabilities of SALMA to perform accurate positioning in typical indoor environments. However, it has also highlighted a number of challenges and open questions that we will elaborate in detail in this section.

Sensitivity to chosen anchor position. Due to SALMA's principle, just one anchor per room is required. In our evaluations, we examined two typical choices for anchor positions, namely, in the vicinity of the room center (Room A) and in the corner of the room (Room B). Both variants have pros and cons: in Room A we have a full candidate point circle for many ranges, which increases the risk of ambiguities, especially for SALMA-light. However, SALMA-full can take full advantage of the beampatterns in all directions to stay relatively unaffected (as we have demonstrated in Sec. 6.2). In Room B, we have, at most, a quarter circle of candidate points: on the one hand, this reduces possible ambiguities. However, on the other hand, this results in higher ranges with reduced signal strength and reduced benefit from the angular information. We also pointed out some difficult positions in both rooms where LOS and the strongest reflection come from the same direction, resulting in a plateau in the likelihood which leads to a dilution of precision. These cases exist no matter what anchor position is chosen.

The anchor orientation can be set arbitrarily, but it has to be fixed and known to correctly weight the amplitudes.

Including the third dimension. SALMA is designed specifically to perform 2D positioning. This choice is rather pragmatic: physical and algorithmic setup of the system are simplified dramatically, enabling a practical implementation with short setup time and efforts, while only using a single anchor. Also, many applications (e.g., navigation tasks) do not require any height information. In principle, the methods can be extended to the third dimension: (i) determining the VAs can be done by mirroring at plane surfaces, (ii) for the MPC angle one needs to take the elevation beampatterns into account and (iii) the candidate points are placed on a range sphere rather than circle. However, this drastically increases the computational complexity and makes the position likelihood even more multimodal. A 3D model of the environment could help to avoid ambiguities due to floor or ceiling reflections, however, our antennas exhibit a fairly narrow elevation pattern, hence, the impact of ceiling, floor, and other reflections is limited significantly.

Effect of wall materials. The main setup effort for SALMA is the determination of reflecting surfaces in the considered environment. However, additional care has to be taken with regard to the material of the surfaces. Preferably, materials such as glass and metal enable good reflectors and including them in the models enhances the position estimate. On the other hand, plaster boards or wooden surfaces, even if they are flat and smooth, give little to no contribution in terms of specular reflections and can in fact decrease the performance. For example, in Room B, the eastern wall, even though close to the anchor, is made out of plasterboard and does not contribute with a specular multipath component. Thus, it should not be included in the signal model as a source of a virtual anchor.

9 RELATED WORK

Indoor localization technologies. Many RF technologies have been investigated for indoor localization, such as Wi-Fi [9], Bluetooth [1, 3], and IEEE 802.15.4 [29, 44]. However, these systems hardly achieve an accuracy below 1 m, require a high amount of reference nodes, and typically come with a high deployment effort. Optical systems are among the most accurate indoor localization systems, but cannot inherently operate in NLOS conditions [34]. SALMA, instead, reaches a median error of 15 cm and a 90% error of 30 cm even in obstructed LOS conditions. Acoustic systems can also achieve decimeter-level accuracy, but their biggest enemy – multipath propagation – is SALMA’s best friend [30, 37].

UWB indoor localization systems. UWB-based systems can also achieve decimeter-level accuracy [35, 57]. Recently, several systems have been implemented using low-cost UWB transceivers [15, 23, 24, 31, 49, 52]. However, these systems require a high amount of anchors, typically between eight [31, 52] and fifteen [24]. SALMA, instead, uses a single anchor and – to the best of our knowledge – no comparable solution exists. In terms of accuracy, as discussed in Sect. 7.3, it is fair to say that SALMA can compete and even outperform existing systems, despite using just a single anchor.

Multipath-assisted localization systems. Theoretical works have discussed the performance bounds of multipath-assisted indoor localization via simulation [10, 17, 55] and using very expensive, bulky and wired-synchronized equipment [25, 32, 39]. Instead,

with SALMA, we are the first to enable the exploitation of multipath reflections for low-cost, low-power wireless localization systems.

Directional antennas to enable single-anchor systems. Several works have exploited electronically steerable or switchable antenna systems to enable single-anchor localization using narrow-band technologies [2, 48]. However, SALMA outperforms all these systems due to the exploitation of the position-related information provided by the MPCs. Sun et al. [51] presented a UWB-based system claimed to achieve decimeter-level accuracy. However, their measurement setup is vague and a thorough analysis of the system performance is missing. Quing et al. [46] and Zhang et al. [56] presented similar systems, but solely based on simulation.

Although not exploiting directional antennas, also Chronos [53] requires just a single access point to estimate the position of another device. In particular, Chronos uses an omni-directional antenna array and emulates a wideband radio on commodity Wi-Fi systems. Still, due to the position-related information provided by the MPCs, SALMA outperforms Chronos in terms of accuracy. Furthermore, by using the license-free ISM bands, Chronos interferes and is prone to the interference of other devices using the 2.4 GHz band.

10 CONCLUSIONS AND FUTURE WORK

In this paper, we present SALMA, a low-cost UWB-based indoor localization system that exploits multipath reflections to tear down the position estimation to a unique solution while only using a single anchor. Besides a crude floor plan and the position of the anchor, the system does not need any prior calibration or training phase. By using directional antennas, we increased the robustness of SALMA against overlapping MPCs. We extensively evaluated the performance of SALMA under LOS and NLOS conditions, as well as during a 24 h stress-test to challenge SALMA in dynamic settings. Under LOS, SALMA achieved a median error below 8 cm and an error below 20 cm for 90% of the position estimates. Even under obstructed LOS and in a highly dynamic environment SALMA sustains a high accuracy.

Our aim in this paper was to show the outstanding capabilities of SALMA without using a tracking filter and solely utilizing single-shot single-anchor measurements. In future work, we will combine SALMA with a particle filter and an inertial measurement unit to benefit from past position estimates. Moreover, we will perform an exhaustive evaluation of SALMA in mobile environments.

ACKNOWLEDGMENTS

This work has been performed within the LEAD project “Dependable Internet of Things in Adverse Environments” funded by Graz University of Technology. This work was also partially funded by the SCOTT project. SCOTT (<http://www.scott-project.eu>) has received funding from the Electronic Component Systems for European Leadership Joint Undertaking under grant agreement No 737422. This joint undertaking receives support from the European Unions Horizon 2020 research and innovation programme and Austria, Spain, Finland, Ireland, Sweden, Germany, Poland, Portugal, Netherlands, Belgium, Norway. SCOTT is also funded by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under the program “ICT of the Future” between May 2017 and April 2020. More information at <https://iktderzukunft.at/en/>.

REFERENCES

- [1] M.S. Bargh and R. de Groote. 2008. Indoor Localization based on Response Rate of Bluetooth Inquiries. In *Proc. of the 1st ACM Int. Workshop on Mobile Entity Localization and Tracking in GPS-less Environments (MELT)*.
- [2] A. Cidronali, S. Maddio, G. Giorgetti, and G. Manes. 2010. Analysis and Performance of a Smart Antenna for 2.45-GHz Single-Anchored Indoor Positioning. *IEEE Transactions on Microwave Theory and Techniques* 58, 1 (Jan. 2010).
- [3] G. Conte et al. 2014. BlueSentinel: a First Approach using iBeacon for an Energy Efficient Occupancy Detection System. In *Proc. of the 1st ACM Conf. on Embedded Systems for Energy-Efficient Buildings (BuildSys)*.
- [4] D. Dardari, P. Closas, and P. M. Djurić. 2015. Indoor Tracking: Theory, Methods, and Technologies. *IEEE Transactions on Vehicular Technology* 64, 4 (April 2015).
- [5] Decawave Ltd. 2016. *DW1000 Datasheet. Version 2.12*.
- [6] Decawave Ltd. 2016. *DW1000 User Manual. Version 2.10*.
- [7] D. Espes, A. Daher, Y. Autret, E. Radoi, and P. Le Parc. 2013. Ultra-Wideband Positioning for Assistance Robots for Elderly. In *Proc. of the 10th International Conference on Signal Processing, Pattern Recognition and Applications (SPPRA)*.
- [8] K. Finkenzerler. 2008. *RFID Handbuch – Grundlagen und praktische Anwendungen von Transpondern, kontaktlosen Chipkarten und NFC*. Hanser.
- [9] S. Gansemer et al. 2010. RSSI-based Euclidean Distance Algorithm for Indoor Positioning Adapted for the use in Dynamically Changing WLAN Environments and Multi-level Buildings. In *Proc. of the 2010 IPIN Conference*.
- [10] C. Gentner, T. Jost, W. Wang, S. Zhang, A. Dammann, and U. Fiebig. 2016. Multipath Assisted Positioning with Simultaneous Localization and Mapping. *IEEE Transactions on Wireless Communications* 15, 9 (Sept 2016), 6104–6117.
- [11] B. Großwindhager et al. 2017. Demo Abstract: UWB-based Single-anchor Low-cost Indoor Localization System. In *Proc. of the 15th ACM Int. Conf. on Embedded Networked Sensor Systems (SenSys), demo session*.
- [12] B. Großwindhager et al. 2018. Dataset: Single-Anchored Indoor Localization with Decawave DW1000 and Directional Antennas. In *Proc. of the 1st Workshop on Data Acquisition To Analysis (DATA'18)*.
- [13] B. Großwindhager et al. 2018. Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications. In *Proc. of the 19th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*.
- [14] K. Guo, Z. Qiu, C. Miao, A. H. Zaini, C.-L. Chen, W. Meng, and L. Xie. 2016. Ultra-Wideband-Based Localization for Quadcopter Navigation. *Unmanned Systems Journal* 4, 1 (Jan. 2016).
- [15] F. Hartmann et al. 2015. Design of an Embedded UWB Hardware Platform for Navigation in GPS Denied Environments. In *Proc. of the IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*.
- [16] S. He and S.-H. G. Chan. 2016. Wi-Fi Fingerprint-based Indoor Positioning: Recent Advances and Comparisons. *IEEE Comm. Surveys & Tutorials* 18, 1 (Aug. 2016).
- [17] B. Hu, Z. Shi, and Y. Wang. 2018. Single-Sensor Based Indoor Localisation by Exploiting Multipath Propagation. *Electronics Letters* 54, 3 (Feb. 2018), 179–181.
- [18] X. Huang, W. Zhu, and D. Lu. 2010. Underground Miners Localization System based on ZigBee and WebGIS. In *Proc. of the 18th Int. Conf. on Geoinformatics*.
- [19] J. T. Isaacs, D. J. Klein, and J. P. Hespanha. 2009. Optimal Sensor Placement for Time Difference of Arrival Localization. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC)*.
- [20] SM R. Islam et al. 2015. The Internet of Things for Health Care: a Comprehensive Survey. *IEEE Access* 3 (June 2015).
- [21] L. Jiang, L. N. Hoe, and L. L. Loon. 2010. Integrated UWB and GPS Location Sensing System in Hospital Environment. In *Proc. of the 5th Int. Conf. on Industrial Electronics and Applications (ICIEA)*.
- [22] J. Karedal, S. Wyne, P. Almers, F. Tufvesson, and A. F. Molisch. 2007. A Measurement-Based Statistical Model for Industrial Ultra-Wideband Channels. *IEEE Transactions on Wireless Communications* 6, 8 (Aug. 2007), 3028–3037.
- [23] B. Kempke et al. 2016. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proc. of the 14th Int. Conf. on Embedded Network Sensor Systems (SenSys)*.
- [24] B. Kempke, P. Pannuto, and P. Datta. 2015. PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization. In *Proc. of the 2nd Int. Workshop on Hot Topics in Wireless (HotWireless)*.
- [25] Y. Kuang, K. Åström, and F. Tufvesson. 2013. Single Antenna Anchor-Free UWB Positioning Based on Multipath Propagation. In *Proceedings of the International Conference on Communications (ICC)*.
- [26] J. Kulmer et al. 2017. Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning. In *Proc. of the 5th Int. Workshop on Advances in Network Localization and Navigation (ANLN)*.
- [27] J. Kulmer, S. Grebien, M. Rath, and K. Witralsal. 2018. On the Unimportance of Phase-Coherent Measurements for Beam-pattern-Assisted Positioning. In *Proc. of the Int. Conf. on Wireless Comm. and Networking (WCNC)*.
- [28] J. Kulmer, E. Leitinger, S. Grebien, and K. Witralsal. 2018. Anchorless Cooperative Tracking using Multipath Channel Information. *IEEE Transactions on Wireless Communications* PP, 99 (Jan. 2018).
- [29] J. Larranaga, L. Muguira, J. M. Lopez-Garde, and J. I. Vazquez. 2010. An Environment Adaptive ZigBee-based Indoor Positioning Algorithm. In *Proc. of the Int. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*.
- [30] P. Lazik and A. Rowe. 2012. Indoor Pseudo-Ranging of Mobile Devices using Ultrasonic Chirps. In *Proc. of the 10th Int. Conf. on Embedded Network Sensor Systems (SenSys)*.
- [31] A. Ledergerber, M. Hamer, and R. D'Andrea. 2015. A Robot Self-Localization System using One-Way Ultra-Wideband Communication. In *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*.
- [32] E. Leitinger, M. Fröhle, P. Meissner, and K. Witralsal. 2014. Multipath-Assisted Maximum-Likelihood Indoor Positioning using UWB Signals. In *Proc. of the Int. Conf. on Comm. Workshops (ICC)*.
- [33] R. Leser, A. Schleindlhuber, K. Lyons, and A. Baca. 2014. Accuracy of an UWB-based Position Tracking System used for Time-Motion Analyses in Game Sports. *European Journal of Sport Science* 14, 7 (Feb. 2014).
- [34] S. Liu and T. He. 2017. SmartLight: Light-weight 3D Indoor Localization Using a Single LED Lamp. (Nov. 2017).
- [35] D. Lymberopoulos and J. Liu. 2017. The Microsoft Indoor Localization Competition: Experiences and Lessons Learned. *IEEE Signal Processing Magazine* 34, 5 (Sept. 2017).
- [36] S. Marañón, W. M. Gifford, H. Wymeersch, and M. Z. Win. 2010. NLOS Identification and Mitigation for Localization Based on UWB Experimental Data. *IEEE Journal on Selected Areas in Communications* 28, 7 (Sept. 2010).
- [37] R. Mautz. 2009. The Challenges of Indoor Environments and Specification on some Alternative Positioning Systems. In *Proc. of the 6th Int. Workshop on Positioning, Navigation and Communication (WPNC)*.
- [38] R. Mautz. 2012. *Indoor positioning technologies*. Ph.D. Dissertation, ETH Zürich.
- [39] P. Meissner, C. Steiner, and K. Witralsal. 2010. UWB Positioning with Virtual Anchors and Floor Plan Information. In *Proc. of the 7th Int. Workshop on Positioning, Navigation and Communication (WPNC)*.
- [40] R. Mirza, A. Tehseen, and AV J. Kumar. 2012. An Indoor Navigation Approach to Aid the Physically Disabled People. In *Proc. of the 18th Int. Conf. on Computing, Electronics and Electrical Technologies (ICCEET)*.
- [41] P. Mowlaee et al. 2016. *Single Channel Phase-Aware Signal Processing in Speech Communication: Theory and Practice*. John Wiley & Sons.
- [42] V. Namboodiri et al. 2012. An Extensive Study of Slotted Aloha-based RFID Anti-collision Protocols. *Computer Communications* 35, 16 (Sept. 2012).
- [43] C. Nerguizian, C. Despins, and S. Affes. 2006. Geolocation in Mines with an Impulse Response Fingerprinting Technique and Neural Networks. *IEEE Transactions on Wireless Communications* 5, 3 (March 2006).
- [44] C. W. Ou et al. 2017. A ZigBee Position Technique for Indoor Localization Based on Proximity Learning. In *Proc. of the International Conference on Mechatronics and Automation (ICMA)*.
- [45] B. Perrat et al. 2015. Quality Assessment of an Ultra-Wide Band Positioning System for Indoor Wheelchair Court Sports. *Proc. of the Inst. of Mechanical Engineers, Part P: Journal of Sports Eng. and Techn.* 229, 2 (April 2015).
- [46] X. Qing, Zhi Ning Chen, and T. S. P. See. 2009. Sectorized Antenna Array for Indoor Mono-Station UWB Positioning Applications. In *Proc. of the 3rd European Conf. on Antennas and Propagation*. 822–825.
- [47] T. H. Riehle, P. Lichter, and N. A. Giudice. 2008. An Indoor Navigation System to Support the Visually Impaired. In *Proc. of the 30th Int. Conf. of the IEEE Engineering in Medicine and Biology Society*.
- [48] M. Rzymowski, P. Woznica, and L. Kulas. 2016. Single-Anchored Indoor Localization Using ESPAR Antenna. *IEEE Antennas and Wireless Prop. Letters* 15 (Nov. 2016).
- [49] B. Silva, Z. Pang, J. Åkerberg, J. Neander, and G. Hancke. 2014. Experimental Study of UWB-based High Precision Localization for Industrial Applications. In *Proc. of the Int. Conf. on Ultra-WideBand (ICUWB)*.
- [50] IEEE Computer Society. 2015. Standard for Low-Rate Wireless Networks. (2015).
- [51] X. Sun, Y. Ma, J. Xu, J. Zhang, and J. Wang. 2008. A high Accuracy Mono-Station UWB Positioning System. In *Proc. of the Int. Conf. on Ultra-WideBand (ICUWB)*.
- [52] J. Tiemann, F. Eckermann, and C. Wietfeld. 2016. ATLAS - An Open-Source TDOA-based Ultra-Wideband Localization System. In *Proc. of the Int. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*.
- [53] D. Vasisht, S. Kumar, and D. Katabi. 2016. Decimeter-level Localization with a Single WiFi Access Point. In *Proc. of the 13th Usenix Conf. on Networked Systems Design and Implementation (NSDI)*.
- [54] K. Witralsal et al. 2016. High-Accuracy Localization for Assisted Living: 5G systems will Turn Multipath Channels from Foe to Friend. *IEEE Signal Processing Magazine* 33, 2 (March 2016).
- [55] K. Witralsal and P. Meissner. 2012. Performance Bounds for Multipath-Assisted Indoor Navigation and Tracking (MINT). In *Proc. of the Int. Conf. on Comm. (ICC)*.
- [56] H. Zhang, X. Cui, B. An, and T. A. Gulliver. 2013. A Distance and Angle Estimated Method based on Single UWB Station. In *Proc. of the Int. Conf. on Sig. Proc., Comm. and Computing (ICSPCC)*.
- [57] J. Zhang, P. V. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney. 2009. UWB Systems for Wireless Sensor Networks. *Proc. IEEE* 97, 2 (Feb. 2009).

Paper C

B. Grosswindhager, C.A. Boano, M. Rath, and K. Römer. **Concurrent Ranging with Ultra-Wideband Radios: From Experimental Evidence to a Practical Solution.** In *Proceedings of the 38th IEEE International Conference on Distributed Computing System (ICDCS '18)*, pages 1460-1467, Vienna, Austria. July 2018.

©2018 IEEE

ISBN: 978-1-5386-6871-9

DOI: 10.1109/ICDCS.2018.00149

Link: <https://ieeexplore.ieee.org/document/8416412>

Abstract. To enable future location-aware Internet of Things (IoT) applications, Ultra-wideband (UWB) technology provides centimeter-accurate distance estimations. In the common case of a non-synchronized network, at least $N \cdot (N - 1)$ message exchanges are required to derive the distance between N nodes. Enabling concurrent ranging between an initiator and an arbitrary number of responders can drastically reduce the amount of necessary transmissions and hence increases the efficiency of UWB systems. Although the feasibility of concurrent ranging has been proven experimentally, several key challenges still need to be addressed to practically implement concurrent ranging in real-world UWB systems, such as the automatic detection of multiple responses, the identification of a responder, as well as the detection of overlapping responses (especially in the presence of multipath components). In this paper, we provide a concurrent ranging solution tackling the aforementioned challenges. Among others, our solution enables (i) to detect responses in the CIR reliably, (ii) to encode the responder ID in the CIR to allow personalized ranging, as well as (iii) to mitigate the impact of overlapping responses and multipath components. We further show how the proposed solution increases the scalability of concurrent ranging in real-world UWB-based distributed systems.

My contribution. I am the main author of this publication and developed all the presented concepts and algorithms. Furthermore, I have carried out all the experiments to show the effectiveness of the techniques to enable a practical solution for concurrent ranging. Carlo Boano helped me significantly in writing and structuring the paper and with Michael Rath I had several fruitful technical discussions. I presented the paper at ICDCS'18.

Concurrent Ranging with Ultra-Wideband Radios: From Experimental Evidence to a Practical Solution

Bernhard Großwindhager, Carlo Alberto Boano, Michael Rath, and Kay Römer
Faculty of Electrical and Information Engineering, Graz University of Technology, Austria
E-mail: {grosswindhager, cboano, mrath, roemer}@tugraz.at

Abstract—To enable future location-aware Internet of Things (IoT) applications, Ultra-wideband (UWB) technology provides centimeter-accurate distance estimations. In the common case of a non-synchronized network, at least $N \cdot (N - 1)$ message exchanges are required to derive the distance between N nodes. Enabling concurrent ranging between an initiator and an arbitrary number of responders can drastically reduce the amount of necessary transmissions and hence increases the efficiency of UWB systems. Although the feasibility of concurrent ranging has been proven experimentally, several key challenges still need to be addressed to practically implement concurrent ranging in real-world UWB systems, such as the automatic detection of multiple responses, the identification of a responder, as well as the detection of overlapping responses (especially in the presence of multipath components). In this paper, we provide a concurrent ranging solution tackling the aforementioned challenges. Among others, our solution enables (i) to detect responses in the CIR reliably, (ii) to encode the responder ID in the CIR to allow personalized ranging, as well as (iii) to mitigate the impact of overlapping responses and multipath components. We further show how the proposed solution increases the scalability of concurrent ranging in real-world UWB-based distributed systems.

Index Terms—Channel impulse response, concurrent ranging, Decawave DW1000, Multipath components, Ultra-wideband.

I. INTRODUCTION

Ultra-wideband (UWB) radios allow for precise ranging and localization thanks to the high bandwidth (≥ 500 MHz). Theoretical work on UWB technology dates back to the late 1990s and early 2000s [1], [2]. However, UWB was not commercially successful [3] until (i) the release of the IEEE 802.15.4a amendment adding a UWB physical layer [4] and (ii) the commercialization of the first low-cost IEEE 802.15.4-compliant UWB transceiver, the Decawave DW1000 [5]. These two key drivers have allowed UWB to become a key technology enabling location-aware IoT applications with centimeter-level positioning accuracy [6].

Scheduled ranging. The distance between two UWB nodes in a non-synchronized network is estimated by carrying out a two-way ranging scheme, i.e., a pair-wise exchange of at least two messages between two nodes: an initiator and a responder. To estimate the distance between all N nodes in a network, one typically needs to schedule the exchange of $N \cdot (N - 1)$ messages. This procedure is not only time-consuming; it is especially energy-inefficient, considering that the DW1000 radio draws up to 155mA and 90mA in receive and transmit mode, respectively. This is significantly higher than for other low-power wireless technologies such as BLE and LoRa [7].

Concurrent ranging. The short pulses transmitted by UWB radios reduce multipath fading and, hence, allow to resolve individual multipath components. Recent work has used this capability to build multipath-assisted indoor localization systems [8], [9]. This property can also be used to extract the simultaneous responses of an arbitrary number of responders. This principle is called *concurrent ranging* and its feasibility was shown experimentally by Corbalán and Picco [10]. Instead of scheduling several ranging operations between an initiator and other responders in the network, the latter reply simultaneously to a single broadcast message. The different responses (i.e., the transmitted pulses of the preamble from each responder) are visible in the channel impulse response (CIR) of the initiator. Hence, by detecting these pulses (i.e., by identifying the signal peaks associated to the responders in the channel impulse response estimated by off-the-shelf UWB transceivers), it is possible to estimate the path delay and distance to all responders concurrently.

Open challenges. Although the feasibility of concurrent ranging was shown experimentally by Corbalán and Picco [10], several key challenges still need to be addressed to practically implement concurrent ranging in real-world UWB systems.

I. Automatic detection of multiple responses. To practically implement concurrent ranging, it is necessary to automatically detect the responses of several nodes in the received CIR. In other words, an initiator should be able to process the estimated CIR *at run-time* and reliably identify the signal peaks associated to the different responders. Although existing work has discussed possible approaches to achieve this goal [10], a practical implementation working at run-time is still missing.

II. Identifying responders. A key challenge hindering the feasibility of concurrent ranging in real-world systems is the impossibility of associating a distance estimate to a specific responder, i.e., the anonymity of ranging. Previous work investigating concurrent ranging applied artificial setups where all nodes are placed in a line [10], which gives the initiator prior knowledge about the order in which the signal peaks associated to the responders are received in the CIR. In practical situations, however, the relative locations of nodes are typically unknown. Approaches making use of cross-correlation between the CIR acquired with concurrent responders and a previously-obtained CIR for each isolated responder [10] are also not applicable, as the channel impulse response for an isolated responder varies depending on its position and on the surrounding environment.




Fig. 1: (a) Floor plan showing line-of-sight (LOS) path and first-order reflections (MPC1 – MPC4); (b) Theoretical multipath reflections with a bandwidth of 900 MHz (top) and 50 MHz (bottom).

III. Detecting overlapping responses. If several responders are at a similar distance from the initiator, their responses will overlap, making it difficult to extract meaningful information from the channel impulse response. To date, no practical solution to decode overlapping responses has been presented.

IV. Mitigating the impact of multipath reflections. Another open challenge is how to differentiate between a response and a strong multipath component from another responder. Corbalán et al. [10] suggested to use power boundaries based on the Friis equation to differentiate between main CIR peaks and disturbing multipath components. This principle has three main issues limiting its applicability in real-world networks: (i) the Friis equation is idealized and does not hold true in typical UWB operational areas; (ii) in the case of an attenuated direct path due to obstacles, it is likely that multipath reflections have higher amplitudes than the direct path; (iii) the amplitude of the peaks in a CIR derived from low-cost UWB transceivers is highly varying. Hence, there is a need for algorithms that operate independently of the absolute amplitude of the signal peaks associated to the different responders.

V. Scalable concurrent ranging. Existing work has not yet focused on how to maximize the number of responders performing concurrent ranging. In current solutions, an initiator can obtain responses from multiple responders, but only as long as the latter are physically far away from each other, which limits the applicability of concurrent ranging in typical indoor UWB application settings.

Contributions. In this paper, we tackle all aforementioned challenges and provide a solution that allows the practical implementation of concurrent ranging on off-the-shelf UWB devices. We hence significantly advance the state-of-the-art in UWB concurrent ranging by addressing the yet open challenges highlighted by recent feasibility studies [10]. First, we describe a method to let initiators automatically detect responses in the channel impulse response at run-time *independently* of their absolute amplitude, hence making

concurrent ranging feasible in real-world UWB applications (Sect. IV). Second, we present a novel technique based on *pulse shaping* that allows to associate a distance measurement to a specific responder, so that ranging is no longer anonymous (Sect. V). Third, we show that the employed algorithm allows to detect the signal peaks associated to the different responders reliably even in the case of overlapping responses (Sect. VI). Fourth, we propose a method to prevent the overlap of responses and strong multipath components from other responders by employing response position modulation (Sect. VII). Finally, we show that, by combining response position modulation and pulse shaping, we can provide a *scalable* concurrent ranging solution that can be practically implemented on off-the-shelf UWB devices (Sect. VIII). Before presenting our contributions in detail, we provide the reader with basic information about ultra-wideband technology (Sect. II) and concurrent ranging (Sect. III).

II. ULTRA-WIDEBAND BASICS

The use of a high bandwidth (and consequently very short pulses) in UWB transceivers allows to resolve individual multipath components (MPC) as illustrated in Figure 1. In particular, Figure 1a shows a rectangular floor plan with a transmitter (TX) sending pulse signals to a receiver (RX). Due to the omni-directional wave propagation and to the reflections from walls, multiple versions of the same pulse arrive at the receiver. Figure 1a shows the line-of-sight (solid) as well as the first-order MPCs (dashed), i.e., the pulse is reflected only at a single object. Figure 1b shows the (theoretically) received pulses. The top figure shows pulses with a bandwidth of 900 MHz (i.e., the maximum bandwidth of the DW1000 radio [5]), whilst the bottom figure shows pulses sent with a bandwidth of 50 MHz. Due to the steep rising edge at 900 MHz, the precision of the distance estimation in a system based on time-of-flight (ToF) is increased.




Fig. 2: Estimated CIR obtained from the DW1000 radio in an indoor environment. The path delay of multipath components is marked with τ_k , where τ_0 denotes the LOS component.

Furthermore, the multipath reflections at 50 MHz are highly overlapping. This makes narrowband systems more susceptible to multipath fading and – in contrast to UWB systems – it is not possible to extract multipath components. This is an important observation, as widely-used low-power wireless technologies such as BLE have a bandwidth that is even smaller than 5 MHz.

At a reasonably high bandwidth, Figure 1b resembles the channel impulse response (CIR), i.e., information about the multipath propagation consisting of reflections from walls and scattering from other objects. The channel impulse response $h(t)$ can be modeled as [8]:

$$h(t) = \sum_{k=0}^K \alpha_k \delta(t - \tau_k) + \nu(t) \quad (1)$$

with α_k and τ_k denoting, respectively, the complex-valued amplitude and path delay of K deterministic multipath components resulting from specular reflections from walls, windows, or doors. The last term $\nu(t)$ is diffuse or non-deterministic multipath, i.e., higher-order reflections or signal components due to scattering. UWB transceivers such as the Decawave DW1000 provide a channel impulse response estimation to precisely estimate the arrival time of a packet by detecting the first pulse in the CIR. The DW1000 radio estimates the arrival time with a precision of 15.65ps (using a 63.9 GHz sampling clock), which results in a distance resolution of 4.69mm [11]. Figure 2 shows an exemplary estimated channel impulse response obtained with the DW1000 radio. It depicts the line-of-sight (LOS) component (τ_0) and significant multipath reflections ($\tau_1 - \tau_5$). Besides estimating the arrival time of a signal, the CIR information can be used to enable multipath-assisted UWB localization [8], [9]. Furthermore, the CIR can be used to detect a degrading channel as well as any change of the surrounding environment: this can be exploited to adapt UWB physical layer (PHY) parameters and increase communication performance [7]. In this paper, instead, we use the CIR to receive simultaneous responses from several neighbors in a single packet and to estimate the distance of a wireless node to each of the neighbors simultaneously, as shown in Sect. III.




Fig. 3: Principle of single-sided two-way ranging (SS-TWR) and concurrent ranging schemes.

III. CONCURRENT RANGING

The single-sided two-way ranging (SS-TWR) scheme traditionally used to derive the distance between two users in a non-synchronized network requires the exchange of two messages (*INIT*, *RESP*) between an initiator and a responder (see Figure 3). Thus, in a network with N users, each node requires $N - 1$ transmissions and receptions.

Given the high current draw of UWB transceivers (especially in receive mode [7]), it is important to reduce the number of messages exchanged in a network to make UWB systems feasible for building location-aware IoT applications. Furthermore, scheduling the distance estimation to each neighbor requires a significant amount of time, which increases channel utilization and traffic load, as well as the inaccuracy of ranging in mobile settings.

In a *concurrent ranging* scheme, instead, the initiator broadcasts the *INIT* message to all neighbors (responders), who reply simultaneously with a *RESP* message after a constant delay Δ_{RESP} . Consequently, the *RESP* messages (containing the timestamps $t_{rx,i}$ and $t_{tx,i}$ in the payload) sent by all responders are overlapping in time. In narrowband transceivers, this leads to (unusable) severely overlapping pulses, as shown in Figure 1b. Using ultra-wideband radios, instead, it is possible to identify the signal peaks associated to each responder in the estimated channel impulse response. By employing concurrent ranging, the total number of messages needed to estimate distances to all neighbors is hence reduced from $N \cdot (N - 1)$ to N . In fact, the initiator has to broadcast just one message and, more importantly, to receive just a single message that aggregates all responses.

Figure 3 also shows the UWB PHY frame structure according to the IEEE 802.15.4 standard. It consists of the preamble, start-of-frame delimiter (SFD), physical layer header (PHR), as well as a payload. The channel impulse response, used to derive responses from multiple neighbors, is estimated solely from the preamble and is hence independent from the payload.

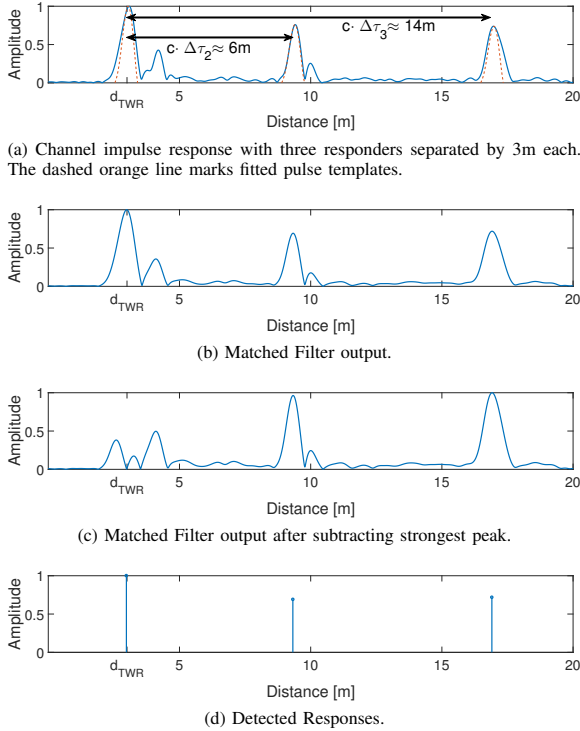


Fig. 4: Principle of the proposed response detection algorithm.

Thus, the extracted responses from the CIR correspond to the pulses transmitted in the preamble (consisting of a pre-defined sequence of single pulses). The payload, instead, is not sent as single pulses, but as bursts of pulses [7].

The IEEE 802.15.4 standard defines the timestamp of a frame as the beginning of the first symbol of the PHR (RMARKER, marked with arrow in Figure 3) [12]. Hence, the minimum delay Δ_{RESP} consists of the duration of PHR and payload of *INIT* message, as well as the preamble and SFD of the *RESP* message. Using a data rate of $DR = 6.8$ Mbps, a pulse repetition frequency $PRF = 64$ MHz, and a preamble symbol repetition $PSR = 128$, this results in a duration of $178.5\mu s$. Additionally, we evaluated experimentally the minimum time necessary to switch the DW1000 radio from receive to transmit mode, which is less than $100\mu s$. Including a safety gap, we hence set the delay Δ_{RESP} to $290\mu s$.

Figure 4a shows an acquired (normalized) channel impulse response from the *RESP* message when three neighbors are responding concurrently in a hallway. The three responders are placed at a distance from the initiator of $d_1 = 3m$, $d_2 = 6m$, and $d_3 = 10m$, respectively. Three significant peaks are visible in the CIR shown in Figure 4a, representing the strongest path component of each neighbor. The distance between the initiator and responder 1 is derived from the SS-TWR scheme (see Figure 3), as it is still possible to decode

one of the concurrently transmitted payloads containing the required timestamps [10]. The formula to calculate the distance between initiator and responder 1 is:

$$d_{TWR} = \frac{(t_{rx,init} - t_{tx,init}) - (t_{tx,1} - t_{rx,1})}{2} \cdot c \quad (2)$$

with c denoting the propagation speed in air. However, the distance between the initiator and the remaining responders is derived from the CIR. Due to the longer time-of-flight, the responding peaks of responder 2 and 3 arrive with a delay of $\Delta\tau_2 = 2 \cdot (\tau_2 - \tau_1)$ and $\Delta\tau_3 = 2 \cdot (\tau_3 - \tau_1)$ at the initiator, with τ_i ($i \in \{1, \dots, 3\}$) denoting the path delay between the initiator and each responder. The resulting delays $\Delta\tau_2$ and $\Delta\tau_3$ are due to the *INIT* as well as the *RESP* message: hence, they have to be halved to correctly estimate the distance. The resulting distance between initiator and responders is thus $d_2 = d_{TWR} + \frac{c \cdot \Delta\tau_2}{2} = 6m$ and $d_3 = d_{TWR} + \frac{c \cdot \Delta\tau_3}{2} = 10m$, respectively.

Limited TX timestamp resolution. The Decawave DW1000 UWB transceiver has the useful feature of *delayed transmission*. The latter enables to set a future timestamp at which the transceiver sends a packet. This allows to align a pre-calculated timestamp with the real transmit timestamp and embed it in the message being transmitted ($t_{tx,i}$ in Figure 3). Unfortunately, the Decawave DW1000 ignores the low-order 9 bits of the delay transmit value, limiting the transmission timestamp resolution to approximately 8ns [11, p. 26]. This is not an issue in the classical SS-TWR scheme as the real transmit timestamp is anyway embedded in the message, but it has a severe impact on the precision of the concurrent ranging scheme, as it negatively affects the concurrency of the *RESP* messages of the neighbors. Given that this issue is hardware-dependent and could be solved in the next-generation UWB transceivers, this problem is out of scope of this paper.

IV. RELIABLE RESPONSE DETECTION

To make concurrent ranging feasible in real-world systems, it is essential to detect responses from the neighbors reliably in the channel impulse response. To this end, we propose a scheme based on the *search and subtract* algorithm [13]. This algorithm employs a matched filter computing the correlation between the received CIR and a pulse template with duration T_p transmitted by the UWB radio. To detect the $N-1$ strongest responses in the CIR, we use the following procedure:

- 1) We first upsample the CIR using *fast Fourier transform* in order to obtain a smoother signal. Furthermore, to correct for the unknown time offset of the CIR derived from the DW1000 radio [8], the channel impulse response is aligned with the distance d_{TWR} (see Eq. 2 and Figure 4). This step is not necessarily required, as the differences of the peaks are relevant for concurrent ranging, but it enhances visualization of the CIR and simplifies plausibility checks of the result.

- 2) Denoting the pulse template as $s(t)$, we can define the time-discrete impulse response of the matched filter as the time-reversed pulse template:

$$\mathbf{h}_{MF} = [s((N_p - 1) \cdot T_s), s((N_p - 2) \cdot T_s), \dots, s(0 \cdot T_s)]$$

with T_s marking the sampling period, and $N_p = T_p/T_s$ the number of samples of the pulse. The output of the matched filter \mathbf{y} is computed as the discrete convolution (*) between the impulse response of the matched filter \mathbf{h}_{MF} and the derived CIR denoted as \mathbf{r}

$$\mathbf{y} = \mathbf{h}_{MF} * \mathbf{r}. \quad (3)$$

Figure 4b shows the matched filter output of the CIR depicted in Figure 4a. It is evident that the matched filter increases the signal-to-noise ratio (SNR) of the channel impulse response.

- 3) We identify the sample corresponding to the maximum of the matched filter output \mathbf{y} , indicating the index of the strongest path l_k . The latter relates to the path delay with $\tau_k = l_k \cdot T_s$ and to the path length with $d_k = \tau_k \cdot c$.
- 4) To reduce complexity, instead of the least squares solution suggested in [13], we calculate the amplitude of the strongest path $\hat{\alpha}_k$ as the amplitude of \mathbf{y} at sample l_k .
- 5) The estimated neighbor response ($\hat{\alpha}_k s(t - \tau_k)$) is subtracted from the received signal \mathbf{r} . Figure 4c shows the matched filter output of the remaining signal after subtracting the strongest peak/neighbor response.
- 6) We repeat steps 2 to 5 with the remaining output signal until the $N - 1$ strongest paths are detected. Figure 4d shows the $N - 1 = 3$ strongest peaks corresponding to the responses from the three neighbors.
- 7) Independently of their amplitude $\hat{\alpha}_k$, the responses ($\hat{\alpha}_k, \tau_k$) are arranged in ascending order starting with the one of the closest neighbor. Being $\hat{\alpha}_1$ and τ_1 the amplitude and path delay of responder 1, respectively, the distance of responder i is calculated as:

$$d_i = d_{TWR} + \frac{c \cdot (\tau_i - \tau_1)}{2}. \quad (4)$$

Due to lack of information from Decawave regarding the transmitted pulse used in the DW1000 radio, we identified the pulse shape $s(t)$ used in step 2 of our detection algorithm with a measurement campaign. In particular, we connected transmitter and receiver with a SMA cable and a 60 dB attenuator to avoid reflections and saturation of the transceiver, respectively. We then let the receiver log 1000 CIRs and, in a post-processing step, cut the direct path component from the channel impulse response, and calculate the average pulse shape. Figure 5a shows the default pulse shape at Channel 7 (900 MHz bandwidth).

Following the algorithm described in this section, we are hence able to reliably detect the responses of all neighbors in the CIR to perform concurrent ranging.




Fig. 5: Pulse shape $s_i(t)$ for different values of the TC_PGDELAY register. 0x93 represents the default value.

V. ENCODING RESPONDER ID IN THE CIR

To make concurrent ranging usable in real-world applications, it is required to encode the ID of the responder in the channel impulse response. The preamble, indeed, consists of a fixed sequence of pulses (see Sect. III) and does not embed the identity of the sender. Hence, in the concurrent ranging scheme, the responses of the neighbors derived from the CIR do not contain any information to associate the responses to the corresponding neighbors. Consequently, distance estimations are typically anonymous in a concurrent ranging scheme. To remedy this problem, we suggest to use *pulse shaping* to associate a peak in the channel impulse response to a responder, i.e., to change the shape of the transmitted pulses as a function of the responder ID.

Pulse shaping. The off-the-shelf DW1000 radio provides the ability to change the width of the transmitted pulses via the 8-bit register TC_PGDELAY. Changing the value of this register effectively alters the output bandwidth [11, p. 148]. While making the pulse narrower (i.e., increasing the bandwidth) is not an option due to the regulatory spectral mask, making the pulse wider, instead, does not violate the regulations. Figure 5 exemplarily shows the pulse shape¹ $s_i(t)$ obtained when using the same settings as in Sect. IV (i.e., Channel 7, $PRF = 64$ MHz), and when configuring TC_PGDELAY with values 0x93 (s_1), 0xC8 (s_2), 0xE6 (s_3), and 0xF0 (s_4). Given that the default value (0x93) is the lower limit, up to 108 different pulse shapes are supported, which limits the theoretical number of possible responders differentiable with the proposed pulse shaping technique.

¹Please note that (i) the amplitude of the pulses shown in Figure 5 are different due to scaling the pulses to unit energy, and that (ii) 0x93 is the default value of the TC_PGDELAY register for the employed configuration.

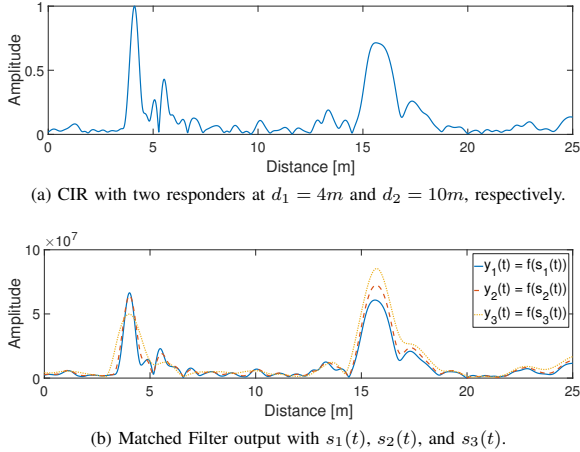


Fig. 6: CIR and matched filter output in the case of two responders replying with different pulse shapes.

No impact on ranging performance. Increasing the default value of `TC_PGDELAY` results in a wider pulse, which reduces the bandwidth and hence, in theory, also the ranging precision. To evaluate if the change in pulse shape has any impact on ranging performance, we place two UWB nodes three meters apart from each other in an office environment, and perform 5000 SS-TWR operations with three different pulse shapes (s_1, s_2, s_3 in Figure 5). Our results show that the standard deviation of the difference between the true distance and the estimated distance for the three pulse shapes is $\sigma_1 = 0.0228m$ (s_1), $\sigma_2 = 0.0221m$ (s_2), and $\sigma_3 = 0.0283$ (s_3), respectively. Therefore, changing the pulse shape by configuring `TC_PGDELAY` has a negligible impact on the ranging precision and can be safely used to encode the responder ID in a concurrent ranging scheme.

Identifying pulse shapes. Figure 6a shows the acquired channel impulse response when a neighbor at a distance $d_1 = 4m$ answers using the default pulse $s_1(t)$ (0×93 , see Figure 5a), and a second neighbor at a distance of $d_2 = 10m$ responds with a wider pulse $s_3(t)$ ($0 \times E6$, see Figure 5c). The different pulse shape is clearly visible in the channel impulse response. Performing the algorithm described in Sect. IV with $N_{PS} = 3$ possible pulse templates $s_i(t)$ (with $i \in \{1, \dots, N_{PS}\}$) results in the matched filter outputs $y_i(t)$ shown in Figure 6b. In each case, the responses of the neighbors are easily detectable, independently of the pulse template. To decode the transmitted pulse shape of the responders and hence their ID, we compare the estimated amplitudes of the neighbor responses $\hat{\alpha}_{k,i}$ (with k denoting the number of the response) of all N_{PS} matched filter outputs $y_i(t)$. The pulse shape i at which the amplitude $\hat{\alpha}_{k,i}$ is the maximum determines the pulse shape used by the responder. Therefore, in Figure 6b, the first response

d_2 [m]	6	7	8	9	10
$s_2(t)$ ($0 \times C8$) [%]	99.9	99.5	99.8	100	99.8
$s_3(t)$ ($0 \times E6$) [%]	99.2	99.7	99.9	100	100

TABLE I: Percentage of pulse shapes identified correctly.

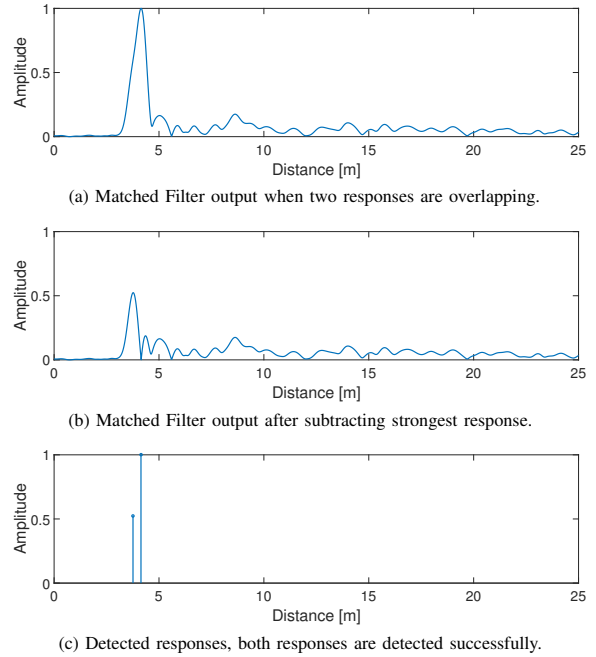


Fig. 7: Performance of the proposed algorithm in the case of overlapping responses from multiple responders.

corresponds to the pulse template $s_1(t)$ (blue, solid) and the second to the pulse template $s_3(t)$ (yellow, dotted). Please note that, in the provided example, the number of pulse templates is set to $N_{PS} = 3$, but, in principle, up to 108 concurrent responders can be supported.

Evaluation. We evaluate the performance of the proposed technique by placing an initiator and a responder at a fixed distance $d_1 = 3m$. Another responder is placed at a variable distance $d_2 \in \{6, 7, 8, 9, 10\}m$ from the initiator. Responder 1 uses the default pulse shape $s_1(t)$ (see Figure 5a). Responder 2 uses either $s_2(t)$ (see Figure 5b) or $s_3(t)$ (see Figure 5c). For each distance and pulse shape, we perform 1000 concurrent ranging operations. Table I shows how many pulses could successfully be identified. Regardless of the pulse shape and of the distance, a responder could successfully be identified in more than 99.2% of the cases, showing that *pulse shaping* can effectively encode the identity information of a responder in a concurrent ranging scheme.




Fig. 8: Combining response position modulation with pulse shaping. In this example, nine users perform concurrent ranging by employing *four* slots and *three* different pulse shapes. The responders with ID 0, 1, 2 make use of pulse shape $s_1(t)$, $s_2(t)$, and $s_3(t)$, respectively.

VI. DETECTION OF OVERLAPPING RESPONSES

So far, we assumed that the responses are nicely separated in time, i.e., that responses from responders placed at similar distances from the initiator do not overlap. We study next the performance of the detection algorithm proposed in Sect. IV in the presence of overlapping responses and show that it reliably detects responses from nodes at similar distances.

In particular, we compare the performance of our proposed algorithm with a *threshold-based algorithm* as proposed by Falsi et al. [13]. The threshold-based algorithm compares the channel impulse response with a defined threshold. If the CIR crosses this threshold, the maximum of the following N_p samples, i.e., the pulse duration, is derived. This operation is repeated until $N - 1$ peaks are detected.

Both the *search and subtract* and the *threshold-based* algorithm exhibit a good performance when the responses are well-separated from each other. However, as soon as the responses are overlapping due to a similar distance (and hence time-of-flight) of several responders, the algorithm proposed in Sect. IV outperforms threshold-based algorithms.

We show this by acquiring 2000 concurrent ranges from two responders placed at the same distance $d_1 = d_2 = 4m$ from the initiator. The responses of the two nodes are highly overlapping: as a result, only one peak is visible in Figure 7a. Figure 7b shows that, after subtracting the strongest response with our proposed algorithm (step 5 described in Sect. IV), we obtain an easily detectable second response. Our evaluation shows that the threshold-based algorithm detects both responses in only 48% of the trials, whilst our algorithm is successful in 92.6% of the cases. Therefore, the algorithm that we propose in Sect. IV can be also used to detect the presence of responders positioned at similar distances from the initiator.

Please note that, even if the nodes are physically positioned at the same distance, the responses may still not overlap consistently. The reason is the limited transmission timestamp resolution of the Decawave DW1000 discussed in Sect. III, i.e., even if the distance is the same and two responders reply simultaneously, there might be an offset of $\pm 8ns$. For this reason, in our performance evaluation, we have considered only trials in which the responses are actually overlapping.

VII. MITIGATING THE IMPACT OF STRONG MULTIPATH

In multipath-rich environments it is likely that several strong multipath components are received. In fact, it might be even the case that a received MPC is stronger than the direct path component due to blocked and attenuated line-of-sight. In these situations, it can hence be challenging to differentiate between a response from a neighbor and a dominant MPC.

Response position modulation. The most effective solution in these scenarios is to avoid that multipath components are overlapping with responses from other responders by separating them in time. Hence, we propose *response position modulation*, i.e., to modulate the response delay Δ_{RESP} . For this purpose, we introduce $\Delta'_{RESP} = \Delta_{RESP} + \delta_i$, where δ_i sets an additional individual delay for each responder i . This reduces the probability of overlapping responses, as the latter are spread over a wider range of the CIR. In total, the CIR estimate provided by the DW1000 has a length of 1016 samples (for $PRF = 64MHz$) with a sampling period of $T_s = 1.0016ns$. Thus, the maximum offset is $\delta_{max} \approx 1017ns$, which relates to a maximum distance offset of $\delta_{max} \cdot c \approx 307m$. Knowing the maximum communication range and an estimate of the delay spread allows to define the number of non-overlapping responses fitting in the CIR.

VIII. COMBINING RESPONSE POSITION MODULATION WITH PULSE SHAPING FOR A HIGHER SCALABILITY

The response position modulation proposed in Sect. VII poses a stringent limitation on the maximum number of users due to its dependency on the communication range r_{max} . Indeed, considering that the range of UWB transceivers can easily extend to $r_{max} > 75\text{m}$ [7], [14], just up to $N_{RPM} = \frac{\delta_{max} \cdot c}{r_{max}} \approx 4$ responders are supported to ensure non-overlapping responses. Due to this limitation, we suggest to combine response position modulation with pulse shaping to increase the maximum amount of users that can make simultaneous use of concurrent ranging, and hence its scalability.

To this end, we use response position modulation to split the channel impulse response into N_{RPM} slots separated by $\delta = \frac{\delta_{max}}{r_{max}}$, hence reducing the number of responders which possibly interfere with each other. This allows to mitigate the impact of overlapping responses and multipath reflections. Within each slot, we use pulse shaping to identify the responder, as described in Sect. V. The number of pulse shapes N_{PS} defines the number of supported users per slot. Each responder is assigned to a slot as well as a pulse shape depending on its responder ID. The used slot is $n_{RPM} = ID \% N_{RPM}$, whilst the used pulse shape is $n_{PS} = \lfloor ID/N_{RPM} \rfloor$. Depending on the slot number n_{RPM} , we set the additional delay $\delta_i = n_{RPM} \cdot \delta$.

In Figure 8 we divide the CIR into $N_{RPM} = 4$ slots using response position modulation. Within a slot, up to $N_{PS} = 3$ responders are active. In the example in Figure 8, three responders make use of slot 0 and slot 3, two responders make use of slot 1, and one responder makes use of slot 2. Therefore, the total number of responders is $N = 9$. The maximum number of responders that can make use of concurrent ranging in this scenario is $N_{max} = N_{RPM} \cdot N_{PS} = 12$. One can increase N_{max} (and hence the scalability of the system) by increasing the amount of pulse shapes N_{PS} and by increasing the number of slots N_{RPM} . As discussed in Sect. V the maximum number of pulse shapes is approximately 100. Assuming that the communication range r_{max} is limited manually by adapting the physical layer settings to 20m (which is sufficient for typical indoor applications [15]), the number of supported responders becomes more than 1500. Thus, using the proposed technique, the initiator requires just a single transmit and receive operation to estimate the distance to all neighbors simultaneously. Using classical two-way ranging schemes, instead, the initiator would need to send a packet to and to receive a packet from all the 1499 neighbors, respectively. This emphasizes the impact of the presented solutions, especially when having a high number of neighbors.

IX. CONCLUSIONS AND FUTURE WORK

In this paper, we developed a practical solution for concurrent ranging enabling an efficient distance estimation to multiple users in parallel. We first proposed an algorithm to detect responses within the estimated channel impulse response automatically, and showed that this technique is

highly-effective also in the presence of overlapping responses. Second, we encoded the responder ID in the CIR using different pulse shapes, making it possible to assign distance estimations to specific responders. Third, we introduced a technique called response position modulation to mitigate the impact of strong multipath components. Finally, we combined response position modulation and pulse shaping to increase the number of supported users and the scalability of the proposed concurrent ranging scheme.

In future work, we plan to use concurrent ranging to build an efficient cooperative or anchor-based localization system. Furthermore, in the work presented in this paper, we have neglected the impact of non-line-of-sight situations on the performance of concurrent ranging. We will hence investigate this impact thoroughly in the next months.

ACKNOWLEDGMENTS

This work was performed within the LEAD-Project ‘‘Dependable Internet of Things in Adverse Environments’’, funded by Graz University of Technology (Graz, Austria).

REFERENCES

- [1] M. Z. Win and R. A. Scholtz, ‘‘Impulse radio: How it works,’’ *IEEE Communications letters*, vol. 2, no. 2, pp. 36–38, 1998.
- [2] D. Porcino and W. Hirt, ‘‘Ultra-wideband radio technology: potential and challenges ahead,’’ *IEEE Communications Magazine*, vol. 41, no. 7, pp. 66–74, Jul. 2003.
- [3] P. Catherwood and W. Scanlon, ‘‘Ultrawideband communications - an idea whose time has still yet to come?’’ *IEEE Antennas and Propagation Magazine*, vol. 57, no. 2, pp. 38–43, May 2015.
- [4] IEEE, *Standard for Local and metropolitan area networks. Part 15.4. Amendment 1: Add Alternate PHYs*, Std. 802.15.4a, 2007.
- [5] Decawave Ltd., *DW1000 Datasheet. Version 2.12*, 2016.
- [6] J. Zhang, P. V. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney, ‘‘UWB Systems for Wireless Sensor Networks,’’ *Proceedings of the IEEE*, vol. 97, no. 2, Feb. 2009.
- [7] B. Grobwindhager, M. Rath, C. A. Boano, and K. Römer, ‘‘Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications,’’ in *Proceedings of the 19th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Jun. 2018.
- [8] J. Kulmer, S. Hinteregger, B. Grobwindhager, M. Rath, M. S. Bakr, E. Leitinger, and K. Witrissal, ‘‘Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning,’’ in *Proc. of the Intern. Conf. on Commun. (ICC) Workshops*, 2017, pp. 1239–1245.
- [9] B. Grobwindhager, M. Rath, J. Kulmer, S. Hinteregger, M. S. Bakr, C. A. Boano, K. Witrissal, and K. Römer, ‘‘UWB-based Single-anchor Low-cost Indoor Localization System,’’ in *Proceedings of the 15th ACM International Conference on Embedded Networked Sensor Systems (SenSys), demo session*. ACM, Nov. 2017.
- [10] P. Corbalan and G. P. Picco, ‘‘Concurrent Ranging in Ultra-wideband Radios: Experimental Evidence, Challenges, and Opportunities,’’ in *Proceedings of the 15th International Conference on Embedded Wireless Systems and Networks (EWSN)*, Feb. 2018.
- [11] Decawave Ltd., *DW1000 User Manual. Version 2.10*, 2016.
- [12] IEEE Computer Society, *Standard for Low-Rate Wireless Networks*, Std. 802.15.4, 2015.
- [13] C. Falsi, D. Dardari, L. Mucchi, and M. Z. Win, ‘‘Time of Arrival Estimation for UWB Localizers in Realistic Environments,’’ *EURASIP Journal on Advances in Signal Processing*, no. 1, Jul. 2006.
- [14] Decawave Ltd., *Application Note APS017. Maximising Range in DW1000 Based Systems. Version 1.0*, 2014.
- [15] D. Dardari, P. Closas, and P. M. Djurić, ‘‘Indoor Tracking: Theory, Methods, and Technologies,’’ *IEEE Transactions on Vehicular Technology*, vol. 64, no. 4, Apr. 2015.

Paper D

B. Grosswindhager, M. Stocker, M. Rath, C.A. Boano, and K. Römer. **SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags.** In *Proceedings of the 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN'19)*, pages 61–72, Montreal, Canada. April 2019. **Best Paper Award.**

©2019 ACM

ISBN: 9781450362849

DOI: 10.1145/3302506.3310389

Link: <https://doi.org/10.1145/3302506.3310389>

Abstract. A large body of work has shown that ultra-wideband (UWB) technology enables accurate indoor localization and tracking thanks to its high time-domain resolution. Existing systems, however, are typically designed to localize only a limited number of tags, and involve the exchange of several messages following a given schedule. As a result, the scalability of current solutions in terms of tag density is limited, as well as their efficiency and responsiveness. In this paper, we present SnapLoc, a UWB-based indoor localization system that allows an unlimited number of tags to self-localize at a theoretical upper bound of 2.3 kHz. In SnapLoc, a tag obtains the responses from multiple anchors simultaneously. Based on these signals, the tag derives the time difference of arrival between anchors and estimates its position. Therefore, SnapLoc does not require tags to actively transmit packets, but to receive only a single message. This allows tags to passively localize themselves and ensures that the performance of SnapLoc does not degrade with high node densities. Moreover, due to the (quasi-)simultaneous responses, a tight clock synchronization between anchors is not needed. We have implemented SnapLoc on a low-cost platform based on the Decawave DW1000 radio and solved limitations in the transceiver's timestamp resolution to sustain a high localization accuracy. An experimental evaluation shows that SnapLoc exhibits a 90% error and median error of 33 cm and 18 cm, respectively, hence enabling decimeter-level accuracy at fast update rates for countless tags.

My contribution. I am the main author of this publication and conceived the ideas and concepts. Michael Stocker put a tremendous effort in supporting me with the implementation of SnapLoc, especially with designing the firmware of the anchors and the timestamp correction methods. Furthermore, we have carried out the evaluations together. Carlo Boano helped me significantly in writing and structuring the paper and also Michael Rath supported in the writing process and to discuss some ideas. I presented the paper at IPSN'19 and won the Best Paper Award.

SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags

Bernhard Großwindhager, Michael Stocker, Michael Rath, Carlo Alberto Boano, and Kay Römer
Department of Electrical and Information Engineering, Graz University of Technology, Austria
{grosswindhager, michael.stocker, mrath, cboano, roemer}@tugraz.at

ABSTRACT

A large body of work has shown that ultra-wideband (UWB) technology enables accurate indoor localization and tracking thanks to its high time-domain resolution. Existing systems, however, are typically designed to localize only a *limited number* of tags, and involve the exchange of *several messages* following a given schedule. As a result, the scalability of current solutions in terms of *tag density* is limited, as well as their efficiency and responsiveness. In this paper, we present SnapLoc, a UWB-based indoor localization system that allows an unlimited number of tags to self-localize at a theoretical upper bound of 2.3 kHz. In SnapLoc, a tag obtains the responses from multiple anchors *simultaneously*. Based on these signals, the tag derives the time difference of arrival between anchors and estimates its position. Therefore, SnapLoc does *not* require tags to actively transmit packets, but to receive only a single message. This allows tags to passively localize themselves and ensures that the performance of SnapLoc does not degrade with high node densities. Moreover, due to the (quasi-)simultaneous responses, a tight clock synchronization between anchors is not needed. We have implemented SnapLoc on a low-cost platform based on the Decawave DW1000 radio and solved limitations in the transceiver's timestamp resolution to sustain a high localization accuracy. An experimental evaluation shows that SnapLoc exhibits a 90% error and median error of 33 cm and 18 cm, respectively, hence enabling decimeter-level accuracy at fast update rates for countless tags.

CCS CONCEPTS

• **Computer systems organization** → *Embedded and cyber-physical systems*; • **Networks** → *Location based services*;

KEYWORDS

Localization, TDOA, ultra-wideband, channel impulse response

ACM Reference Format:

Bernhard Großwindhager, Michael Stocker, Michael Rath, Carlo Alberto Boano, and Kay Römer. 2019. SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags. In *The 18th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2019) (IPSN '19)*, April 16–18, 2019, Montreal, QC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3302506.3310389>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IPSN '19, April 16–18, 2019, Montreal, QC, Canada
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6284-9/19/04...\$15.00
<https://doi.org/10.1145/3302506.3310389>

1 INTRODUCTION

Ultra-wideband (UWB) technology is becoming increasingly popular thanks to its robustness and outstanding localization accuracy. Spreading the signal over a wide bandwidth, indeed, results in: (i) greater immunity to multipath fading, (ii) better interference mitigation, (iii) higher throughput, as well as (iv) an improved timing resolution allowing for accurate localization and tracking [13].

Such a high time-domain resolution allows UWB-based solutions to significantly outperform narrowband RF technologies like Bluetooth Low Energy and Wi-Fi in terms of localization accuracy. These technologies, indeed, can hardly achieve a sub-meter accuracy [5, 18], and are hence unable to satisfy the requirements of location-aware Internet of Things (IoT) applications such as assisted living [42], robot navigation [15, 22], and smart manufacturing [20].

Because of the aforementioned properties, several works have used UWB technology to build indoor localization and tracking systems [1]. These works have shown that UWB-based solutions can achieve a localization accuracy up to a few cm [30], even in challenging conditions [24] and despite the use of a single anchor [14].

Existing solutions do not scale. Unfortunately, most of the existing solutions based on UWB technology focus on achieving a high localization accuracy, often disregarding properties such as *multi-tag support* and *high update rates* [33]. As a result, current systems typically support only a few tags and do not scale in terms of *tag density*, due to (i) the large number of messages exchanged, and (ii) the use of scheduling techniques for collision avoidance [33].

Large message overhead. A large number of UWB systems are indeed based on time of flight (TOF) techniques and make use of two-way ranging (TWR) schemes or a variant of it [24, 32]. These systems require the exchange of several consecutive messages, such that a mobile tag can derive the distance from multiple anchors and unambiguously determine its position. The large number of messages exchanged to carry out each distance estimation limits the overall update rate [25] and requires a tag to be heavily involved in the communication, which increases its energy-consumption. As mobile tags are typically battery-powered, their radio-on time should, instead, be minimized, in order to preserve their limited energy budget. Furthermore, sequentially estimating the distance to each anchor leads to inconsistent measurements in mobile and highly-dynamic settings (as one combines distances estimated at slightly different times), which limits the achievable accuracy.

Use of collision avoidance techniques. To reduce message overhead and avoid exchanging consecutive messages, a few UWB systems employ time difference of arrival (TDOA) techniques and allow a tag to broadcast only one message per position estimate [34, 41]. The latter is received from synchronized anchors, which compute the TDOA and communicate back the estimated position to the tag.

Whilst this allows to minimize the number of transmissions carried out by a tag and to shift the computational burden to more powerful anchors, one still needs to allocate *specific timeslots* to each tag in order to avoid collisions. Such scheduling techniques, which are also needed in systems based on TOF [24, 32] and single anchors [14], however, limit the number of tags that can be supported and, consequently, the scalability of a localization system.

Need for a tight synchronization. TDOA-based systems typically require anchors to be synchronous. For example, in anchor-initiated solutions [28], the tags estimate their position based on signals received from synchronized anchors. Whilst this approach allows tags to carry out self-localization without the need to transmit information, it requires a tight (ns-range) synchronization between anchors. However, this results in an overhead [41, 44] and is challenging [43]. Furthermore, the anchors still send messages sequentially, which requires also the tag to compensate for clock deviations [28].

Concurrent ranging still immature. Recent work on concurrent ranging has the potential to significantly reduce message overhead by exploiting simultaneous responses to a ranging request issued by an initiator [6]. However, concurrent ranging is still inapplicable in practical UWB systems due to: (i) the inability to identify responders and to discern them from strong multipath components [12], (ii) the high amount of payloads lost when responders are located at similar distances, as well as (iii) the limited transmit timestamp resolution of off-the-shelf UWB transceivers (see Sect. 2.3).

This state of affairs represents a significant problem, as increasing the density of tags in existing UWB systems results in a significant reduction of the localization update rate [26, 34, 37], due to the large message overhead and the use of collision avoidance techniques. In order to create UWB-based indoor localization systems that scale regardless of the number of tags, one would ideally (i) address the aforementioned limitations of concurrent ranging, and (ii) apply the latter to TDOA-based anchor-initiated approaches, in such a way that anchors are not required to be tightly synchronized. This would keep tags away from actively transmitting messages, minimize their radio-on time, and avoid the use of collision avoidance schemes. More importantly, this would enable an *unlimited number* of tags to passively self-localize at *fast update rates*.

Contributions. In this paper we present SnapLoc, a UWB-based indoor localization system that achieves exactly this. SnapLoc solves the limitations of existing concurrent ranging techniques and allows tags to obtain responses from multiple anchors *simultaneously*. Instead of scheduled sequential anchor messages [28], in SnapLoc multiple anchors reply (*quasi*-)simultaneously to an initialization message sent by a reference anchor. Based on these responses, a tag can quickly derive the time difference of arrival between anchor pairs and accurately estimate its position.

Therefore, in SnapLoc, a tag does not require to actively transmit messages, and its radio-on time can be reduced to a *single* read operation. This removes the need for a tight clock synchronization between anchors and eliminates the clock correction at the tag completely. Furthermore, SnapLoc's approach allows tags to passively localize themselves, ensuring that the performance does not degrade with high node densities. Theoretically, SnapLoc requires just 434 μ s to provide the tag with all the necessary information to

estimate its location. Thus, SnapLoc enables an unlimited number of tags to self-localize at position update rates up to 2.3 kHz.

Besides the reception of a single packet, a key property of SnapLoc is the use of information that is only contained in a packet's preamble for computing the actual position (see Sect. 3). In fact, a tag extracts the (quasi-)simultaneous responses from the anchors by analyzing the estimated channel impulse response (CIR) provided by standard-compliant UWB transceivers upon reception of a preamble. This avoids the need to correctly receive a payload and bypasses an intrinsic limitation of concurrent ranging.

To associate each response in the CIR to the correct anchor and to counteract the impact of strong multipath components, SnapLoc assigns an individual delay in the nanosecond range to each anchor (see Sect. 4). This, however, limits the maximum number of anchors due to the finite length of the estimated CIR. To overcome this restriction, we propose to partition the area of operation in multiple cells and let each cell operate using orthogonal *preamble codes*.

We have implemented SnapLoc on a low-cost platform based on the Decawave DW1000 radio (see Sect. 4.5). As discussed in Sect. 2.3, this transceiver constrains the transmit timestamp resolution to 8 ns, which severely limits the accuracy. We devise two techniques to restore a resolution of 15.65 ps and 1 ns, respectively, allowing SnapLoc to sustain a high localization accuracy (see Sect. 5).

An experimental evaluation in a common office and in a larger laboratory classroom shows that SnapLoc exhibits a 90% error and a median error in the order of 33.7 cm and 18.4 cm, respectively. Therefore, SnapLoc enables decimeter-level localization accuracy at fast update rates for an unlimited number of tags.

After describing the limitations of existing approaches in Sect. 2, this paper makes the following contributions:

- We introduce the design of SnapLoc, a UWB-based indoor localization system that allows an unlimited number of tags to self-localize at very high update rates (Sect. 3);
- We describe SnapLoc's principle, and detail on: how to reliably detect the anchors' responses, how to associate each response to the corresponding anchor, as well as how to derive the TDOA and compute a tag's position (Sect. 4);
- We implement SnapLoc on a low-cost platform based on the DW1000 radio and propose a technique to overcome the transceiver's limited transmit timestamp resolution (Sect. 5);
- We evaluate SnapLoc in common office environments and show that it enables decimeter-level localization accuracy at fast update rates also for high tag densities (Sect. 6);

After describing related work in Sect. 7, we conclude our paper in Sect. 8, along with a discussion on future work.

2 LIMITATIONS OF EXISTING APPROACHES

We begin our discussion by describing existing UWB-based localization approaches and highlighting their limitations. We first point out how solutions based on TOF and two-way ranging schemes typically incur a large message overhead (Sect. 2.1). We then discuss how TDOA-based approaches require a tight synchronization across anchors, which is complex to attain and may introduce errors lowering the localization accuracy (Sect. 2.2). Finally, we discuss recent work on concurrent ranging, and elaborate on the limitations that make it still inapplicable to practical UWB systems (Sect. 2.3).




Figure 1: Example of CIRs estimated by a UWB transceiver. Whilst the first path component is prominent in an environment with limited multipath (a), the same does not necessarily hold true in the presence of strong multipath components (b).

2.1 TOF-based Approaches

Time-based localization systems rely on measuring the travel time of a radio signal between two nodes (typically an *anchor* and a *tag*). Among time-based systems, the most popular approaches are *time of flight* (TOF) and *time difference of arrival* (TDOA).

TOF-based systems – often also referred to as time of arrival based (TOA-based) systems – determine the absolute distance between sender and receiver by measuring the time of flight of a packet multiplied with the propagation speed (i.e., speed of light in air c). Allowing a tag to passively self-localize in a TOF scheme (i.e., performing only one-way communications from anchors to tag) requires anchors and tags to be tightly synchronized, which implies a significant overhead and is often infeasible [44]. To avoid this, several UWB systems let a tag estimate the TOF from multiple anchors by making use of two-way ranging (TWR) or similar schemes [24, 32]. These schemes do not require synchronization between anchors and tags, but envisage the exchange of multiple messages, such that a tag can derive its distance to several anchors.

Limitations. Such an approach incurs a large communication overhead, considering that at least three (2D) or four (3D) anchors are necessary to unambiguously determine a tag’s position. Furthermore, common systems typically make use of up to eight [28, 41] or fifteen [22] anchors to increase the redundancy and robustness of localization, e.g., to mitigate non-line-of-sight (NLOS) conditions. This may result in tens of messages transmitted and received by a tag for each localization attempt, which limits the achievable update rate [25] and heavily affects its energy consumption.

2.2 TDOA-based Approaches

TDOA approaches do not require the absolute time of flight of a packet, but exploit the difference Δt in the arrival time of a signal at two reference points. Based on Δt , the difference in distance Δd between tag and reference points can be calculated as $\Delta d = \Delta t \cdot c$.

The advantage of TDOA approaches is that sender and receiver do not need to be synchronized, as the offset of the tag’s clock is canceled out [8]. This simplifies the system design and removes the need of exchanging several messages between tags and anchors.

Indeed, in most UWB-based TDOA localization systems, a tag broadcasts only one message per position estimate [34, 41]. The tag’s position is then estimated in a central localization engine computing the TDOA, which allows to shift the computational burden from the tag to other (more powerful) devices [35]. This is especially advantageous in applications where a central entity monitors the

position of all users, and where the tags do not necessarily need to know their own position, e.g., asset- or sports tracking [4, 29].

Limitations. Whilst such an approach minimizes the number of transmissions carried out by tags, one still needs to allocate *specific timeslots* to each tag in order to avoid collisions. Such scheduling techniques, however, limit the number of tags that can be supported and, consequently, the scalability of the system. TDOA-based approaches allowing tags to carry out passive self-localization exist [28], in which synchronized anchors subsequently send signals that are received by a tag to estimate the time differences. However, besides the need for a tight ns-range synchronization between anchors (which is hard to achieve [43], and increases message overhead [41, 44]), one needs to correct the tags’ clock skew due to the long reception phase of the sequential messages.

2.3 Concurrent Ranging

Corbalán and Picco [6] have recently introduced the concurrent ranging primitive, which enables the simultaneous distance estimation between an initiator and an arbitrary number of responders. By doing so, concurrent ranging potentially allows to reduce the number of messages required to estimate the distance from N neighbors to a single transmit and receive operation.

Channel impulse response (CIR). Concurrent ranging exploits the CIR estimated by standard-compliant UWB transceivers, such as the Decawave DW1000, to extract simultaneous responses from an arbitrary number of nodes. Fig. 1a shows an exemplary CIR estimated with the DW1000 radio in an environment with limited multipath: one can clearly note the first path or line-of-sight (LOS) component. The latter is typically used to precisely estimate the arrival time of a packet (and consequently the distance between two nodes). A CIR further contains information about the multipath propagation consisting of reflections from surfaces as well as scattering. This feature has been exploited, among others, to derive the presence of NLOS conditions [31], destructive interference [13], as well as to perform multipath-assisted single-anchor localization [14, 27].

Principle of operation. To perform concurrent ranging, an initiator broadcasts an INIT message to all its neighbors (responders), who answer simultaneously with a RESP message after a constant delay Δ_R , as shown in Fig. 2a. The principle foresees the computation of the distance to the closest neighbor using single-sided two-way ranging. Thus, it is assumed that the timestamp included in the payload of the closest neighbor’s RESP message is reliably detected. After completing this step, one can estimate the distance to all other responders by analyzing the CIR of the received RESP message.

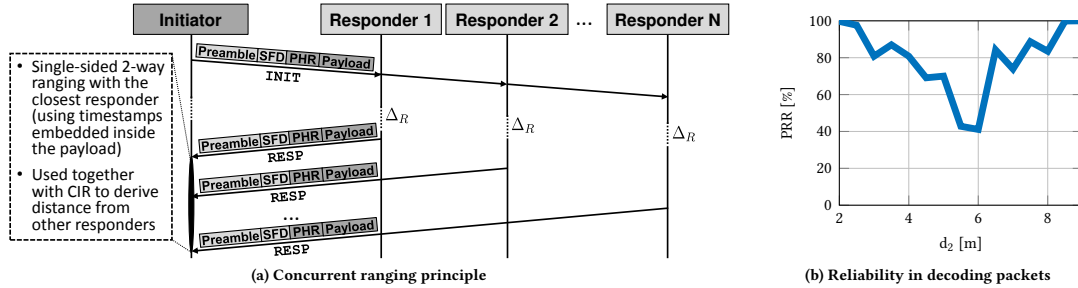


Figure 2: Concurrent ranging principle (a) and reliability in decoding packets when keeping a first responder at a fixed distance $d_1 = 5$ meters and varying the distance d_2 of a second responder (b). Up to 60% of the packets are not received correctly.

Limitations. Unfortunately, concurrent ranging is currently inapplicable in practical UWB systems due to: (i) the inability to identify responders and to discern them from strong multipath components, (ii) the high amount of payloads lost when responders are located at similar distances, as well as (iii) the limited transmit timestamp resolution of common UWB radios, as discussed next.

Identifying responders. One of the key challenges hindering the feasibility of concurrent ranging in real-world systems is the inability to associate a distance estimate to a specific responder. Corbalán and Picco have shown the feasibility of concurrent ranging in artificial setups where the initiator had prior knowledge about the order in which the signal peaks associated to the responders are received in the CIR [6]. In practical situations, however, one does not know the relative locations of nodes beforehand [12].

Discerning responses from strong multipath components. In multipath-rich indoor environments, several strong multipath components (MPCs) may appear in the CIR and overlap with the concurrent responses. Fig. 1b, for example, shows a CIR estimated by the DW1000 radio in a University office while a single responder answers with a RESP message. One can clearly identify five peaks, four of which correspond to MPCs that have an amplitude similar to the LOS component. The presence of such strong MPCs makes it impossible to differentiate between a desired response and a strong multipath component due to reflections from walls and solid surfaces.

Unreliability of correctly decoding timestamps. Concurrent ranging assumes that the timestamp included in the payload of the closest neighbor's RESP message is reliably received. However, the probability to lose a packet or to decode a corrupted payload is very high when one or more responders are located at similar distances. To illustrate this problem, we perform an evaluation in a corridor using one initiator node and two concurrent responders R_1 and R_2 , all using the DW1000 radio. R_1 is placed at a fixed position whose distance from the initiator is $d_1 = 5$ m. We execute different measurements while varying the distance of R_2 between $d_2 = 2, \dots, 9$ m in steps of 50 cm. For each step, we perform 1000 concurrent rangings and log the number of RESP messages successfully decoded at the initiator, which we denote as packet reception rate (PRR). Fig. 2b shows the PRR as a function of d_2 : concurrent ranging as suggested in [6] does not perform reliably when two responders are close to each other. In practice, the PRR would decrease even further if more than two responders are located at a similar distance.

Limited transmit timestamp resolution. As shown in Fig. 2a, all responders dispatch a RESP message after a constant delay Δ_R . To this end, one can use the *delayed transmission* feature of the Decawave DW1000 radio. The latter enables to set a future timestamp at which the transceiver actually sends a RESP message. This allows to align a pre-calculated timestamp with the real transmit timestamp and embed it in the message being transmitted. Unfortunately, the Decawave DW1000 ignores the low-order 9 bits of the timestamp, limiting the transmission resolution to approximately 8 ns [9, p. 26]. This is not an issue in the classical single-sided two-way ranging scheme, as the real transmit timestamp is anyway embedded in the message. However, this aspect has a severe impact on the precision of concurrent ranging, as it negatively affects the concurrency of the RESP messages of the neighbors.

SnapLoc mitigates the aforementioned limitations of concurrent ranging and applies a modification of the latter to a TDOA-based approach, allowing the creation of an indoor localization system that scales regardless of the tags density, as elaborated in Sect. 3.

3 SNAPLOC: DESIGN RATIONALE

In SnapLoc, we tackle the limitations of concurrent ranging and allow tags to reliably obtain and identify simultaneous responses from multiple static anchors. To this end, we assign an individual delay in the nanosecond range to each anchor, which avoids misclassification of responses due to overlapping responses or multipath components (Sect. 3.1). This allows tags to derive the TDOA between anchors by *only* reading and analyzing the CIR. Hence, it removes the need to carry out a single-sided two way ranging and to correctly receive the timestamp embedded in a RESP payload (Sect. 3.2). We finally show how embedding these key principles into a TDOA-based anchor-initiated approach allows to create a scalable UWB-based localization system (Sect. 3.3).

3.1 Correctly Identifying Multiple Responses

As discussed in Sect. 2.3, concurrent ranging fails in situations where responders are located at a similar distance from the tag. Furthermore, in multipath-rich indoor environments, several strong MPCs may be present and overlap with responses from the anchors, making it hard to correctly recognize desired anchor responses.

To address this problem, instead of making use of just a fixed Δ_R as in Fig. 2a, we set an additional individual delay δ_i for each




Figure 3: Tag initiating the position estimation by sending an INIT message to the surrounding anchors, who respond simultaneously (a). Based *only* on the CIR embedded in the response, the tag can derive the TDOA between anchors (b).




Figure 4: Introducing an additional individual delay δ_i for each responder A_i allows to identify responses and discern them from strong multipath components.

responder A_i in the nanosecond range. Consequently, the anchors do not respond simultaneously, but rather quasi-simultaneously. This allows to obtain responses that are separated in time and spread over a wider range of the CIR, as well as to avoid the overlap of MPCs and desired responses. Fig. 4 shows the resulting CIR with four responders: although the first response exhibits a peak due to a strong MPC, it is possible to distinguish it from the other responses, thanks to the additional individual delay δ_i .

3.2 Exploiting CIR Information Only

The approach described in Sect. 3.1 allows tags to seamlessly derive the TDOA between anchors by *only* reading and analyzing a *single* CIR – a novel approach allowing ultra-fast TDOA estimations. Fig. 3a illustrates a scenario with four anchors A_i ($i = 1 \dots 4$) and one tag T . The latter broadcasts an INIT message that is received by all anchors (solid arrows), which simultaneously respond with a RESP message after a constant delay $\Delta_R + \delta_i$ (dashed arrows).

Due to the individual delay δ_i and the different TOF, the responses in the CIR are separated in time, as shown conceptually in Fig. 3b. The distances of the responses $\Delta\tau_{i,j}$ in the CIR contain position-related information of the tag, namely the TDOA between the anchors A_i and A_j ($i \neq j$):

$$\Delta\tau_{i,j} = \delta_j - \delta_i + 2 \cdot (t_{j,T} - t_{i,T}). \quad (1)$$

Given that the individual delay δ_j of anchor A_j is known, the time difference of arrival $\Delta t_{i,j}$ follows as:

$$\Delta t_{i,j} = t_{j,T} - t_{i,T} = \frac{\Delta\tau_{i,j} - (\delta_j - \delta_i)}{2}. \quad (2)$$

Note that this approach removes the need to carry out a single-sided two way ranging and to correctly receive the timestamp embedded in a RESP payload – one of the key limitations outlined in Sect. 2.3. Therefore, as highlighted in Fig. 3b, one can estimate the TDOA between anchors using *only* information that is contained in the CIR estimated from a *single* read operation.

3.3 Allowing the System to Scale

The novel approach described in Sect. 3.2 allows an ultra-fast estimation using only information contained in a single CIR. In principle, by having the tag initiating the localization, this approach allows tags to trigger a position update individually and aperiodically¹. However, it requires a tag to initiate the location estimate by actively sending an INIT message. In order to avoid collisions between tags, one would hence still need to allocate *specific timeslots* to each tag, as well as elect one anchor responsible to periodically broadcast the fixed position and ID of all involved anchors. This would decrease the scalability of the system, as described in Sect. 2.2.

Therefore, we design SnapLoc as an *anchor-initiated* approach in which an anchor is selected to act as the initiator broadcasting the INIT message (called *reference anchor* or A_{ref} in the remainder of this paper). The key advantage of such an anchor-initiated approach is that the tag is not actively involved in the communication and thus no scheduling between multiple tags is required. Furthermore, similar to Global Navigation Satellite Systems (GNSS), this approach allows *passive self-localization*. This enables tags to remain anonymous and maximize their privacy, as well as to achieve a high scalability regardless of the tag density.

4 SNAPLOC: INNER WORKING MECHANISMS

SnapLoc consists of two types of nodes: *anchors* and *tags*. N anchors are placed at known positions $\mathbf{a}^{(i)} \in \mathbb{R}^3$ (with $i = 1, \dots, N$) to

¹Furthermore, by overhearing the INIT message and the anchors' responses it is possible to compute the position of other tags or the position of all tags at a central entity, which is valuable for smart factories as well as people- and asset-tracking [21].

localize N_t tags located at unknown positions $\mathbf{p}^{(n)} \in \mathbb{R}^3$ (with $n = 1, \dots, N_t$). One of the anchors, A_{REF} , is selected as reference to broadcast the INIT message, as described in Sect. 3.3.

We discuss next how to estimate the unknown positions of the tags $\mathbf{p}^{(n)}$. We first assign an individual delay δ_i to anchors in order to avoid misclassification of responses (Sect. 4.1). We then present a mechanism to reliably detect responses within a CIR (Sect. 4.2), show how to derive the TDOA from the detected responses (Sect. 4.3), and how to use this information to estimate the position of the tags $\mathbf{p}^{(n)}$ (Sect. 4.4). We finally describe SnapLoc's implementation on a low-cost UWB platform and present a clock correction scheme to support constrained anchors (Sect. 4.5).

4.1 Setting Individual Anchor Delays

To avoid the overlap of anchor responses and MPCs in the CIR, we suggest to use an individual delay δ_i at each anchor to separate the responses in time, as discussed in Sect. 3.1. Due to the limited length of the CIR register in common UWB transceivers, there is a trade-off between how much the anchors' responses can be separated in time (i.e., the ability to avoid overlaps between strong MPCs and actual responses), and the supported number of anchors. For example, the DW1000 radio limits the CIR to a maximum length of 1016 samples with a sampling period of $T_s = 1.0016$ ns [9].

In SnapLoc, we set the individual delay $\delta_i = (i - 1) \cdot \alpha$, where α represents the size of the slot assigned to each anchor. We suggest to use $\alpha = 128$ ns, which relates to a distance offset of ≈ 38.5 m and makes it very unlikely that a strong MPC of an earlier response interferes with the current response². This allows to use up to eight anchors when using the DW1000 transceiver [9]. In case this anchor density is insufficient, one needs to reduce α to increase the number of supported anchors. In multipath-rich environments, this is not advisable, and we suggest instead to support an unlimited number of anchors using a cellular approach similar to the one employed in mobile networks, e.g., GSM. Instead of multiple frequencies, one can use orthogonal *preamble codes* between neighboring cells, which enables the re-use of slots in the channel impulse response³.

Note that the use of an individual delay δ_i to separate the responses of a CIR in time is, in spirit, similar to the one proposed by Großwindhager et al. [12]. However, in that solution, slots are assigned to mobile tags: this highly limits the number of users that can be supported and hence the scalability of the system, even when using techniques such as *pulse shaping*. In SnapLoc, instead, we allocate slots to anchors, whose number is limited and known beforehand, which allows to keep the overall design simple.

4.2 Reliable Response Detection

Reliably detecting anchor responses in the CIR is key to achieve a high performance. To this end, in SnapLoc we follow these steps:

- (1) Upsample the estimated CIR denoted as \mathbf{r} using fast Fourier transform by a factor of $L = 30$. This improves the time granularity for further processing.

²This holds true also for large areas as the MPCs are attenuated due to path- and reflection loss and will have a negligible impact on the response of the next anchor.

³According to the IEEE 802.15.4 standard [39], up to 24 orthogonal preamble codes can be used to extend SnapLoc with this approach. The implementation of such an extension is, however, out of the scope of this paper and left as future work.




Figure 5: In SnapLoc, a reference anchor A_{REF} broadcasts an INIT message, to which all surrounding anchors reply (quasi-)simultaneously with an empty RESP message.

- (2) Use the estimated channel impulse response to compute the matched filter output $\mathbf{y} = \mathbf{h}_{MF} * \mathbf{r}$, where $*$ marks the discrete convolution and \mathbf{h}_{MF} is the time-discrete impulse response of the matched filter. The latter is defined as the time-reversed transmitted pulse shape $s(t)$ [10], which is derived in a measurement campaign according to [14]. This operation optimizes the signal-to-noise ratio of \mathbf{r} .
- (3) Within each slot i defined by the individual delay δ_i described in Sect. 4.1, the first sample m_i of the matched filter output \mathbf{y} exceeding a certain threshold TH indicates the first path of each anchor response. The threshold TH is chosen experimentally as the 10-fold power of the noise floor.
- (4) The estimated time difference of the responses $\Delta \hat{\tau}_{i,j}$ is determined by $\Delta \hat{\tau}_{i,j} = (m_j - m_i) \cdot (T_s/L)$.

4.3 Deriving Time Difference of Arrival

As discussed in Sect. 3.3, SnapLoc employs an *anchor-initiated* approach, where a reference anchor A_{REF} broadcasts an INIT message. The remaining anchors (marked as $A_1 \dots A_4$ in Fig. 5) respond simultaneously with a RESP message after a delay $\Delta_R + \delta_i$ (with $i = 1 \dots 4$). A nearby tag T can listen to the signals sent from the anchors and detect the responses in the CIR using the algorithm described in Sect. 4.2. Similarly to the approach discussed in Sect. 3.2, the responses encode information related to the time difference of arrival between the anchors. However, due to the different setup, the distances of the responses $\Delta \tau_{i,j}$ follow as:

$$\Delta \tau_{i,j} = (\delta_j - \delta_i) + (t_{R,j} + t_{j,T}) - (t_{R,i} + t_{i,T}). \quad (3)$$

Due to the static nature of the anchors, $t_{R,i}$ and $t_{R,j}$, respectively, are known, and the TDOA $\Delta \tau_{i,j}$ of the anchors A_i and A_j follows as:

$$\Delta \tau_{i,j} = t_{j,T} - t_{i,T} = \Delta \tau_{i,j} - (\delta_j - \delta_i) - t_{R,j} + t_{R,i}. \quad (4)$$

Selection of reference anchor. In principle, any anchor within the communication range and optimally in line-of-sight of all other anchors in the same area can be selected as reference (A_{REF}). The selection of an anchor as initiator allows tags to *self-localize*, as discussed in Sect. 3.3. Furthermore, it also increases the robustness of the localization system. Indeed, anchors are typically installed

in corners and well above objects in a room. Thus, it is less likely that there is a degraded link between anchors. For this reason, the probability to lose the *INIT* message is lower in the *anchor-initiated* approach (Fig. 5) than with the *tag-initiated* one (Fig. 3).

Broadcast anchor positions. As for every anchor-based system allowing self-localization of tags, also in SnapLoc a tag needs to know the ID and location of the anchors to compute its position. To avoid the need of additional infrastructure or packet exchanges, we propagate (i) the ID of the anchors, (ii) their individual delay δ_i , and (iii) their position within the *INIT* message sent by the reference anchor. Furthermore, the *INIT* message contains the initialization interval T_{init} between two consecutive *INIT* messages, as well as transmit timestamp correction values as discussed in Sect. 5.

4.4 Localization Algorithm

As described in Sect. 4.3, the time difference of arrivals between the anchors are derived from the CIR. Based on these estimates, we are able to directly derive the unknown position of the tags $\mathbf{p}^{(n)}$ using TDOA trilateration. For simplification and due to space limitations, we tackle in this section just the two-dimensional case (\mathbb{R}^2) and a single tag (i.e., $N_t=1$) at position $\mathbf{p}^{(1)} = \mathbf{p} = [x, y]^T$. The anchor nodes are positioned at $\mathbf{a}^{(i)} = [x_i, y_i]^T$ (with $i = 1, \dots, N$). The distance d_i between the tag and an anchor A_i is defined by:

$$d_i = \sqrt{(x_i - x)^2 + (y_i - y)^2} \quad (5)$$

Therefore, the distance differences between anchors Δd_{ij} (with $i \neq j$) – derived by multiplying the time difference of arrivals $\Delta t_{i,j}$ with the propagation speed c – is:

$$\Delta d_{i,j} = \sqrt{(x_i - x)^2 + (y_i - y)^2} - \sqrt{(x_j - x)^2 + (y_j - y)^2}. \quad (6)$$

The use of N anchors results in $N - 1$ non-redundant nonlinear equations. In the two-dimensional space, at least $N = 4$ anchors are required, i.e., three non-redundant equations, for finding the unambiguous position of a tag [2]. Even with $N = 4$ anchors, just with zero measurement noise we are guaranteed to get a single solution, which corresponds to the real tag position \mathbf{p} . Adding white Gaussian measurement noise \mathbf{n} results in the signal model in vector notation in Eq. (7), expressing the relationship between the position of anchor/tag and the estimated time difference of arrivals $\Delta \hat{t}_{i,j}$. We obtain the latter by applying equation (4) to the estimated time differences of the responses $\hat{\tau}_{i,j}$ derived from the CIR (see Sect. 4.2). Please note that we relate the TDOA estimates to the first anchor.

$$\hat{\mathbf{d}} = \mathbf{s}(\mathbf{p}) + \mathbf{n} \quad (7)$$

with

$$\mathbf{s}(\mathbf{p}) = \begin{bmatrix} \Delta d_{2,1} \\ \vdots \\ \Delta d_{N,1} \end{bmatrix} = \begin{bmatrix} \sqrt{(x_2 - x)^2 + (y_2 - y)^2} - \sqrt{(x_1 - x)^2 + (y_1 - y)^2} \\ \vdots \\ \sqrt{(x_N - x)^2 + (y_N - y)^2} - \sqrt{(x_1 - x)^2 + (y_1 - y)^2} \end{bmatrix} \quad (8)$$

and the observation vector

$$\hat{\mathbf{d}} = c \cdot [\Delta \hat{t}_{2,1}, \Delta \hat{t}_{3,1}, \dots, \Delta \hat{t}_{N,1}]^T.$$




Figure 6: Low-cost UWB platform based on a Decawave DW1000 radio with an omni-directional dipole antenna.

Based on Eq. (7), the nonlinear least squares (NLS) cost function $J_{NLS}(\hat{\mathbf{p}})$ follows as [38]:

$$J_{NLS}(\hat{\mathbf{p}}) = (\hat{\mathbf{d}} - \mathbf{s}(\hat{\mathbf{p}}))^T (\hat{\mathbf{d}} - \mathbf{s}(\hat{\mathbf{p}})). \quad (9)$$

Therefore, the NLS position estimate follows as:

$$\hat{\mathbf{p}} = \arg \min_{\mathbf{p}} J_{NLS}(\mathbf{p}) = \arg \min_{\mathbf{p}} (\hat{\mathbf{d}} - \mathbf{s}(\mathbf{p}))^T (\hat{\mathbf{d}} - \mathbf{s}(\mathbf{p})). \quad (10)$$

To find $\hat{\mathbf{p}}$, we use the quasi-Newton method [36] with an initial position estimate chosen at the center of the room.

4.5 Implementation

We implement SnapLoc on a low-cost UWB platform described in Sect. 4.5.1. In general, the hardware employed to build a localization system severely affects the minimum response delay ΔR that can be used, as discussed in Sect. 4.5.2. This affects the update rate that can be achieved by SnapLoc, and raises the need for a simple clock correction scheme at the anchor, which we present in Sect. 4.5.3.

4.5.1 Hardware. We employ a self-made low-cost UWB platform based on the IEEE 802.15.4-compliant DW1000 transceiver for both anchors and tags. The platform, shown in Fig. 6, is based on the STM32 Nucleo-64 board, which employs an ultra-low power ARM Cortex-M3 based STM32L152RE microcontroller. The RF front-end contains a low-cost EPSON TSX-3225 oscillator with a frequency of 38.4 MHz, a tolerance of 10 ppm, and with no temperature compensation. To overcome quartz imperfections and varying temperatures we used a built-in functionality of the DW1000 to tune the clock of the anchors depending on the *INIT* message received from the reference anchor. Every node uses an off-the-shelf omni-directional UWB dipole antenna. The DW1000 is configured to use channel 4 (i.e., a bandwidth of 900 MHz and a carrier frequency of 3.9936 GHz), maximum data rate (6.8 Mbps), a pulse repetition frequency of 64 MHz, as well as a preamble symbol repetition of 128.

4.5.2 Minimum response delay. In SnapLoc, anchors respond to an *INIT* message broadcasted by the reference anchor after a delay $\Delta R + \delta_i$. Since the first symbol of the physical header (PHR) determines the transmit timestamp [9], the minimum applicable response delay ΔR_{min} is defined by the duration of PHR and payload of the *INIT* message, as well as the duration of preamble and start-of-frame-delimiter (SFD) of the *RESP* message. Overall, this corresponds to a delay of 178.5 μs . Additionally, we also need to account for the minimum time necessary to switch the DW1000 radio from receive to transmit mode. We have evaluated this minimum switching time (due to the SPI communication overhead and delays introduced by the processing of a packet) experimentally using

two different platforms. The powerful Decawave EVB1000 board, which embeds the STM32F105 ARM Cortex M3 microcontroller, exhibits a minimum switching time of roughly $100 \mu\text{s}$, which results in $\Delta_{R,min} = 278.5 \mu\text{s}$. When employing a more constrained microcontroller with lower CPU and SPI speed to control the DW1000 (e.g., the UWB platform presented in Sect. 4.5.1), the minimum response delay increases to $\Delta_{R,min} = 850 \mu\text{s}$. This delay affects the update rate that can be achieved by SnapLoc, as shown in Sect. 6.2.

4.5.3 Clock correction. SnapLoc requires that all anchor nodes send their RESP message at well-defined time instances. Therefore, variations of the response delay Δ_R due to imperfections of low-cost oscillators driving the UWB transceiver can potentially degrade the performance of the system. This problem is exacerbated when using a highly constrained hardware causing a large minimum response delay $\Delta_{R,min}$, as discussed in Sect. 4.5.2. Thus, to allow a flexible selection of the response delay Δ_R , we suggest a simple technique to correct the response time $t_{k,i}^{TX,R}$ at each anchor A_i .

Consider that the oscillators of the reference anchor and a fixed anchor A_i are running at different speeds due to imperfections, i.e., also the reported time $C_{REF}(t_k)$ and $C_i(t_k)$ vary. The relative skew $a_{REF,i}$ between them can be calculated as [17]:

$$a_{REF,i} = \frac{C_i(t_{k+1}) - C_i(t_k)}{C_{REF}(t_{k+1}) - C_{REF}(t_k)}, \quad (11)$$

where $C_{REF}(t_k) = t_k^{TX,I}$ is the transmission time of the k^{th} INIT message and $C_i(t_k) = t_{k,i}^{RX,I}$ denotes the reception time of the k^{th} INIT message at the anchor A_i neglecting the time of flight. In SnapLoc, the reference anchor broadcasts the INIT message with the interval $T_{init} = t_{k+1}^{TX,I} - t_k^{TX,I}$. Eq. (11) hence follows as:

$$a_{REF,i} = \frac{t_{k+1}^{RX,I} - t_k^{RX,I}}{T_{init}}. \quad (12)$$

The common response time Δ_R and the individual anchor delay δ_i are defined in the common time of the reference node. Thus, they have to be brought into the time domain of the corresponding anchor A_i using the relative skew $a_{REF,i}$ between them. The corrected transmit time $t_{k,i}^{TX,R}$ of the RESP message at A_i follows as:

$$t_{k,i}^{TX,R} = t_{k,i}^{RX,I} + a_{REF,i} \cdot (\Delta_R + \delta_i). \quad (13)$$

5 IMPROVING TIMESTAMP RESOLUTION

To implement SnapLoc on the UWB Decawave DW1000 transceiver, we employ the *delayed transmission* feature. The latter allows to program a future timestamp in a register and lets the DW1000 initiate a packet transmission at this defined timestamp. In SnapLoc, this allows each anchor to set the timestamp at which the RESP message needs to be transmitted upon reception of the INIT message. Although the DW1000 radio represents receive (RX) and transmit (TX) timestamps as 40-bit values with a resolution of 15.65 ps [9], it ignores the lower 9-bit when performing delayed transmissions. This lowers the effective transmission resolution from (theoretical) 15.65 ps to $4/(499.2 \cdot 10^6) \approx 8 \text{ ns}$. Without correction, in SnapLoc, this transmission uncertainty results in a uniformly distributed and memoryless error $e_{TS} \sim \mathcal{U}(-8 \text{ ns} \cdot c, 0)$. Considering that an error of 1 ns in the time domain results in an error of $\approx 30 \text{ cm}$ in the distance domain, it is evident that this error highly affects the

localization performance, as we show experimentally in Sect. 6.3. Thus, to sustain a decimeter-level accuracy in SnapLoc, we propose two techniques to increase the transmit timestamp resolution.

Wired correction. We first propose an optimal correction scheme that tracks the lost 9-bit at each anchor and sends these correction values back via a wired backbone to the reference anchor A_{REF} . Such a wired connection is typically available in localization systems, in order to power, reprogram, and reconfigure the anchors. In this scheme, the reference anchor broadcasts the missing transmit timestamp information in the next INIT message to all tags. The latter then correct the timestamps of the anchor responses derived from the previous CIR. In this way, the correction does not require additional messages to be transmitted, as the correction values are embedded in the INIT message. Nonetheless, the tag applies the correction values sent in the latest INIT message to correct the timestamps of the previous position estimate, which causes a delay by one initialization interval T_{init} . Due to the high update rate of SnapLoc, this trade-off is tolerable, as discussed in Sect. 6.2.

Wireless correction. In case a backbone network is not available, we propose a second scheme to increase the timestamp resolution that does not require a wired connection between anchors and the reference anchor. In principle, so far, the latter was used to initiate a position estimation by sending an INIT message and could act as a regular anchor by responding to its own initialization message. In the *wireless correction* scheme, instead, the reference anchor listens to the responses of the anchors and derives the estimated CIR. As the anchors are static and their positions are known, the distance information estimated from the CIR can be compared with the true values. Deviations of the estimations from the true values are treated as errors due to ignoring the least significant 9-bits in the transmit timestamp. To recover the lost precision, we differentiate between the correction at anchor A_1 and the remaining anchors. This is due to the fact that, in SnapLoc, anchor A_1 has an individual time delay $\delta_1 = 0$ and its response hence corresponds to the first peak in the CIR. Thus, the timestamp of its response $t_{RESP,1}^{RX}$ is detected with the highest possible resolution of 15.65 ps by the embedded leading edge detection of the DW1000 [9]. Instead, the resolution of the remaining anchor responses is limited by the sampling period $T_s = 1.0016 \text{ ns}$ of the CIR (see Sect. 4.1). For anchor A_1 , we define the transmit error due to the limited timestamp resolution $e_{A_1}^{TX}$ as the difference between the true round trip time t_{RT} of A_1 and A_{REF} and the estimated one \hat{t}_{RT} :

$$e_{A_1}^{TX} = t_{RT} - \hat{t}_{RT}. \quad (14)$$

The true round trip time t_{RT} is defined by

$$t_{RT} = 2 \cdot t_{ref,1} + \Delta_R + \delta_1 + 2 \cdot \Theta_a, \quad (15)$$

where $t_{ref,1}$ is the time of flight between reference node and A_1 , Δ_R the common response delay at all anchor nodes, δ_1 the individual response delay of A_1 , and Θ_a an antenna delay. The latter is required to correct for delays introduced by the antenna, PCB, and internal and external components [9, p.205 ff.]. To measure the antenna delay Θ_a , we have performed 5000 two-way ranging trials between two nodes placed 3 m apart from each other. The antenna delay Θ_a is calibrated such that the difference between the reported distance and the true distance $d_0 = 3 \text{ m}$ is minimized. The estimated

round trip time \hat{t}_{RT} is determined by the difference between the timestamp $t_{RESP,1}^{RX}$ of A_1 's response and the transmission time of the INIT message at the reference anchor t_{INIT}^{TX} . Therefore, the TX timestamp error of A_1 follows as:

$$e_{A_1}^{TX} = (2 \cdot t_{ref,1} + \Delta_R + \delta_1 + 2 \cdot \Theta_a) - (t_{RESP,1}^{RX} - t_{INIT}^{TX}). \quad (16)$$

The transmit timestamp resolution error of the remaining anchors $e_{A_i}^{TX}$ ($i = 2, \dots, N$) is defined as the true TDOA $\Delta t_{i,1}$ between A_i and A_1 and the one estimated from the CIR $\hat{\Delta t}_{i,1}$:

$$e_{A_i}^{TX} = \Delta t_{i,1} - \hat{\Delta t}_{i,1}. \quad (17)$$

The true TDOA $\Delta t_{i,1}$ is derived from the known positions of the reference node and anchors and follows as:

$$\Delta t_{i,1} = t_{ref,i} - t_{ref,1} \quad (18)$$

where $t_{ref,i}$ is the time of flight between A_{REF} and A_i . The estimated TDOA $\hat{\Delta t}_{i,1}$ is derived from the CIR according to (2) and has to be corrected by the previously acquired transmit error of A_1 $e_{A_1}^{TX}$. Thus, the resulting error of the anchor A_i is:

$$e_{A_i}^{TX} = 2 \cdot (t_{ref,i} - t_{ref,1}) - (\hat{\Delta t}_{i,1} + e_{A_1}^{TX}). \quad (19)$$

As discussed, the resolution of the error value $e_{A_i}^{TX}$ is restricted by the sampling period of the CIR $T_s = 1.0016$ ns. Thus, 3-bits in the INIT message broadcasted by the reference anchor are enough to represent the error correction value. Therefore, the overhead due to a longer packet size is slightly shorter in the *wireless correction* method compared to the *wired correction*.

6 EVALUATION

We evaluate SnapLoc experimentally in a challenging office environment (Room A, see Fig. 7a) and a larger laboratory classroom (Room B, see Fig. 7b). We describe the experimental setup in Sect. 6.1, followed by an analysis of the energy consumption in terms of over-the-air time and the potential update rate in Sect. 6.2. We then extensively evaluate the performance of SnapLoc in Sect. 6.3, showing that it can achieve decimeter-level localization accuracy.

6.1 Experimental Setup

To evaluate SnapLoc in a realistic indoor environment, we use a common office for three employees with a size of 5.2×6.03 m ≈ 31.36 m² (see Fig. 7a) and a larger laboratory classroom with 6.05×10 m = 60.5 m² (see Fig. 7b). The rooms contain several scattering and reflecting objects such as monitors, desks, and chairs. The reference anchor (magenta square) and the remaining anchors (blue squares) are placed on tripods at known positions. The height of all tripods is 1.60 m, which puts all nodes in the same 2D plane. For all evaluations, we employ just the *minimum amount of anchors necessary*, i.e., $N = 4$: this allows to examine the performance of SnapLoc using just minimal infrastructure. The number of evaluation points ($N_{EP} = 28$ in Room A and $N_{EP} = 14$ in Room B) are randomly distributed in the rooms to evaluate the performance of SnapLoc. At each evaluation point, $N_P = 500$ position estimates are derived. The absolute error of each trial is calculated as the Euclidean distance between the position of the evaluation point \mathbf{p}_{EP} and the i -th position estimate $\hat{\mathbf{p}}_i$:

$$\text{Err}_i = \|\hat{\mathbf{p}}_i - \mathbf{p}_{EP}\|. \quad (20)$$

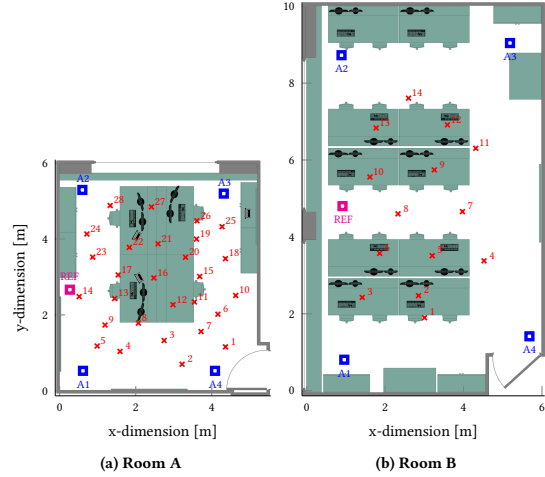


Figure 7: Evaluation setup: we consider two different environments with 28 and 14 evaluation points, respectively.

6.2 Position Update Rate and Efficiency

Due to the high current consumption of the DW1000 in the transmit and especially in the receive mode [13], it is critical for UWB-based localization systems to minimize the radio-on time at the tag. Due to the simultaneous acquisition of all the anchor signals, SnapLoc excels in this regard. Indeed, the tag does not have to send any packet, but just listens to a single message. Thus, the number of anchors does not affect the system's energy consumption in terms of packet reception and transmission. This is in contrast to state-of-the-art UWB-based localization systems, where the energy consumption increases – typically linearly – with the number of anchors [24, 28]. We measure the energy consumption of SnapLoc with the settings described in Sect. 4.5.1 with a Keysight MSOS-254A oscilloscope. Acquiring the simultaneous anchor responses requires only approximately 82.4 μ J. Besides a low energy consumption, simultaneously responding anchors also highly affect the achievable position update rate, as the latter relates to the total time needed to provide the tag with the necessary information to estimate its position. In SnapLoc, this total time consists of the duration of INIT and RESP messages, as well as the time to switch between receive and transmit mode at the anchors. As discussed in Sect. 4.5.2, this switching time is approximately 100 μ s when using the Decawave EVB1000 board and the duration of the two messages is roughly 334 μ s. Thus, deriving the information to estimate the tag's position just takes 434 μ s overall. Theoretically, this enables an update rate of more than 2.3 kHz for SnapLoc, without any limitation on the number of tags. Even when using the highly constrained microcontroller with low SPI and CPU speeds described in Sect. 4.5.1, we still achieve an update rate of about 996 Hz. This high update rate makes SnapLoc highly suitable for feedback control systems and enables the precise tracking of highly-dynamic objects. Note that the update rate is also influenced by (i) streaming the CIR via SPI from the DW1000, (ii) deriving the actual TDOAs, as well as (iii) executing the algorithm to estimate the tag's position. However, these values are strongly




Figure 8: Error ellipses showing the bias (blue circles) and the standard deviation (black ellipses) of the position estimation without correction of the transmit timestamp (a), with the wireless correction (b), and with the wired correction (c).

hardware-dependent and could significantly be reduced by integrating a UWB transceiver together with a performant microcontroller in a system on chip solution. Furthermore, when using the techniques proposed in Sect. 5 to increase the timestamp resolution of the DW1000, the uncertainties of the timestamps have to be either sent back via wire (*wired correction*) or estimated at the reference anchor (*wireless correction*), which decreases the update rate.

6.3 Localization Accuracy

We evaluate next the performance of SnapLoc and the effectiveness of the methods to overcome the limited transmit timestamp resolution proposed in Sect. 5.

Individual evaluation points. We start by investigating the localization accuracy of SnapLoc in a smaller room (Fig. 7a) using $N_{EP} = 28$ evaluation points. Fig. 8 shows the impact of the transmit timestamp correction techniques presented in Sect. 5. The mean (blue circle) and the standard deviation (black ellipses) for $N_P = 500$ position estimates are shown for each evaluation point. Fig. 8a shows the accuracy of SnapLoc's position estimation *without transmit timestamp correction*. Fig. 8b shows the accuracy of SnapLoc's position estimation with the *wireless correction*, whilst Fig. 8c with the *wired correction*. As expected, the latter performs best, as it recovers the least significant 9-bits of the transmit timestamp at all anchors. The *wireless correction*, instead, restores a time resolution of 15.56 ps for anchor A_1 and a resolution of 1 ns for the remaining anchors; thus, its performance is slightly worse compared to the one obtained with the *wired correction*. Without any correction, each of the anchor timestamps has a resolution of just 8 ns, which induces a high error, as shown by the larger ellipses in Fig. 8a. Moreover, it is noticeable in Fig. 8b and 8c that the evaluation points within a distance of 1.5 m to an anchor ($EP \in \{1, 5, 24, 25, 26, 28\}$) perform worse than those located further away from the anchors. This is due to the high signal strength of the close anchor, which causes the CIR register to saturate. As the amplitude of the other anchors' responses remains relatively low, a correct response detection is

impaired. Thus, when deploying SnapLoc, a distance of at least 1.5 m between the tag and the anchors should be ensured. This is often already the case in indoor localization systems, as anchors are typically mounted close to the ceiling.

Overall localization accuracy and precision. To investigate the overall performance of SnapLoc, we derive its accuracy and precision statistically using the cumulative distribution function (CDF) over the error Err_i of all position estimates. Due to the saturation effects at tag positions close to the anchors, we have ignored the corresponding evaluation points $EP \in \{1, 5, 24, 25, 26, 28\}$ for this analysis. Fig. 9a shows the performance of SnapLoc depending on the used method to correct the limited TX timestamp resolution of the Decawave DW1000. *Without correction* (solid orange line), a 90% error of 1.15 m and a median error of 0.68 m was achieved. Instead, the use of *wireless correction* allows to reduce the 90% error to 55.8 cm and the median error to 25.4 cm (dashed blue line) and the *wired correction* even reaches a 90% error of just 33.7 cm and a median error of 18.4 cm (magenta dash dotted line). Thus, by using the proposed correction methods, SnapLoc easily achieves decimeter-level accuracy despite the limited transmit timestamp resolution of 8 ns and the CIR resolution of about 1 ns.

Performance in larger room. To validate the accuracy of SnapLoc also in other environments, we carry out an evaluation in a laboratory classroom (Fig. 7b) that is significantly larger than the previously employed office room ($31.36 m^2$ vs. $60.5 m^2$). Fig. 9b shows the CDF of all position estimates in the $N_{EP} = 14$ evaluation points shown in Fig. 7b. Without using a transmit timestamp correction, the 90% error is at 1.30 m and the median error at 0.73 m. The *wireless correction* allows SnapLoc to sustain a 90% error of 74 cm and a median error of 22.3 cm. With the *wired correction*, the median error is reduced to 17 cm and the 90% error to 35.2 cm. The slight differences compared to the evaluation in room A are due to the presence of a few more outliers with a position error above 0.5 m, as shown in Fig. 9b. Still, the results are consistent to the evaluation in Room A despite the use of a larger area.




Figure 9: Performance of SnapLoc depending on the method used to correct the limited transmit resolution of the DW1000 transceiver in the two rooms used in our evaluation.




Figure 10: Performance of SnapLoc when a subject follows a pre-defined track (red solid line). The positions estimated by SnapLoc are marked with light blue crosses.

Free movement. The previous evaluations were performed at randomly chosen, but static evaluation points to deliver reproducible results. To investigate the performance of SnapLoc also while moving around freely, we mounted a tag on a rolling stand and asked a subject to follow a marked line in a slow but continuous fashion. While moving, we have continuously estimated the tag’s position using SnapLoc combined with the *wired correction* method. Fig. 10 shows the results of the experiment. It is evident that the position estimates (light blue crosses) follow the pre-defined track (red solid line). Especially in the middle of the room, SnapLoc shows reliable results due to equally strong responses of the anchors. Instead, close to the anchors and at the border of the envelope curve spanned by the anchors, the number of outliers increases. This is coherent with the observations made in the previous evaluations. Please note that we did *not* use any (tracking) filter on the measurements, such as Kalman filter, particle filter, or moving average. The results are solely *raw position estimates*. Due to unavailability of an optical tracking system to provide ground truth data, we could not determine the absolute error properly. Still, estimating the shortest distance to the desired track reveals a mean deviation of just 14.8 cm.

7 RELATED WORK

UWB localization systems. Ultra-wideband technology enables decimeter-level localization accuracy in multipath-rich indoor environments without the need of extensive infrastructure [30, 47]. Several practical implementations using low-cost UWB radios exist, e.g., based on the Decawave DW1000 [3, 14, 24, 28, 32, 37, 41], on Time Domain’s PulsOn module [11, 45], or on self-made hardware [23]. The main objective of these systems is to achieve a *high localization accuracy*: as a result, the update rate at which the position can be computed has often not been discussed. Amongst works explicitly mentioning the supported update rates, Kempke et al. [24] achieve a 99% error in 3D of 53 cm with an update rate of 12 Hz. However, the latter is divided by the number of supported tags (e.g., 6 Hz for two tags). Silva et al. [37] report average errors between 5 and 40 cm in 2D, and achieve an update rate of 10 Hz for a single tag. Hartmann et al. [16] report an average error of 27 cm in 2D and update a single tag every 50 Hz. SnapLoc achieves similar accuracies (90% error of 33 cm), but *at much higher update rates* and with the ability to support an *unlimited number of tags*.

Passive self-localization. One of the main features of SnapLoc is that it gives tags the ability to carry out passive self-localization and remain anonymous (i.e., tags are not actively transmitting data). This allows to build localization systems that scale regardless of the density of tags, given that an unlimited number of tags can, in principle, localize at the same time. Passive self-localization is the same principle adopted by GNSS systems [19] and one of the key reasons for their long-lasting and enduring success. However, GNSS satellites require the use of atomic clocks to maintain synchronization of anchors. An UWB-based system comparable to GNSS has been presented in [28], but it requires a tight synchronization at the anchors and clock skew correction at the tag due to the use of sequential messages, which is often hard to attain [43, 44]. SnapLoc, instead, *removes the need for tight synchronization* and *does not need a correction at the tag* due to the use of (quasi-)simultaneous responses, which ultimately enables very high update rates.

Concurrent passive localization. Similarly to SnapLoc, also Chorus [7] exploits the concept of concurrent transmissions to perform passive localization. Both works are developed independently and published simultaneously in the same venue: differences include implementation details, evaluation methodology, and the (complementary) slant of the contribution. While Chorus focuses on

modeling the impact of the limited timestamp resolution, SnapLoc proposes a technique to overcome this limitation and implements it on platforms making use of state-of-the-art UWB transceivers (Sect. 5): this allows to achieve decimeter-level accuracy, as demonstrated experimentally (Sect. 6) and showcased at public events [40]. SnapLoc also counteracts the clock drift between INIT and RESP messages, enabling also highly-constrained devices (such as the low-cost UWB platform presented in Sect. 4.5.1) to make use of the proposed TDOA-based localization concept.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we present SnapLoc, an ultra-fast localization system for an unlimited number of tags – actually faster than a finger snap, which typically takes 1 to 3 ms [46]. SnapLoc derives simultaneously all the information required to estimate a tag's position, which is enabled by extracting concurrent anchor responses from a single estimated CIR. Based on the detected responses, SnapLoc estimates the TDOA between anchors, removing the need to derive the distance to the closest anchor using multiple messages as in previous solutions, and allowing tags to anonymously self-localize. Furthermore, in contrast to classical TDOA systems, SnapLoc does neither require tight synchronization of anchors, nor correction of clock deviations at the tag. We implement and evaluate SnapLoc experimentally on a low-cost platform based on the Decawave DW1000 UWB radio, as well as mitigate the intrinsic limited transmit timestamp resolution of this transceiver. Our results show that SnapLoc sustains decimeter-level positioning accuracy, with a 90% error of 33.4 cm and a median error of 18.4 cm, and that it is highly suited for supporting mobile applications.

Future work includes the evaluation of SnapLoc in three dimensions, as well as the installation of SnapLoc in a multi-room multi-level building to investigate its performance when introducing multiple cells operating with orthogonal preamble codes.

9 ACKNOWLEDGMENTS

This work was supported by the TU Graz LEAD project "Dependable Internet of Things in Adverse Environments".

REFERENCES

- [1] A. Alarifi et al. 2016. Ultra Wideband Indoor Positioning Technologies: Analysis and Recent Advances. *Sensors* 16, 5 (2016).
- [2] R. M. Buehrer and S. Venkatesh. 2012. Fund. of Time-of-Arrival-Based Position Locations. *Handbook of Position Location: Theory, Practice, and Adv.* (2012).
- [3] W. Chantaweksomboon et al. 2016. On Performance Study of UWB Real Time Locating System. In *Proceedings of the 7th IC-ICTES Conference*.
- [4] Y.K. Cho et al. 2010. Error Modeling for an Untethered Ultra-Wideband System for Construction Indoor Asset Tracking. *Automation in Construction* 19, 1 (2010).
- [5] G. Conte et al. 2014. BlueSentinel: a First Approach using iBeacon for an Energy Efficient Occupancy Detection System. In *Proc. of the 1st ACM BuildSys Conf.*
- [6] P. Corbalán et al. 2018. Concurrent Ranging in UWB Radios: Experimental Evidence, Challenges, and Opportunities. In *Proc. of the 15th EWSN Conference*.
- [7] P. Corbalán et al. 2019. Chorus: UWB Concurrent Transmissions for GPS-like Passive Localization of Countless Targets. In *Proc. of the 18th IPSN '19 Conference*.
- [8] W. Dargie and C. Poellabauer. 2010. *Fundamentals of Wireless Sensor Networks: Theory and Practice*.
- [9] Decawave Ltd. 2017. *DW1000 User Manual. Version 2.13*.
- [10] Chiara Falsi et al. 2006. Time of Arrival Estimation for UWB Localizers in Realistic Environments. *EURASIP Journal on Advances in Signal Processing* 1 (2006).
- [11] Javier González et al. 2009. Mobile Robot Localization based on Ultra-Wide-Band Ranging: A Particle Filter Approach. *Robotics and Auton. Syst.* 57, 5 (2009).
- [12] B. Großwindhager et al. 2018. Concurrent Ranging with UWB Radios: From Experimental Evidence to a Practical Solution. In *Proceedings of the 38th International Conference on Distributed Computing Systems (ICDCS)*.
- [13] B. Großwindhager et al. 2018. Enabling Runtime Adaptation of PHY Settings for Dependable UWB Communications. In *Proceedings of the 19th Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*.
- [14] B. Großwindhager et al. 2018. SALMA: UWB-based Single-Anchored Localization System using Multipath Assistance. In *Proceedings of the 16th ACM International Conference on Embedded Networked Sensor Systems (SenSys)*.
- [15] K. Guo et al. 2016. Ultra-Wideband-Based Localization for Quadcopter Navigation. *Unmanned Systems Journal* 4, 1 (2016).
- [16] F. Hartmann et al. 2015. Design of an Embedded UWB Hardware Platform for Navigation in GPS Denied Environments. In *Proceedings of the IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*.
- [17] J. He et al. 2014. Time Synchronization in WSNs: A Maximum-Value-Based Consensus Approach. *IEEE Trans. Automat. Control* 59, 3 (2014).
- [18] Suining He and S-H Gary Chan. 2016. Wi-Fi Fingerprint-based Indoor Positioning: Recent Advances and Comparisons. *IEEE Comm. Surveys & Tutorials* 18, 1 (2016).
- [19] Bernhard Hofmann-Wellenhof, Herbert Lichtenegger, and Elmar Wasle. 2007. *GNSS—global navigation satellite systems: GPS, GLONASS, Galileo, and more*.
- [20] S. Huang et al. 2017. A Real-time Location System Based on RFID and UWB for Digital Manufacturing Workshop. *Procedia CIRP* 63, 1 (2017).
- [21] L. Jiang et al. 2010. Integrated UWB and GPS Location Sensing System in Hospital Environment. In *Proceedings of the 5th ICIEA Conference*.
- [22] B. Kempke et al. 2015. PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization. In *Proc. of the 2nd HotWireless Workshop*.
- [23] B. Kempke et al. 2016. Harmonium: Asymmetric, Bandstitched UWB for Fast, Accurate, and Robust Indoor Localization. In *Proceedings of the 15th International Conference of Information Processing in Sensor Networks (IPSN)*. IEEE.
- [24] B. Kempke et al. 2016. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proceedings of the 14th ACM Int. Conf. on Embedded Network Sensor Systems (SenSys)*.
- [25] Hakyong Kim. 2009. Performance Comparison of Asynchronous Ranging Algorithms. In *Proceedings of the Global Telecommunications Conference (GLOBECOM)*.
- [26] M. J. Kuhn et al. 2011. A Multi-Tag Access Scheme for Indoor UWB Localization Systems Used in Medical Environments. In *Proc. of the International Conference on Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireless)*.
- [27] J. Kulmer et al. 2017. Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning. In *Proc. of the 5th IEEE ANLN Workshop*.
- [28] A. Ledergerber et al. 2015. A Robot Self-Localization System using One-Way Ultra-Wideband Communication. In *Proceedings of the IROS Conference*.
- [29] R. Leser et al. 2014. Accuracy of an UWB-based Position Tracking System Used for Time-Motion Analyses in Game Sports. *Eur. Journ. of Sport Sc.* 14, 7 (2014).
- [30] D. Lymberopoulos et al. 2017. The Microsoft Indoor Localization Competition: Experiences and Lessons Learned. *IEEE Signal Processing Magazine* 34, 5 (2017).
- [31] S. Marañó et al. 2010. NLOS Identification and Mitigation for Localization Based on UWB Experimental Data. *IEEE J-SAC* 28, 7 (2010).
- [32] Y. Qin et al. 2015. A Distributed UWB-based Localization System in Underground Mines. *Journal of Networks* 10, 3 (2015).
- [33] M. Ridolfi et al. 2018. Analysis of the Scalability of UWB Indoor Localization Solutions for High User Densities. *Sensors* 18, 6 (2018).
- [34] N. C. Rowe et al. 2013. A UWB Transmit-Only Based Scheme for Multi-Tag Support in a Millimeter Accuracy Localization System. In *Proceedings of the IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*.
- [35] J. L. Rullan-Lara et al. 2013. Indoor Localization of a Quadrotor Based on WSN: A Real-Time Application. *Int. Journal of Adv. Robotic Systems* 10, 1 (2013).
- [36] David F Shanno. 1970. Conditioning of Quasi-Newton Methods for Function Minimization. *Math. Comp.* 24, 111 (1970).
- [37] B. Silva et al. 2014. Experimental Study of UWB-based High Precision Localization for Industrial Applications. In *Proc. of the Int. Conf. on Ultra-WideBand (ICUWB)*.
- [38] Hing Cheung So. 2011. Source Localization: Algorithms and Analysis. *Handbook of Position Location: Theory, Practice, and Advances* (2011).
- [39] IEEE Computer Society. 2015. Standard for Low-Rate Wireless Networks.
- [40] M. Stocker et al. 2019. Demo Abstract: SnapLoc: An Ultra-Fast UWB-Based Indoor Localization System for an Unlimited Number of Tags. In *Proceedings of the 18th ACM/IEEE IPSN Conference*.
- [41] J. Tiemann et al. 2016. Atlas: An Open-Source TDOA-based Ultra-Wideband Localization System. In *Proceedings of the IPIN Conference*.
- [42] K. Witrisal et al. 2016. High-Accuracy Local. for Assisted Living: 5G systems will Turn Multipath Channels from Foe to Friend. *IEEE Signal Proc. Mag.* 33, 2 (2016).
- [43] Y. Wu et al. 2011. Clock Synchronization of Wireless Sensor Networks. *IEEE Signal Processing Magazine* 28, 1 (2011).
- [44] B. Xu et al. 2013. High-Accuracy TDOA-Based Localization Without Time Synchronization. *IEEE Transactions of Parallel Distributed Systems* 24, 8 (2013).
- [45] Wang Yan et al. 2015. The Designing of Indoor Local. System Based on Self-Organized WSN Using PulsON UWB Sensors. In *Proc. of the 2nd ICISCE Conf.*
- [46] Yanni Yang et al. 2016. Person Authentication Using Finger Snapping – A New Biometric Trait. *Biometric Recognition* (2016).
- [47] Jinyun Zhang et al. 2009. UWB Systems for Wireless Sensor Networks. *Proc. of the IEEE* 97, 2 (2009).

Bibliography

- [1] 3db Access AG. Technical articles. Impulse radio UWB principles and regulation. <https://3db-access.com/article/17>. Accessed: 2020-01-20.
- [2] Z. Abu-Shaban, H. Wymeersch, T. Abhayapala, and G. Seco-Granados. Single-Anchor Two-Way Localization Bounds for 5G mmWave Systems. *arXiv preprint arXiv:1805.02319*, 2018.
- [3] G. R. Aiello and G. D. Rogerson. Ultra-wideband wireless systems. *IEEE Microwave Magazine*, 4(2):36–47, June 2003.
- [4] A. Alarifi, A. Al-Salman, M. Alsaleh, A. Alnafessah, S. Al-Hadhrami, M. A. Al-Ammar, and H. S. Al-Khalifa. Ultra Wideband Indoor Positioning Technologies: Analysis and Recent Advances. *Sensors*, 2016.
- [5] H. Arslan. Adaptation Techniques and Enabling Parameter Estimation Algorithms for Wireless Communications Systems. *Adaptive Signal Proc. in Wireless Communications*, 2009.
- [6] D. D. Arumugam. Single-Anchor 2-D Magnetoquasistatic Position Sensing for Short to Long Ranges Above Ground. *IEEE Antennas and Wireless Propagation Letters*, 15:1325–1328, 2016.
- [7] L. Atzori, A. Iera, and G. Morabito. The Internet of Things: A survey. *Computer Networks*, 2010.
- [8] N. Baccour, A. Koubâa, L. Mottola, M. A. Zúñiga, H. Youssef, C. A. Boano, and M. Alves. Radio link quality estimation in wireless sensor networks: A survey. *ACM Transactions on Sensor Networks*, 8(4):1–33, September 2012.
- [9] M. S. Bakr, B. Großwindhager, M. Rath, J. Kulmer, I. C. Hunter, R. A. Abd-Alhameed, K. Witrisal, C. A. Boano, K. Römer, and W. Bösch. A Compact Broadband Frequency Selective Microstrip Antenna and Its Application to Indoor Positioning Systems for Wireless Networks. *IET Microwaves, Antennas and Propagation*, March 2019.
- [10] K. Balachandran, S. R. Kadaba, and S. Nanda. Channel quality estimation and rate adaptation for cellular mobile radio. *IEEE Journal on Selected Areas in Communications*, 17(7):1244–1256, 1999.

- [11] M. Bargh and R. de Groot. Indoor Localization based on Response Rate of Bluetooth Inquiries. In *Proceedings of the 1st ACM International Workshop on Mobile Entity Localization and Tracking in GPS-less Environments (MELT)*, pages 49–54, September 2008.
- [12] A. Basiri, E. S. Lohan, T. Moore, A. Winstanley, P. Peltola, C. Hill, P. Amirian, and P. F. e Silva. Indoor location based services challenges, requirements and usability of current solutions. *Computer Science Review*, 24, 2017.
- [13] A. Batra et al. Multi-band OFDM Physical Layer Proposal for IEEE 802.15 Task Group 3a. http://grouper.ieee.org/groups/802/15/pub/2003/Jul03/03268r1P802-15_TG3a-Multi-band-CFP-Document.doc, 2003.
- [14] B. Baumann. NetLoc: A UWB-based Localization System for the Internet of Things. Master’s thesis, TU Graz, 2018.
- [15] BBC News. New cars ‘can be broken into in 10 seconds’. <https://www.bbc.com/news/business-49273028>. Accessed: 2020-01-21.
- [16] A. H. Behzadan, Z. Aziz, C. J. Anumba, and V. R. Kamat. Ubiquitous location tracking for context-specific information delivery on construction sites. *Automation in Construction*, 17(6):737–748, 2008.
- [17] A. Benini, A. Mancini, and S. Longhi. An IMU/UWB/Vision-based Extended Kalman Filter for Mini-UAV Localization in Indoor Environment using 802.15.4a Wireless Sensor Network. *Journal of Intelligent & Robotic Systems*, 70(1):461–476, April 2013.
- [18] A. Bensky. *Short-range wireless communication*. Newnes, 2019.
- [19] P. Bihler, P. Imhoff, and A. B. Cremers. Smartguide—a smartphone museum guide with ultrasound control. *Procedia Computer Science*, 5:586–592, 2011.
- [20] J. Blankenbach and A. Norrdine. Position estimation using artificial generated magnetic fields. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, September 2010.
- [21] M. Blösch, S. Weiss, D. Scaramuzza, and R. Siegwart. Vision based MAV navigation in unknown and unstructured environments. In *2010 IEEE International Conference on Robotics and Automation*, pages 21–28. IEEE, 2010.
- [22] M. Bocca, O. Kaltiokallio, N. Patwari, and S. Venkatasubramanian. Multiple Target Tracking with RF Sensor Networks. *IEEE Transactions on Mobile Computing*, 13(8):1787–1800, 2014.
- [23] F. Boochs, R. Schütze, C. Simon, F. Marzani, H. Wirth, and J. Meier. Increasing the accuracy of untaught robot positions by means of a multi-camera system. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Sept. 2010.

- [24] S. Brasche and W. Bischof. Daily time spent indoors in German homes – Baseline data for the assessment of indoor exposure of German occupants. *International Journal of Hygiene and Environmental Health*, 208(4):247–253, 2005.
- [25] S. Brienza, M. Roveri, D. D. Guglielmo, and G. Anastasi. Just-in-time adaptive algorithm for optimal parameter setting in 802.15.4 WSNs. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 10(4), 2016.
- [26] T. Casey, B. Guimond, and J. Hu. Underwater Vehicle Positioning Based on Time of Arrival Measurements from a Single Beacon. In *OCEANS 2007*, 2007.
- [27] Q. Chang, Q. Li, Z. Shi, W. Chen, and W. Wang. Scalable Indoor Localization via Mobile Crowdsourcing and Gaussian Process. *Sensors*, 16(3), 2016.
- [28] J. Chen, D. Steinmetzer, J. Classen, E. Knightly, and M. Hollick. Pseudo lateration: Millimeter-wave localization using a single RF chain. In *2017 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2017.
- [29] Y. Cheng, K. Chen, H. Sun, Y. Zhang, and F. Tao. Data and knowledge mining with big data towards smart production. *Journal of Industrial Information Integration*, 9, 2018.
- [30] A. Cidronali, S. Maddio, G. Giorgetti, and G. Manes. Analysis and Performance of a Smart Antenna for 2.45-GHz Single-Anchor Indoor Positioning. *IEEE Transactions on Microwave Theory and Techniques*, 58(1):21–31, January 2010.
- [31] G. Conte et al. BlueSentinel: a First Approach using iBeacon for an Energy Efficient Occupancy Detection System. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings (BuildSys)*, pages 11–19, November 2014.
- [32] A. Conti, S. Mazuelas, S. Bartoletti, W. C. Lindsey, and M. Z. Win. Soft Information for Localization-of-Things. *Proceedings of the IEEE*, 2019.
- [33] P. Corbalán and G. P. Picco. Concurrent Ranging in Ultra-wideband Radios: Experimental Evidence, Challenges, and Opportunities. In *Proceedings of the 2018 International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 55–66, 2018.
- [34] P. Corbalán, G. P. Picco, and S. Palipana. Chorus: UWB Concurrent Transmissions for GPS-like Passive Localization of Countless Targets. In *Proceedings of the 18th Intern. Conf. on Information Processing in Sensor Networks (IPSN)*, pages 133–144, 2019.
- [35] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Z. Win. Ranging With Ultrawide Bandwidth Signals in Multipath Environments. *Proc. of the IEEE*, 97(2):404–426, 2009.
- [36] A. P. Das and S. M. Thampi. Single Anchor Node Based Localization in Mobile Underwater Wireless Sensor Networks. In *International Conference on Algorithms and Architectures for Parallel Processing*, pages 757–770, 2015.

- [37] DecaWave Ltd. *APR001 Application Note. UWB Regulations. A Summary of Worldwide Telecomm. Regulations governing the use of Ultra-Wideband radio. Version 1.2*, 2015.
- [38] DecaWave Ltd. *DW1000 Datasheet. Version 2.12*, 2016.
- [39] Decawave Ltd. *DW1000 User Manual. Version 2.13*, 2017.
- [40] Decawave Ltd. *Product Datasheet: DWM1001-DEV. Version 1.3*, 2017.
- [41] M.-G. Di Benedetto and G. Giancola. *Understanding Ultra Wide Band Radio Fundamentals*. Pearson Education, Inc., 2004.
- [42] M. Di Francesco, G. Anastasi, M. Conti, S. K. Das, and V. Neri. Reliability and Energy-Efficiency in IEEE 802.15.4/ZigBee Sensor Networks: An Adaptive and Cross-Layer Approach. *IEEE Journal on Selected Areas in Communications*, 29(8):1508–1524, 2011.
- [43] A. Dutta and S. Sarin. Performance optimization in GSM networks through dynamic power control. In *5th IEEE International Conference on High Speed Networks and Multimedia Communication (Cat. No.02EX612)*, pages 197–201, 2002.
- [44] EETimes. Volkswagen and NXP Show First Car Using UWB To Combat Relay Theft. https://www.eetimes.com/document.asp?doc_id=1335062. Accessed: 2019-09-20.
- [45] EETimes. Xtreme Spectrum demos UWB chip set. <https://www.eetimes.com/xtreme-spectrum-demos-uw-b-chip-set/>. Accessed: 2019-12-27.
- [46] V. Erdélyi, T.-K. Le, B. Bhattacharjee, P. Druschel, and N. Ono. Sonoloc: Scalable Positioning of Commodity Mobile Devices. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys '18*, 2018.
- [47] ETSI. Short Range Devices (SRD) using Ultra Wide Band technology (UWB); Harmonised Standard covering the essential requirements of article 3.2 of the Directive 2014/53/EU; Part 1: Requirements for Generic UWB applications. *ETSI EN 302 065-1*, v2.1.1, November 2016.
- [48] ETSI. Short Range Devices (SRD) using Ultra Wide Band technology (UWB); Harmonised Standard covering the essential requirements of article 3.2 of the Directive 2014/53/EU; Part 2: Requirements for UWB location tracking. *ETSI EN 302 065-2*, v2.1.1, November 2016.
- [49] ETSI. Short Range Devices (SRD) using Ultra Wide Band technology (UWB); Harmonised Standard covering the essential requirements of article 3.2 of the Directive 2014/53/EU; Part 3: Requirements for UWB devices for ground based vehicular applications. *ETSI EN 302 065-3*, v2.1.1, Nov. 2016.
- [50] European Commission. Commission decision on allowing the use of the radio spectrum for equipment using ultra-wideband technology in a harmonised manner in the community. *Official Journal of the European Union*, 2007/131/EC, February 2007.

- [51] C. Falsi, D. Dardari, L. Mucchi, and M. Z. Win. Time of Arrival Estimation for UWB Localizers in Realistic Environments. *EURASIP Journal on Advances in Signal Processing*, 2006(1), July 2006.
- [52] R. Faragher and R. Harle. Location Fingerprinting With Bluetooth Low Energy Beacons. *IEEE Journal on Selected Areas in Communications*, 33(11):2418–2428, Nov. 2015.
- [53] Federal Communications Commission (FCC). Revision of Part 15 of the Commission’s Rule Regarding UWB Transmission Systems. *FCC 10-151*, 2002.
- [54] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele. Low-Power Wireless Bus. In *Proc. of the 10th ACM Conf. on Embedded Network Sensor Systems (SenSys)*, 2012.
- [55] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient network flooding and time synchronization with Glossy. In *Proc. of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 73–84, April 2011.
- [56] A. Freedman, N. Levanon, and S. Gabbay. Perfect periodic correlation sequences. *Signal Processing*, 41(2):165–174, 1995.
- [57] A. Furuskar, S. Mazur, F. Muller, and H. Olofsson. EDGE: Enhanced data rates for GSM and TDMA/136 evolution. *IEEE personal communications*, 6(3):56–66, 1999.
- [58] S. Gansemer, U. Großmann, and S. Hakobyan. RSSI-based Euclidean Distance Algorithm for Indoor Positioning Adapted for the use in Dynamically Changing WLAN Environments and Multi-level Buildings. In *Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, September 2010.
- [59] M. García Fernández. *Contributions to the 3D ionospheric sounding with GPS data*. Universitat Politècnica de Catalunya, 2004.
- [60] C. Gentner, T. Jost, W. Wang, S. Zhang, A. Dammann, and U. Fiebig. Multipath Assisted Positioning with Simultaneous Localization and Mapping. *IEEE Transactions on Wireless Communications*, 15(9):6104–6117, September 2016.
- [61] T. Gigl. *Low-Complexity Localization using Standard-Compliant UWB Signals*. PhD thesis, TU Graz, 2010.
- [62] T. Gigl, F. Troesch, J. Preishuber-Pfluegl, and K. Witrisal. Ranging Performance of the IEEE 802.15.4a UWB Standard Under FCC/CEPT Regulations. *JECE*, 2012, 2012.
- [63] G. Giorgetti, A. Cidronali, S. K. Gupta, and G. Manes. Single-anchor indoor localization using a switched-beam antenna. *IEEE Communications Letters*, 13(1):58–60, 2009.
- [64] B. Großwindhager. The Potential of Low-cost Millimeter-Wave Sensors for Mobile Radar Applications. In *Marshall Plan Scholarship Papers*, 2019.

- [65] B. Großwindhager, M. S. Bakr, M. Rath, F. Gentili, W. Bösch, K. Witrisal, C. A. Boano, and K. Römer. Poster: Switchable Directional Antenna System for UWB-based Internet of Things Applications. In *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 210–211, February 2017.
- [66] B. Großwindhager, C. A. Boano, M. Rath, and K. Römer. Concurrent Ranging with Ultra-Wideband Radios: From Experimental Evidence to a Practical Solution. In *IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pages 1460–1467, July 2018.
- [67] B. Großwindhager, C. A. Boano, M. Rath, and K. Römer. Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications. In *2018 IEEE 19th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, June 2018.
- [68] B. Großwindhager et al. Dependable Internet of Things for Networked Cars. *International Journal of Computing*, 16(4):226–237, 2017.
- [69] B. Großwindhager, M. Rath, M. S. Bakr, P. Greiner, C. A. Boano, K. Witrisal, F. Gentili, J. Grosinger, W. Bösch, and K. Römer. *Dependable Wireless Communication and Localization in the Internet of Things*, pages 209–256. Springer Intern. Publishing, 2019.
- [70] B. Großwindhager, M. Rath, J. Kulmer, M. S. Bakr, C. A. Boano, K. Witrisal, and K. Römer. SALMA: UWB-based Single-Anchor Localization System Using Multipath Assistance. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 132–144, November 2018.
- [71] B. Großwindhager, M. Rath, J. Kulmer, S. Hinteregger, M. Bakr, C. A. Boano, K. Witrisal, and K. Römer. UWB-based Single-anchor Low-cost Indoor Localization System. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 34:1–34:2, November 2017.
- [72] B. Großwindhager, M. Stocker, M. Rath, C. A. Boano, and K. Römer. SnapLoc: An Ultra-fast UWB-based Indoor Localization System for an Unlimited Number of Tags. In *Proceedings of the 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 61–72, April 2019.
- [73] S. Grzonka, G. Grisetti, and W. Burgard. Towards a navigation system for autonomous indoor flying. In *2009 IEEE international conference on Robotics and Automation*, pages 2878–2883. IEEE, 2009.
- [74] Y. Gu, A. Lo, and I. Niemegeers. A survey of indoor positioning systems for wireless personal networks. *IEEE Communications Surveys & Tutorials*, 11(1):13–32, 2009.
- [75] A. Guerra, F. Guidi, A. Clemente, R. D’Errico, L. Dussopt, and D. Dardari. Application of transmitarray antennas for indoor mapping at millimeter-waves. In *2015 European Conference on Networks and Communications (EuCNC)*, pages 77–81, 2015.

- [76] F. Gustafsson and F. Gunnarsson. Positioning using time-difference of arrival measurements. In *Proceedings of the 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'03)*, volume 6, pages VI–553. IEEE, 2003.
- [77] I. Guvenc, C. Chong, and F. Watanabe. NLOS Identification and Mitigation for UWB Localization Systems. In *2007 IEEE Wireless Communications and Networking Conference*, pages 1571–1576, March 2007.
- [78] İ. Güvenç, C.-C. Chong, F. Watanabe, and H. Inamura. NLOS Identification and Weighted Least-Squares Localization for UWB Systems Using Multipath Channel Statistics. *EURASIP Journal on Advances in Signal Processing*, 2008(1), August 2007.
- [79] F. Hartmann et al. Design of an Embedded UWB Hardware Platform for Navigation in GPS Denied Environments. In *Proceedings of the IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*, November 2015.
- [80] J. Haverinen and A. Kemppainen. Global indoor self-localization based on the ambient magnetic field. *Robotics and Autonomous Systems*, 57(10):1028–1035, 2009.
- [81] S. He and S.-H. G. Chan. Wi-Fi Fingerprint-based Indoor Positioning: Recent Advances and Comparisons. *IEEE Comm. Surveys & Tutorials*, 18(1):466–490, August 2016.
- [82] G. Heidari. *WiMedia UWB: technology of choice for wireless USB and Bluetooth*. John Wiley & Sons, 2008.
- [83] M. Heidari, N. A. Alsindi, and K. Pahlavan. UDP identification and error mitigation in toa-based indoor localization systems using neural network architecture. *IEEE Transactions on Wireless Communications*, 8(7):3597–3607, July 2009.
- [84] F. Höflinger et al. Acoustic Self-calibrating System for Indoor Smartphone Tracking (ASSIST). In *2012 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2012.
- [85] B. Hofmann-Wellenhof, K. Legat, and M. Wieser. *Navigation: Principles of Positioning and Guidance*. Springer Science & Business Media, 2003.
- [86] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins. *Global Positioning System: Theory and Practice*. Springer Science & Business Media, 2012.
- [87] B. Hofmann-Wellenhof, H. Lichtenegger, and E. Wasle. *GNSS—global navigation satellite systems: GPS, GLONASS, Galileo, and more*. Springer Science & Business Media, 2007.
- [88] B. Hu, Z. Shi, and Y. Wang. Single-Sensor Based Indoor Localisation by Exploiting Multipath Propagation. *Electronics Letters*, 54(3):179–181, February 2018.
- [89] C. Huang, L. Lee, C. C. Ho, L. Wu, and Z. Lai. Real-Time RFID Indoor Positioning System Based on Kalman-Filter Drift Removal and Heron-Bilateration Location Estimation. *IEEE Transactions on Instrumentation and Measurement*, 64(3):728–739, March 2015.

- [90] M. Ibnkahla. *Adaptive Signal Processing in Wireless Communications*. CRC Press, 2008.
- [91] IEEE. P802.15.4z - IEEE Draft Standard for Low-Rate Wireless Networks Amendment: Enhanced High Rate Pulse (HRP) and Low Rate Pulse (LRP) Ultra Wide-Band (UWB) Physical Layers (PHYs) and Associated Ranging Techniques. https://standards.ieee.org/project/802_15_4z.html. Accessed: 2020-01-28.
- [92] IEEE Computer Society. IEEE Standard 802.15.4a-2007. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs). Amendment 1: Add Alternate PHYs. *IEEE Standard for Local and metropolitan area networks*, 2007.
- [93] IEEE Computer Society. IEEE Standard 802.15.4-2011. Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs). *IEEE Standard for Local and metropolitan area networks*, 2011.
- [94] IEEE Computer Society. IEEE Standard 802.15.4f-2012. Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs). Amendment 2: Active Radio Frequency Identification (RFID) System Physical Layer (PHY). *IEEE Standard for Local and metropolitan area networks*, 2012.
- [95] IEEE Computer Society. IEEE Standard 802.15.4-2015 for Low-Rate Wireless Networks. *IEEE Standard for Local and metropolitan area networks*, 2015.
- [96] IEEE Computer Society. IEEE Draft Standard 802.15.4z/D03. Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs). Amendment: Enhanced Ultra Wide-Band (UWB) Physical Layers (PHYs) and Associated Ranging Techniques. *IEEE Standard for Local and metropolitan area networks*, 2019.
- [97] infsoft GmbH. Indoor Positioning and Services. *White Paper*, 2019.
- [98] International Telecommunication Union Radiocommunication Sector (ITU-R). Recommendation ITU-R SM. 1755. *Characteristics of UWB technology*, 2006.
- [99] Y. Jiang and V. C. M. Leung. An Asymmetric Double Sided Two-Way Ranging for Crystal Offset. In *2007 International Symposium on Signals, Systems and Electronics*, pages 525–528, July 2007.
- [100] D. B. Jourdan, D. Dardari, and M. Z. Win. Position error bound for UWB localization in dense cluttered environments. *IEEE Transactions on Aerospace and Electronic Systems*, 44(2):613–628, April 2008.
- [101] S.-Y. Jung, S. Hann, and C.-S. Park. TDOA-based optical wireless indoor localization using LED ceiling lamps. *IEEE Transactions on Consumer Electronics*, 57(4), 2011.
- [102] A. Kakkavas, M. H. C. Garcia, R. A. Strirling-Gallacher, and J. A. Nossek. Performance Limits of Single-Anchor Millimeter-Wave Positioning. *arXiv preprint arXiv:1808.08116*, 2018.

- [103] E. Karapistoli, F. Pavlidou, I. Gragopoulos, and I. Tsetsinas. An overview of the IEEE 802.15.4a Standard. *IEEE Communications Magazine*, 48(1):47–53, January 2010.
- [104] S. Karpischek, F. Michahelles, F. Resatsch, and E. Fleisch. Mobile Sales Assistant - An NFC-Based Product Information System for Retailers. In *2009 First International Workshop on Near Field Communication*, pages 20–23. IEEE, 2009.
- [105] S. M. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Prentice Hall, 1993.
- [106] B. Kempke, P. Pannuto, B. Campbell, and P. Dutta. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proceedings of the 14th International Conference on Embedded Network Sensor Systems (SenSys)*, pages 137–149, November 2016.
- [107] B. Kempke, P. Pannuto, and P. Dutta. PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization. In *Proceedings of the 2nd International Workshop on Hot Topics in Wireless (HotWireless)*, pages 16–20, September 2015.
- [108] H. Kim. Performance Comparison of Asynchronous Ranging Algorithms. In *GLOBE-COM 2009 - 2009 IEEE Global Telecommunications Conference*, November 2009.
- [109] N. E. Klepeis et al. The National Human Activity Pattern Survey (NHAPS): a resource for assessing exposure to environmental pollutants. *Journal of Exposure Science and Environmental Epidemiology*, 11(3), 2001.
- [110] M. Kok, J. D. Hol, and T. B. Schön. Indoor Positioning Using Ultrawideband and Inertial Measurements. *IEEE Transactions on Vehicular Technology*, 64(4):1293–1303, 2015.
- [111] G. Ku and J. M. Walsh. Resource Allocation and Link Adaptation in LTE and LTE Advanced: A Tutorial. *IEEE Communications Surveys Tutorials*, 17(3):1605–1633, 2015.
- [112] J. Kulmer, S. Grebien, M. Rath, and K. Witrisal. On the Unimportance of Phase-Coherent Measurements for Beampattern-Assisted Positioning. In *Proceedings of the International Conference on Wireless Communications and Networking (WCNC)*, April 2018.
- [113] J. Kulmer, S. Hinteregger, B. Großwindhager, M. Rath, M. S. Bakr, E. Leitinger, and K. Witrisal. Using DecaWave UWB transceivers for high-accuracy multipath-assisted indoor positioning. In *2017 IEEE International Conference on Communications Workshops (ICC Workshops)*, pages 1239–1245, May 2017.
- [114] J. Kulmer, E. Leitinger, S. Grebien, and K. Witrisal. Anchorless Cooperative Tracking using Multipath Channel Information. *IEEE Transactions on Wireless Communications*, PP(99), January 2018.
- [115] O. Landsiedel, F. Ferrari, and M. Zimmerling. Chaos: Versatile and Efficient All-to-All Data Sharing and in-Network Processing at Scale. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2013.

- [116] J. Larranaga, L. Muguira, J. M. Lopez-Garde, and J. I. Vazquez. An Environment Adaptive ZigBee-based Indoor Positioning Algorithm. In *Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, September 2010.
- [117] A. Ledergerber, M. Hamer, and R. D’Andrea. A Robot Self-Localization System using One-Way Ultra-Wideband Communication. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, pages 3131–3137, September 2015.
- [118] B. M. Ledvina, R. W. Brumley, and S. Hariharan. Mobile device for communicating and ranging with access control system for automatic functionality, May 2019. US Patent App. 15/983388.
- [119] J. S. Lee, Y. W. Su, and C. C. Shen. A Comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee, and Wi-Fi. In *IECON 2007-33rd Annual Conference of the IEEE Industrial Electronics Society*, pages 46–51, November 2007.
- [120] E. Leitinger, M. Fröhle, P. Meissner, and K. Witrisal. Multipath-Assisted Maximum-Likelihood Indoor Positioning using UWB Signals. In *Proceedings of the International Conference on Communications Workshops (ICC)*, pages 170–175, June 2014.
- [121] E. Leitinger, F. Meyer, F. Hlawatsch, K. Witrisal, F. Tufvesson, and M. Z. Win. A belief propagation algorithm for multipath-based SLAM. *IEEE Transactions on Wireless Communications*, 18(12):5613–5629, 2019.
- [122] P. Leu, M. Singh, M. Roeschlin, K. G. Paterson, and S. Capkun. Message time of arrival codes: A fundamental primitive for secure distance measurement. In *IEEE Symposium on Security and Privacy*, 2019.
- [123] L. Li, Y. Wang, X. Ma, C. Chen, and X. Guan. Dual-tone radio interferometric positioning systems for multi-target localization using a single mobile anchor. *China Communications*, 12(1):25–35, 2015.
- [124] X. Li, L. Wang, A. Timm-Giel, C. Gorg, R. Schelb, and T. Winter. Optimization of Bit Rate Adaptation in UMTS Radio Access Network. In *2007 IEEE 65th Vehicular Technology Conference - VTC2007-Spring*, pages 1051–1055, 2007.
- [125] C. Liu, D. Fang, Z. Yang, H. Jiang, X. Chen, W. Wang, T. Xing, and L. Cai. RSS Distribution-Based Passive Localization and Its Application in Sensor Networks. *IEEE Transactions on Wireless Communications*, 15(4):2883–2895, 2016.
- [126] H. Liu, H. Darabi, P. Banerjee, and J. Liu. Survey of wireless indoor positioning techniques and systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37(6):1067–1080, 2007.
- [127] K. Liu, X. Liu, and X. Li. Guoguo: Enabling fine-grained indoor localization via smartphone. In *Proceedings of the 11th annual International Conference on Mobile systems, applications, and services*, pages 235–248. ACM, 2013.

- [128] S. Liu and T. He. SmartLight: Light-weight 3D Indoor Localization Using a Single LED Lamp. In *Proceedings of the 15th International Conference on Embedded Network Sensor Systems (SenSys)*, November 2017.
- [129] D. Lymberopoulos and J. Liu. The Microsoft Indoor Localization Competition: Experiences and Lessons Learned. *IEEE Signal Proc. Magazine*, 34(5):125–140, Sept. 2017.
- [130] S. Ma and Y. Shi. A Scalable Passive RFID-Based Multi-User Indoor Location System. In *7th Intern. Conf. on Wireless Comm., Networking and Mobile Computing*, 2011.
- [131] S. Maranò, W. M. Gifford, H. Wymeersch, and M. Z. Win. NLOS Identification and Mitigation for Localization Based on UWB Experimental Data. *IEEE Journal on Selected Areas in Communications*, 28(7):1026–1035, September 2010.
- [132] R. Mautz. *Indoor positioning technologies*. Habilitation, ETH Zürich, 2012.
- [133] R. McAulay and E. Hofstetter. Barankin Bounds on Parameter Estimation. *IEEE Transactions on Information Theory*, 17(6):669–676, November 1971.
- [134] C. Medina, J. Segura, and A. De la Torre. Ultrasound Indoor Positioning System Based on a Low-Power Wireless Sensor Network Providing Sub-Centimeter Accuracy. *Sensors*, 13(3):3501–3526, 2013.
- [135] P. Meissner. *Multipath-Assisted Indoor Positioning*. PhD thesis, TU Graz, 2014.
- [136] P. Meissner, C. Steiner, and K. Witrisal. UWB Positioning with Virtual Anchors and Floor Plan Information. In *Proceedings of the 7th International Workshop on Positioning, Navigation and Communication (WPNC)*, pages 150–156, March 2010.
- [137] K. Mikhaylov, J. Petäjajarvi, M. Hämäläinen, A. Tikanmäki, and R. Kohno. Impact of IEEE 802.15.4 Communication Settings on Performance in Asynchronous Two Way UWB Ranging. *Intern. Journal of Wireless Information Networks*, pages 1–16, 2017.
- [138] A. F. Molisch. Ultra-Wide-Band Propagation Channels. *Proceedings of the IEEE*, 97(2):353–371, 2009.
- [139] A. F. Molisch. *Wireless Communications*, volume 34. John Wiley & Sons, 2012.
- [140] A. F. Molisch, K. Balakrishnan, C. chin Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak. IEEE 802.15.4a Channel Model. In *Converging: Technology, Work and Learning*, 2004.
- [141] P. Mowlae, J. Kulmer, J. Stahl, and F. Mayer. *Single Channel Phase-Aware Signal Processing in Speech Communication: Theory and Practice*. John Wiley & Sons, 2016.
- [142] R. Murai, T. Sakai, H. Kawano, Y. Matsukawa, Y. Kitano, Y. Honda, and K. C. Campbell. A novel visible light communication system for enhanced control of autonomous delivery robots in a hospital. In *IEEE/SICE Intern. Symposium on System Integration (SII)*, 2012.

- [143] S. Nanda, K. Balachandran, and S. Kumar. Adaptation techniques in wireless packet data services. *IEEE Communications Magazine*, 38(1):54–64, 2000.
- [144] D. Neiryneck, E. Luk, and M. McLaughlin. An alternative double-sided two-way ranging method. In *2016 13th Workshop on Positioning, Navigation and Communications (WPNC)*, October 2016.
- [145] G. Oberholzer, P. Sommer, and R. Wattenhofer. Spiderbat: Augmenting wireless sensor networks with distance and angle information. In *Proceedings of the 10th ACM/IEEE Intern. Conf. on Information Processing in Sensor Networks (IPSN)*, pages 211–222, 2011.
- [146] I. Oppermann, M. Hämmäläinen, and J. Iinatti. *UWB: theory and applications*. John Wiley & Sons, 2005.
- [147] C. W. Ou et al. A ZigBee Position Technique for Indoor Localization Based on Proximity Learning. In *Proc. of the International Conference on Mechatronics and Automation (ICMA)*, pages 875–880, August 2017.
- [148] M. Passafiume, S. Maddio, G. Collodi, and A. Cidronali. An enhanced algorithm for 2D indoor localization on single anchor RSSI-based positioning systems. In *2017 European Radar Conference (EURAD)*, pages 287–290. IEEE, 2017.
- [149] R. Patachaianand and K. Sandrasegaran. Performance Comparison of Adaptive Power Control in UMTS. In *The 2nd International Conference on Wireless Broadband and Ultra Wideband Communications (AusWireless 2007)*, 2007.
- [150] R. Paucher and M. Turk. Location-based augmented reality on mobile phones. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, pages 9–16, June 2010.
- [151] C. Perera et al. Context aware computing for the internet of things: A survey. *IEEE Communications Surveys & Tutorials*, 16(1):414–454, 2013.
- [152] A. Purohit, Z. Sun, Shijia Pan, and P. Zhang. SugarTrail: Indoor navigation in retail environments without surveys and maps. In *2013 IEEE International Conference on Sensing, Communications and Networking (SECON)*, pages 300–308, June 2013.
- [153] Y. Qi, H. Li, S. Hara, and R. Kohno. Clear Channel Assessment (CCA) with multiplexed preamble symbols for impulse Ultra-wideband (UWB) communications. In *2006 IEEE International Conference on Ultra-Wideband*, pages 675–680, 2006.
- [154] X. Qing, Z. N. Chen, and T. S. P. See. Sectorized Antenna Array for Indoor Mono-Station UWB Positioning Applications. In *Proceedings of the 3rd European Conference on Antennas and Propagation*, pages 822–825, March 2009.
- [155] T. Qu, S. Lei, Z. Wang, D. Nie, X. Chen, and G. Q. Huang. IoT-based real-time production logistics synchronization system under smart cloud manufacturing. *The International Journal of Advanced Manufacturing Technology*, 84(1-4):147–164, 2016.

- [156] D. R. Quiñones, G. Lopes, D. Kim, C. Honnet, D. Moratal, and A. Kampff. HIVE Tracker: a tiny, low-cost, and scalable device for sub-millimetric 3D positioning. In *Proceedings of the 9th Augmented Human International Conference*, pages 1–8, 2018.
- [157] P. Rashidi and A. Mihailidis. A Survey on Ambient-Assisted Living Tools for Older Adults. *IEEE journal of biomedical and health informatics*, 17(3):579–590, 2012.
- [158] M. Ridolfi, S. Van de Velde, H. Steendam, and E. De Poorter. Analysis of the Scalability of UWB Indoor Localization Solutions for High User Densities. *Sensors*, 18(6), 2018.
- [159] R. Roberts. XtremeSpectrum CFP Document. Doc. IEEE P802.15-03/154r3, 2003.
- [160] K. Römer and F. Mattern. The design space of wireless sensor networks. *IEEE Wireless Communications*, 11(6):54–61, December 2004.
- [161] N. C. Rowe, A. E. Fathy, M. J. Kuhn, and M. R. Mahfouz. A UWB transmit-only based scheme for multi-tag support in a millimeter accuracy localization system. In *2013 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*, pages 7–9, 2013.
- [162] J. L. Rullan-Lara, G. Sanahuja, R. Lozano, S. Salazar, R. Garcia-Hernandez, and J. A. Ruz-Hernandez. Indoor Localization of a Quadrotor Based on WSN: A Real-Time Application. *International Journal of Advanced Robotic Systems*, 10(1), 2013.
- [163] M. Rzymowski, P. Woznica, and L. Kulas. Single-Anchor Indoor Localization Using ESPAR Antenna. *IEEE Antennas and Wireless Prop. Letters*, 15:1183–1186, Nov. 2016.
- [164] S. S. Saab and Z. S. Nakad. A standalone RFID indoor positioning system using passive tags. *IEEE Transactions on Industrial Electronics*, 58(5):1961–1970, 2010.
- [165] A. Saeed, A. E. Kosba, and M. Youssef. Ichnaea: A Low-Overhead Robust WLAN Device-Free Passive Localization System. *IEEE Journal of Selected Topics in Signal Processing*, 8(1):5–15, 2014.
- [166] Z. Sahinoglu, S. Gezici, and I. Guvenc. *Ultra-wideband positioning systems: Theoretical limits, ranging algorithms, and protocols*. Cambridge University Press, 2008.
- [167] C. L. Sang, M. Adams, T. Hörmann, M. Hesse, M. Porrman, and U. Rückert. An Analytical Study of Time of Flight Error Estimation in Two-Way Ranging Methods. In *2018 Intern. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*, Sept. 2018.
- [168] R. Scholtz. Multiple access with time-hopping impulse modulation. In *Proceedings of MILCOM'93-IEEE Military Comm. Conference*, volume 2, pages 447–450. IEEE, 1993.
- [169] H. Schweinzer and M. Syafrudin. LOSNUS: An ultrasonic system enabling high accuracy and secure TDoA locating of numerous devices. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2010.

- [170] F. Seco, C. Plagemann, A. R. Jiménez, and W. Burgard. Improving RFID-based indoor positioning accuracy using Gaussian processes. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, September 2010.
- [171] P. Sedlacek, M. Slanina, and P. Masek. An Overview of the IEEE 802.15.4z Standard its Comparison and to the Existing UWB Standards. In *2019 29th International Conference Radioelektronika (RADIOELEKTRONIKA)*, April 2019.
- [172] M. Seifeldin, A. Saeed, A. E. Kosba, A. El-Keyi, and M. Youssef. Nuzzer: A Large-Scale Device-Free Passive Localization System for Wireless Environments. *IEEE Transactions on Mobile Computing*, 12(7):1321–1334, 2013.
- [173] M. Sha, G. Hackmann, and C. Lu. Energy-efficient Low Power Listening for wireless sensor networks in noisy environments. In *2013 ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 277–288, April 2013.
- [174] M. Shahzad and M. P. Singh. Continuous Authentication and Authorization for the Internet of Things. *IEEE Internet Computing*, 21(2):86–90, March 2017.
- [175] D. F. Shanno. Conditioning of Quasi-Newton Methods for Function Minimization. *Mathematics of Computation*, 24(111):647–656, 1970.
- [176] R. C. Shit, S. Sharma, D. Puthal, and A. Y. Zomaya. Location of Things (LoT): A review and taxonomy of sensors localization in IoT infrastructure. *IEEE Communications Surveys & Tutorials*, 20(3):2028–2061, 2018.
- [177] Y. Shu, C. Bo, G. Shen, C. Zhao, L. Li, and F. Zhao. Magicol: Indoor localization using pervasive magnetic field and opportunistic WiFi sensing. *IEEE Journal on Selected Areas in Communications*, 33(7):1443–1457, 2015.
- [178] B. Silva, Z. Pang, J. Åkerberg, J. Neander, and G. Hancke. Experimental Study of UWB-based High Precision Localization for Industrial Applications. In *Proceedings of the International Conference on Ultra-Wideband (ICUWB)*, pages 280–285, September 2014.
- [179] A. Simonsson and A. Furuskar. Uplink power control in LTE-overview and performance, subtitle: principles and benefits of utilizing rather than compensating for SINR variations. In *2008 IEEE 68th Vehicular Technology Conference*. IEEE, 2008.
- [180] M. Slabicki, G. Premsankar, and M. Di Francesco. Adaptive Configuration of LoRa Networks for Dense IoT Deployments. In *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*, 2018.
- [181] P. Smyth. *Mobile and wireless communications: key technologies and future applications*, volume 9. IET, 2004.
- [182] H. C. So and E. M. K. Shiu. Performance of TOA-AOA hybrid mobile location. *IE-ICE transactions on fundamentals of electronics, communications and computer sciences*, 86(8):2136–2138, 2003.

- [183] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. An empirical study of low-power wireless. *ACM Transactions on Sensor Networks (TOSN)*, 6(2):1–49, 2010.
- [184] C. Steiner. *Location Fingerprinting for Ultra-Wideband Systems - The Key to Efficient and Robust Localization*. PhD thesis, ETH Zurich, 2010.
- [185] M. Stocker, B. Großwindhager, C. A. Boano, and K. Römer. SnapLoc: An Ultra-fast UWB-based Indoor Localization System for an Unlimited Number of Tags: Demo Abstract. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 348–349, April 2019.
- [186] R. Stoleru, P. Vicaire, T. He, and J. A. Stankovic. StarDust: A Flexible Architecture for Passive Localization in Wireless Sensor Networks. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 57—70, 2006.
- [187] Strategy Analytics. Report: Global Connected and IoT Device Forecast Update, 2019.
- [188] X. Sun, Y. Ma, J. Xu, J. Zhang, and J. Wang. A high Accuracy Mono-Station UWB Positioning System. In *Proceedings of the International Conference on Ultra-Wideband (ICUWB)*, pages 201–204, September 2008.
- [189] J. Swann, E. Chatre, D. Ludwig, et al. Galileo: Benefits for location based services. *Positioning*, 1(04), 2009.
- [190] Y. Takahashi, N. Honma, J. Sato, T. Murakami, and K. Murata. Accuracy Comparison of Wireless Indoor Positioning Using Single Anchor: TOF only Versus TOF-DOA Hybrid Method. In *2019 IEEE Asia-Pacific Microwave Conf. (APMC)*, pages 1679–1681, 2019.
- [191] H. Tang, S. S. Liao, and S. X. Sun. A prediction framework based on contextual data to support Mobile Personalized Marketing. *Decision Support Systems*, 56:234–246, 2013.
- [192] Thatcham Research. The Relay Attack: Guidance for concerned car owners. <https://www.thatcham.org/relay-attack/>. Accessed: 2020-01-21.
- [193] J. Tiemann, F. Eckermann, and C. Wietfeld. ATLAS - An Open-Source TDOA-based Ultra-Wideband Localization System. In *Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, October 2016.
- [194] J. Tiemann, F. Schweikowski, and C. Wietfeld. Design of an UWB indoor-positioning system for UAV navigation in GNSS-denied environments. In *2015 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2015.
- [195] N. O. Tippenhauer, H. Luecken, M. Kuhn, and S. Capkun. UWB Rapid-Bit-Exchange System for Distance Bounding. In *Proceedings of the 8th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '15*, 2015.
- [196] R. Toledo-Moreo, D. Betaille, F. Peyret, and J. Laneurit. Fusing GNSS, Dead-Reckoning, and Enhanced Maps for Road Vehicle Lane-Level Navigation. *IEEE Journal of Selected Topics in Signal Processing*, 3(5):798–809, October 2009.

- [197] N. Uchitomi, A. Inada, M. Fujimoto, T. Wada, K. Mutsuura, and H. Okada. Accurate indoor position estimation by Swift-Communication Range Recognition (S-CRR) method in passive RFID systems. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2010.
- [198] United Nations. Satellite navigation and location systems. Background paper 4. *Third United Nations Conference on the Exploration and Peaceful Uses of Outer Space, A/CONF.184/BP/4*, 1998.
- [199] USC. How to improve communication between people and smart buildings, 2019.
- [200] H. L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory*. John Wiley & Sons, 2004.
- [201] D. Vasisht, S. Kumar, and D. Katabi. Decimeter-level Localization with a Single WiFi Access Point. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation (NSDI)*, pages 165–178, March 2016.
- [202] S. Venkatesh and R. M. Buehrer. Non-line-of-sight identification in ultra-wideband systems based on received signal statistics. *IET Microwaves, Antennas Propagation*, 1(6):1120–1130, December 2007.
- [203] S. Venkatraman and J. Caffery. Hybrid TOA/AOA techniques for mobile location in non-line-of-sight environments. In *2004 IEEE Wireless Communications and Networking Conference (IEEE Cat. No.04TH8733)*, volume 1, pages 274–278 Vol.1, 2004.
- [204] A. E. Waadt, C. Kocks, S. Wang, G. H. Bruck, and P. Jung. Maximum likelihood localization estimation based on received signal strength. In *2010 3rd Intern. Symposium on Applied Sciences in Biomedical and Communication Techn. (ISABEL 2010)*, Nov. 2010.
- [205] J. Wagner, C. Isert, A. Purschwitz, and A. Kistner. Improved vehicle positioning for indoor navigation in parking garages through commercially available maps. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Sept. 2010.
- [206] T. Wang, H. Zhao, and Y. Shen. An Efficient Single-Anchor Localization Method Using Ultra-Wide Bandwidth Systems. *Applied Sciences*, 10(1), 2020.
- [207] Y. Wang, Xu Yang, Yutian Zhao, Yue Liu, and L. Cuthbert. Bluetooth positioning using RSSI and triangulation methods. In *2013 IEEE 10th Consumer Communications and Networking Conference (CCNC)*, pages 837–842, January 2013.
- [208] R. Wen, C.-K. Chui, S.-H. Ong, K.-B. Lim, and S. K.-Y. Chang. Projection-based visual guidance for robot-aided RF needle insertion. *International Journal of Computer Assisted Radiology and Surgery*, 8(6):1015–1025, November 2013.
- [209] J. Wendel, O. Meister, C. Schlaile, and G. F. Trommer. An integrated GPS/MEMS-IMU navigation system for an autonomous helicopter. *Aerospace Science and Technology*, 10(6):527–533, 2006.

- [210] R. Wilfinger, T. Moder, M. Wieser, and B. Großwindhager. Indoor position determination using location fingerprinting and vehicle sensor data. In *2016 European Navigation Conference (ENC)*, May 2016.
- [211] M. Win, D. Dardari, A. Molisch, W. Wiesbeck, and Jinyun Zhang. History and Applications of UWB. *Proceedings of the IEEE*, 97(2):198–204, Febr. 2009.
- [212] M. Z. Win and R. A. Scholtz. Impulse radio: how it works. *IEEE Communications Letters*, 2(2):36–38, February 1998.
- [213] M. Z. Win and R. A. Scholtz. On the robustness of ultra-wide bandwidth signals in dense multipath environments. *IEEE Communications Letters*, 2(2):51–53, February 1998.
- [214] M. Z. Win and R. A. Scholtz. Ultra-Wide Bandwidth Time-Hopping Spread-Spectrum Impulse Radio for Wireless Multiple-Access Communications. *IEEE Transactions on Communications*, 48(4), 2000.
- [215] Wired. The Biggest iPhone News Is a Tiny New Chip Inside It. <https://www.wired.com/story/apple-u1-chip/>. Accessed: 2019-09-20.
- [216] Wired. The technology behind Amazon’s Go store. <https://www.wired.co.uk/article/amazon-go-seattle-uk-store-how-does-work>. Accessed: 2019-10-28.
- [217] V. Wirz, J. Beutel, S. Gruber, S. Gubler, and R. S. Purves. Estimating velocity from noisy GPS data for investigating the temporal variability of slope movements. *Natural Hazards and Earth System Sciences*, 14(9):2503–2520, 2014.
- [218] K. Witrisal et al. High-Accuracy Localization for Assisted Living: 5G systems will turn multipath channels from foe to friend. *IEEE Signal Processing Magazine*, 33(2):59–70, March 2016.
- [219] K. Witrisal, E. Leitinger, S. Hinteregger, and P. Meissner. Bandwidth Scaling and Diversity Gain for Ranging and Positioning in Dense Multipath Channels. *IEEE Wireless Communications Letters*, 5(4):396–399, August 2016.
- [220] K. Witrisal, G. Leus, G. J. M. Janssen, M. Pausini, F. Troesch, T. Zasowski, and J. Romme. Noncoherent ultra-wideband systems. *IEEE Signal Processing Magazine*, 26(4):48–66, July 2009.
- [221] K. Witrisal and P. Meissner. Performance Bounds for Multipath-Assisted Indoor Navigation and Tracking (MINT). In *Proceedings of the International Conference on Communications (ICC)*, pages 4321–4325, June 2012.
- [222] R. W. Wolcott and R. M. Eustice. Fast LIDAR localization using multiresolution Gaussian mixture maps. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2814–2821, May 2015.

- [223] Y.-C. Wu, Q. Chaudhari, and E. Serpedin. Clock Synchronization of Wireless Sensor Networks. *IEEE Signal Processing Magazine*, 28(1):124–138, 2011.
- [224] M. P. Wylie and J. Holtzman. The non-line of sight problem in mobile location estimation. In *Proceedings of ICUPC-5th International Conference on Universal Personal Communications*, volume 2, pages 827–831. IEEE, 1996.
- [225] H. Wymeersch, J. Lien, and M. Z. Win. Cooperative Localization in Wireless Networks. *Proceedings of the IEEE*, 97(2):427–450, 2009.
- [226] H. Wymeersch, S. Marano, W. M. Gifford, and M. Z. Win. A Machine Learning Approach to Ranging Error Mitigation for UWB Localization. *IEEE Transactions on Communications*, 60(6):1719–1728, June 2012.
- [227] B. Xu, G. Sun, R. Yu, and Z. Yang. High-Accuracy TDOA-Based Localization without Time Synchronization. *IEEE Transactions on Parallel and Distributed Systems*, 24(8):1567–1576, August 2013.
- [228] C. Yang and H. Shao. WiFi-based indoor positioning. *IEEE Communications Magazine*, 53(3):150–157, March 2015.
- [229] A. Yassin, Y. Nasser, M. Awad, A. Al-Dubai, R. Liu, C. Yuen, R. Raulefs, and E. Aboutanios. Recent Advances in Indoor Localization: A Survey on Theoretical Approaches and Applications. *IEEE Comm. Surveys Tutorials*, 19(2):1327–1346, 2017.
- [230] F. Ye, R. Chen, G. Guo, X. Peng, Z. Liu, and L. Huang. A Low-Cost Single-Anchor Solution for Indoor Pos. Using BLE and Inertial Sensor Data. *IEEE Access*, 7, 2019.
- [231] T. Ye, M. Walsh, P. Haigh, J. Barton, A. Mathewson, and B. O’Flynn. An experimental evaluation of IEEE 802.15.4a Ultra Wide Band technology for precision indoor ranging. *International Journal of Ambient Computing and Intelligence (IJACI)*, 4(2):48–63, 2012.
- [232] J. Yick, B. Mukherjee, and D. Ghosal. Wireless sensor network survey. *Computer Networks*, 2008.
- [233] Y. Yoshimura, S. Sobolevsky, C. Ratti, F. Girardin, J. P. Carrascal, J. Blat, and R. Sinatra. An Analysis of Visitors’ Behavior in the Louvre Museum: A Study Using Bluetooth Data. *Environment and Planning B: Planning and Design*, 41(6):1113–1131, 2014.
- [234] K. Yu, I. Sharp, and Y. J. Guo. *Ground-based wireless positioning*, volume 5. John Wiley & Sons, 2009.
- [235] R. Zekavat and R. M. Buehrer. *Handbook of position location: Theory, practice and advances*, volume 27. John Wiley & Sons, 2011.
- [236] H. Zhang, X. Cui, B. An, and T. A. Gulliver. A Distance and Angle Estimated Method based on Single UWB Station. In *Proceedings of the International Conference on Signal Processing, Communication and Computing (ICSPCC)*, pages 1–6, August 2013.

- [237] J. Zhang, R. A. Kennedy, and T. D. Abhayapala. Cramer-Rao lower bounds for the time delay estimation of UWB signals. In *2004 IEEE International Conference on Communications (IEEE Cat. No.04CH37577)*, volume 6, pages 3424–3428 Vol.6, June 2004.
- [238] J. Zhang, P. V. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney. UWB Systems for Wireless Sensor Networks. *Proceedings of the IEEE*, 97(2):313–331, February 2009.
- [239] Y. Zhang, A. K. Brown, W. Q. Malik, and D. J. Edwards. High Resolution 3-D Angle of Arrival Determination for Indoor UWB Multipath Propagation. *IEEE Transactions on Wireless Communications*, 7(8):3047–3055, August 2008.
- [240] Y. Zhuang, J. Yang, Y. Li, L. Qi, and N. El-Sheimy. Smartphone-based indoor localization with bluetooth low energy beacons. *Sensors (Switzerland)*, 16(5), April 2016.
- [241] M. Zimmerling, F. Ferrari, L. Mottola, T. Voigt, and L. Thiele. pTUNES: Runtime parameter adaptation for low-power MAC protocols. In *2012 ACM/IEEE 11th Intern. Conference on Information Processing in Sensor Networks (IPSN)*, pages 173–184, April 2012.
- [242] M. Zimmerling, L. Mottola, and S. Santini. Synchronous Transmissions in Low-Power Wireless: A Survey of Communication Protocols and Network Services. *CoRR*, abs/2001.08557, 2020.
- [243] Zippin. Website. <https://www.getzippin.com/>. Accessed: 2019-10-28.
- [244] J. Ziv and M. Zakai. Some lower bounds on signal parameter estimation. *IEEE Transactions on Information Theory*, 15(3):386–391, 1969.
- [245] T. Zwick, W. Wiesbeck, J. Timmermann, and G. Adamiuk. *Ultra-wideband RF system engineering*. Cambridge University Press, 2013.