



University Politehnica
of Bucharest



Tampere University

LAURA-ȘTEFANIA FLUERĂTORU

**Robust, Energy-Efficient, and
Scalable Indoor Localization with
Ultra-Wideband Technology**

LAURA-ȘTEFANIA FLUERĂTORU

Robust, Energy-Efficient, and Scalable
Indoor Localization with
Ultra-Wideband Technology



ACADEMIC DISSERTATION

This thesis has been done in a double PhD degree agreement between:
University Politehnica of Bucharest, Romania
Tampere University, Finland

The work described in this thesis has been conducted within the project A-WEAR, <http://www.a-wear.eu/>. This project has received funding from the European Union's Horizon 2020 (H2020) Marie Skłodowska-Curie Innovative Training Networks H2020-MSCA-ITN-2018 call, under the Grant Agreement no 813278.

<i>Responsible supervisor and Custos</i>	Prof. Dragoş Niculescu University Politehnica of Bucharest Romania	
<i>Supervisors</i>	Prof. Dragoş Niculescu University Politehnica of Bucharest Romania	Prof. Elena Simona Lohan Tampere University Finland
	Prof. Jari Nurmi Tampere University Finland	
<i>Pre-examiners</i>	Dr. Christian Gentner German Aerospace Center Germany	Dr. Marc Kuhn ZHAW School of Engineering Switzerland
<i>Opponents</i>	Dr. Christian Gentner German Aerospace Center Germany	Prof. Radu Tudor Ionescu University of Bucharest Romania

The originality of this thesis has been checked using the Turnitin Originality Check service.
Copyright © 2023 Author

ISBN 978-952-03-2937-2 (print)

ISBN 978-952-03-2938-9 (online)

<http://urn.fi/URN:ISBN:978-952-03-2938-9>



“Would you tell me, please, which way I ought to go from here?”

“That depends a good deal on where you want to get to,” said the Cat.

“I don’t much care where—” said Alice.

“Then it doesn’t matter which way you go,” said the Cat.

“—so long as I get somewhere,” Alice added as an explanation.

“Oh, you’re sure to do that,” said the Cat, “if you only walk long enough.”

Alice in Wonderland, by LEWIS CARROLL

ACKNOWLEDGMENTS

The thesis is based on research carried out during three years, between November 2019 and November 2022, as part of a double-degree PhD program at University Politehnica of Bucharest (UPB), Romania, and Tampere University (TAU), Finland. I gratefully acknowledge the financial support from European Union’s Horizon 2020 Research and Innovation programme under the Marie Skłodowska-Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints).

I would like to thank my main supervisors, Prof. Dragoș Niculescu and Prof. Simona Lohan, who have offered me invaluable advice and support to grow as a researcher and as a person. The long technical discussions I had with Dragoș helped me uncover new perspectives on my area of research and stay on track with my work even during the challenging Covid-19 pandemic. His inquisitive questions, although difficult to deal with, were always spot-on and helped me identify the gaps in my research. Simona was the heart of our A-WEAR network—without her knowledge and dedication, I believe the connections we created in the network would not have been as strong.

The thesis would not have been possible without my collaborators. A special thank you goes to Boris Danev, Silvan Wehrli, and David Barras from 3db Access, who have always been eager to share their knowledge about UWB with me. I would also like to thank Iuliu Vasilescu, Cristi Pătru, Michele Magno, and Jari Nurmi for our fruitful collaboration.

I am grateful to everyone in the A-WEAR network—you showed me the true power of networking and teamwork. In particular, Viktoriia Shubina has been not only a wonderful colleague but also a dear friend who made me rediscover the meaning of “girl power.” I would like to thank Aleksandr Ometov, who went above and beyond to ensure that every ESR stayed sane during their studies. His direct and honest career advice has always been

refreshing and much-appreciated. Kate Svertoka and Asma Channa have been lovely colleagues with whom I shared many special moments.

My PhD life would have not been the same without my amazing friends and colleagues at both UPB and TAU. I would especially like to thank Mihai Dumitru, who always challenges me to become a better person (and a better debater and jokester at that).

My deepest gratitude goes to my parents, who have offered me unconditional love and support throughout my life. Thank you for always prioritizing my education, even when it required financial sacrifices. Although he cannot read this acknowledgment, a special thanks goes to my cat, Tobi, for lifting my spirits and giving me an incentive to come home early from work.

Last but most important, I would like to thank my husband, Călin, for always believing in me. These three years have been tough, but somehow we succeeded to make the best of them despite the distance, a global pandemic, and the hurdles of PhD life. I am excited for the adventures awaiting us.

Bucharest, November 2022

Laura Fluerătoru

ABSTRACT

Ultra-wideband (UWB) technology has been rediscovered in recent years for its potential to provide centimeter-level accuracy in GNSS-denied environments. The large-scale adoption of UWB chipsets in smartphones brings demanding needs on the energy-efficiency, robustness, scalability, and cross-device compatibility of UWB localization systems. This thesis investigates, characterizes, and proposes several solutions for these pressing concerns.

First, we investigate the impact of different UWB device architectures on the energy efficiency, accuracy, and cross-platform compatibility of UWB localization systems. The thesis provides the first comprehensive comparison between the two types of physical interfaces (PHYs) defined in the IEEE 802.15.4 standard: with low and high pulse repetition frequency (LRP and HRP, respectively). In the comparison, we focus not only on the ranging/localization accuracy but also on the energy efficiency of the PHYs. We found that the LRP PHY consumes between 6.4–100 times less energy than the HRP PHY in the evaluated devices. On the other hand, distance measurements acquired with the HRP devices had 1.23–2 times lower standard deviation than those acquired with the LRP devices. Therefore, the HRP PHY might be more suitable for applications with high-accuracy constraints than the LRP PHY.

The impact of different UWB PHYs also extends to the application layer. We found that ranging or localization error-mitigation techniques are frequently trained and tested on only one device and would likely not generalize to different platforms. To this end, we identified four challenges in developing platform-independent error-mitigation techniques in UWB localization, which can guide future research in this direction.

Besides the cross-platform compatibility, localization error-mitigation techniques raise another concern: most of them rely on extensive data sets for

training and testing. Such data sets are difficult and expensive to collect and often representative only of the precise environment they were collected in. We propose a method to detect and mitigate non-line-of-sight (NLOS) measurements that does not require any manually-collected data sets. Instead, the proposed method automatically labels incoming distance measurements based on their distance residuals during the localization process. The proposed detection and mitigation method reduces, on average, the mean and standard deviation of localization errors by 2.2 and 5.8 times, respectively.

UWB and Bluetooth Low Energy (BLE) are frequently integrated in localization solutions since they can provide complementary functionalities: BLE is more energy-efficient than UWB but it can provide location estimates with only meter-level accuracy. On the other hand, UWB can localize targets with centimeter-level accuracy albeit with higher energy consumption than BLE. In this thesis, we provide a comprehensive study of the sources of instabilities in received signal strength (RSS) measurements acquired with BLE devices. The study can be used as a starting point for future research into BLE-based ranging techniques, as well as a benchmark for hybrid UWB-BLE localization systems.

Finally, we propose a flexible scheduling scheme for time-difference of arrival (TDOA) localization with UWB devices. Unlike in previous approaches, the reference anchor and the order of the responding anchors changes every time slot. The flexible anchor allocation makes the system more robust to NLOS propagation than traditional approaches. In the proposed setup, the user device is a passive listener which localizes itself using messages received from the anchors. Therefore, the system can scale with an unlimited number of devices and can preserve the location privacy of the user. The proposed method is implemented on custom hardware using a commercial UWB chipset. We evaluated the proposed method against the standard TDOA algorithm and range-based localization. In line of sight (LOS), the proposed TDOA method has a localization accuracy similar to the standard TDOA algorithm, down to a 95% localization error of 15.9 cm. In NLOS, the proposed TDOA method outperforms the classic TDOA method in all scenarios, with a reduction of up to 16.4 cm in the localization error.

CONTENTS

1	Introduction	21
1.1	Background and Motivation	21
1.2	Objectives and Scope of the Thesis.	23
1.3	Outline and Structure of the Thesis	24
1.4	Author’s Contributions to the Publications	24
2	Introduction to Indoor Localization with UWB devices	27
2.1	UWB Technology.	27
2.2	Types of measurements	31
2.2.1	Distance Measurements.	31
2.2.2	Time-Difference of Arrival Measurements	33
2.3	Localization Methods	36
3	Different UWB Device Architectures in Localization Systems	39
3.1	State of the Art	39
3.2	LRP vs. HRP	42
3.2.1	Architecture	43
3.2.2	Devices Used for Measurements.	44
3.2.3	Power and Energy Consumption	46
3.2.4	Ranging Accuracy	48
3.2.5	Localization Accuracy.	50
3.3	Challenges in Platform-Independent UWB Localization Systems	53
3.3.1	Same Conditions, Different Distance Errors for Different Devices	55
3.3.2	Same Conditions, Different TOAs	57
3.3.3	Same Conditions, Different CIRs.	59
3.3.4	Different Devices, Different Diagnostics	62

3.4	Summary	63
4	NLOS and Multipath Propagation in Wideband and Narrowband Communications.	65
4.1	State of the Art	66
4.1.1	SOTA on UWB Error Mitigation Techniques	69
4.1.2	SOTA on BLE RSS Variability	72
4.2	Self-learning NLOS Error Detection and Mitigation for UWB Localization Systems	74
4.2.1	Overview of the Proposed Framework	74
4.2.2	Simulation Framework	76
4.2.3	Results	79
4.3	Signal Fluctuations in BLE Communication.	82
4.3.1	Variations in Time	82
4.3.2	Variations Caused by Hardware	83
4.3.3	Variations with Orientation.	85
4.3.4	Variations in LOS and NLOS Propagation	85
4.4	Summary	87
5	TDOA Localization with Flexible Anchor Allocation	89
5.1	State of the Art	91
5.1.1	TOA and TDOA Comparisons.	91
5.1.2	TDOA Localization	91
5.1.3	Clock offset correction	94
5.2	Improved TDMA Scheduling for TDOA Localization	95
5.3	Evaluation Setup.	97
5.4	Results	99
5.4.1	Optimal Number of Responses per Time Slot	99
5.4.2	Channel Diversity	101
5.4.3	NLOS Propagation	104
5.5	Summary	106
6	Conclusions	109
6.1	Main Results	109
6.2	Future Research Directions	113

References	115
[P1]	135
[P2]	145
[P3]	165
[P4]	175
[P5]	185
[P6]	203

List of Figures

2.1	Time resolution and bandwidth in UWB communications	28
2.2	UWB applications.	29
2.3	Practical localization scenarios in everyday life.	30
2.4	Message exchange in the SS- and DS-TWR.	32
2.5	Synchronous DL TDOA	33
2.6	Synchronous UL TDOA	34
2.7	Asynchronous DL TDOA	35
3.1	Peak and average power level in LRP and HRP PHYs.	43
3.2	Setup of ranging measurements in (a) LOS, (b) NLOS with a drywall, and (c) NLOS with a concrete wall. The UWB devices are placed on tripods. The NLOS with human body shadowing setup is identical with the LOS one, except that a person is standing right in front of the transmitter (the device further away)..	48
3.3	Distributions of ranging errors of 3db and DW1000 devices.. . . .	49
3.4	Comparison of ranging errors at individual locations.	50

3.5	Localization setup: the four anchors (A1 to A4) encompass an area of approximately 4.5×3.6 m and the tracking area is on the table. The ground truth was acquired using two HTC Vive base stations (BS1 and BS2) and a tracker that was colocated with the UWB tag, shown on the table.	51
3.6	Examples of environments from the measurement campaign.	55
3.7	CDF of errors in LOS and NLOS for all devices.	56
3.8	Distribution of errors from selected locations.	57
3.9	Stacked CIRs with different alignments.	58
3.10	CIRs acquired with each device at the same location.	61
4.1	Effect of MPP on NB and UWB communication.	67
4.2	BLE and WiFi channels.	73
4.3	The main steps of the proposed error-mitigation technique.	74
4.4	Location estimates with ideal vs. noisy measurements.	75
4.5	(a) Simulated localization setup and (b) data generation process based on the measurement database.	77
4.6	Label generation process using residual analysis.	78
4.7	Model training and testing based on generated labels.	78
4.8	The accuracy of the trained random forest (RF) classifier using data labeled with the residual method.	80
4.9	The pictures of the receiver and transmitter in LOS and NLOS with wall acquired at UPB and TAU (Figure 4.9a to 4.9d). Figure 4.9e to 4.9h show NLOS scenarios with a plexiglass panel, human body, a door, and a wall with a whiteboard at UPB.	83
4.10	The evolution of the BLE RSS distribution over time.	84
4.11	The impact of hardware on the RSS in recordings acquired with different device pairs.	84
4.12	The impact of orientation on the BLE RSS.	85
4.13	Comparison of selected RSS distributions.	86
4.14	Comparison of BLE RSS distributions in multiple LOS/NLOS scenarios at UPB and TAU.	87
5.1	The scheduling of anchor transmissions in classic DL TDOA and FlexTDOA.	90

5.2	TDMA scheduling with changing initiator and responders.	96
5.3	Office setup with ten anchors (A_0 to A_9) and one tag.	97
5.4	Setup of the anchors	98
5.5	TDOA measurement and localization errors depending on the number of responses.	101
5.6	Comparison between fixed vs. changing initiator and/or order of responses for different numbers of anchors and responders. . .	102
5.7	Distributions of the localization errors of classic TDOA (FI-FR) and FlexTDOA (CI-CR), in LOS, aggregated over all the evaluated positions (P1, P2, P3).	103
5.8	Setups in which NLOS measurements were acquired.	104
5.9	Comparison of the classic TDOA, FlexTDOA, and TWR localization in NLOS.	105

List of Tables

3.1	The most popular UWB devices on the market.	40
3.2	The average power consumption of 3db and DW1000 devices. . .	47
3.3	Statistics of distance measurement errors.	51
3.4	Statistics of measurement-based localization errors.. . . .	52
3.5	Statistics of the CIRs.	62
4.1	Comparison of UWB and BLE accuracy in ranging and localization applications.	68
4.2	Performance of the unsupervised residual-based labeling.. . . .	79
4.3	UWB localization error before and after the proposed NLOS mitigation method	81

ABBREVIATIONS

AOA	Angle of arrival
BFSK	Binary phase-shift keying
BLE	Bluetooth Low-Energy
CIR	Channel impulse response
DS-TWR	Double-sided two-way ranging
HRP	High-rate pulse repetition
IC	Integrated circuit
IoT	Internet of things
IQR	Inter-quartile range
LOS	Line of sight
LRP	Low-rate pulse repetition
MPP	Multipath propagation
NLOS	Non-line of sight
PHY	Physical layer
PRF	Pulse repetition frequency
PRR	Packet reception rate
PSD	Power spectral density
PSK	Phase-shift keying
SOTA	State of the art
SS-TWR	Single-sided two-way ranging
TDMA	Time-division multiple access

TDOA	Time-difference of arrival
TOA	Time of arrival
TOF	Time of flight
TWR	Two-way ranging
ULP	Ultra-low power
UWB	Ultra-wideband

ORIGINAL PUBLICATIONS

- [P1] L. Flueratoru, S. Wehrli, M. Magno, and D. Niculescu. “On the energy consumption and ranging accuracy of ultra-wideband physical interfaces”. In: *IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2020, pp. 1–7. doi: 10.1109/GLOBECOM42002.2020.9347984.
- [P2] L. Flueratoru, S. Wehrli, M. Magno, E. S. Lohan, and D. Niculescu. “High-Accuracy Ranging and Localization With Ultrawideband Communications for Energy-Constrained Devices”. In: *IEEE Internet of Things Journal* 9.10 (2021), pp. 7463–7480. doi: 10.1109/JIOT.2021.3125256.
- [P3] L. Flueratoru, E. S. Lohan, and D. Niculescu. “Challenges in Platform-Independent UWB Ranging and Localization Systems”. In: *The 16th ACM Workshop on Wireless Network Testbeds, Experimental evaluation and Characterization (WiNTECH)*. ACM, 2022, pp. 1–8. doi: 10.1109/3556564.3558238.
- [P4] L. Flueratoru, E. S. Lohan, and D. Niculescu. “Self-Learning Detection and Mitigation of Non-Line-of-Sight Measurements in Ultra-Wideband Localization”. In: *2021 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2021, pp. 1–8. doi: 10.1109/IPIN51156.2021.9662532.
- [P5] L. Flueratoru, V. Shubina, E. S. Lohan, and D. Niculescu. “On the High Fluctuations of Received Signal Strength Measurements with BLE Signals for Contact Tracing and Proximity Detection”. In: *IEEE Sensors Journal* 22.6 (2021), pp. 5086–5100. doi: 10.1109/JSEN.2021.3095710.

- [P6] George-Cristian Pătru, Laura Flueratoru, Iuliu Vasilescu, Dragoș Niculescu, and Daniel Rosner. “FlexTDOA: Robust and Scalable Time-Difference of Arrival Localization Using Ultra-Wideband Devices”. In: *IEEE Access* 11 (2023), pp. 28610–28627. doi: 10.1109 / ACCESS.2023.3259320.

1 INTRODUCTION

1.1 Background and Motivation

Navigation applications are integral to our lives, allowing us to reach our destinations faster, interact with our environment, discover new places, or even save our lives in an emergency. While GNSS services have become the *de facto* standard for outdoor positioning, in indoor spaces there are many competing technologies, each with their advantages and drawbacks. In recent years, ultra-wideband (UWB) has emerged as a promising technology that can enable localization with centimeter-level accuracy. The advent of UWB technology was facilitated by its adoption in smartphones, starting with the Apple iPhone 11 in 2019. Since then, more smartphone developers such as Google, Samsung, or Xiaomi have followed the trend and deployed UWB chipsets in their latest smartphone releases. The large-scale adoption of the technology raises the stakes and creates new ambitious needs for future localization applications based on UWB.

With so many UWB solutions on the market, there is a need for comparing the performance of the physical-layer interfaces (PHYs) defined in the IEEE 802.15.4 standard, primarily in terms of localization accuracy and energy-efficiency. While the accuracy has been the main focus of previous research, not much attention had been dedicated to the energy-efficiency of UWB devices. Since UWB devices are most often used in battery-powered devices, it is important to evaluate the trade-off between localization accuracy and energy-efficiency that can satisfy the desired requirements of an application.

Although UWB can deliver centimeter-level accuracy in ideal conditions, the quality of location estimates usually degrades to decimeter-level in the presence of obstructions or multipath propagation (MPP). While there is a large body of research in the area of error mitigation techniques, most such

methods rely on time-consuming and expensive measurement campaigns in order to collect training data [69, 134, 102, 7, 11]. In order for UWB localization systems to be deployed on a large scale, they have to be robust in the presence of environmental noise but, most importantly, easy to deploy and maintain.

Connecting these first two pressing issues in UWB localization, there is a need for error mitigation techniques that are device-agnostic and can therefore be used irrespective of the model of UWB device carried by the user [103]. Although standards [47] and certifications [15] that ensure the interoperability of UWB devices from different developers are currently in place, not much attention has been dedicated to the cross-compatibility of error-mitigation algorithms at the *application* layer.

Given these issues, UWB localization systems that are scalable to an unlimited number of users, highly accurate, energy-efficient, and robust to MPP and non-line-of-sight (NLOS) scenarios *at the same time* are hard to develop. Localization systems based on two-way ranging (TWR) techniques offer centimeter-level accuracy but do not scale well with a high number of anchors and tags [116]. Time-difference of arrival (TDOA) localization can solve the scalability issues to a certain extent but needs specialized scheduling algorithms to avoid synchronizing devices [31] and is usually more prone to measurement noise than its TWR counterpart [118]. Therefore, more research is needed to find a good balance between all the desirable requirements of a UWB localization system.

With so many standardized communication protocols available, out of which the most notable are WiFi and Bluetooth, UWB may not arise as the obvious best solution for indoor localization. In particular, Bluetooth Low Energy (BLE) is often a preferred alternative especially for battery-constrained wearables owing to its ultra-low power consumption [8]. However, being a narrowband (NB) technology, BLE ranging can currently achieve at best only decimeter-level accuracy [5, 77, 86, 3, 26]. It is therefore worthwhile to look more into detail at the advantages and disadvantages of UWB and BLE in ranging and/or localization in challenging indoor environments.

1.2 Objectives and Scope of the Thesis

In line with the identified issues in current indoor localization systems, the purpose of the thesis is to advance the research on robust, energy-efficient, and platform-independent indoor localization systems. To this end, we identified the following research questions:

1. *How do UWB physical interfaces compare in terms of energy efficiency, maximum range, and accuracy?*

We compare in [P1] and [P2] the two PHYs defined in the IEEE 802.15.4 standard: with low- and high-pulse repetition (LRP and HRP, respectively) and their implementations in commercial UWB devices.

2. *What implications do different UWB hardware architectures have on the cross-platform compatibility of localization systems?*

In [P3], we identified four challenges currently standing in the way of platform-independent localization systems based on an extensive measurement campaign performed with three different UWB platforms.

3. *How can we mitigate the impact of NLOS propagation on UWB localization accuracy without manually-collected training data?*

The thesis proposes in [P4] a novel technique for detecting and mitigating measurements acquired in NLOS without the need to perform time-consuming measurement campaigns.

4. *How does UWB compare to complementary low-power technologies, such as BLE, in terms of ranging performance?*

The thesis includes a study [P5] on the variability of BLE received signal strength (RSS) measurements, which can be used as a starting point for this comparison.

5. *How can we achieve scalable and robust UWB localization?*

In [P6], we propose a flexible TDMA scheme for TDOA localization with UWB devices which scales up to an unlimited number of users and is robust to NLOS propagation owing to its enhanced channel diversity compared to existing approaches in the literature.

1.3 Outline and Structure of the Thesis

The present thesis is organized into three main chapters, as follows:

Chapter 2 introduces the basic principles of UWB communication and the most frequently-used ranging and localization methods, which are a prerequisite for presenting our contributions.

Chapter 3 addresses the impact of different UWB device architectures in localization systems, especially concerning the ranging/localization accuracy, the energy-efficiency, and the cross-platform compatibility of the systems. The chapter is based on publications [P1], [P2], and [P3].

Chapter 4 examines the main sources of errors in UWB and BLE localization, which can be viewed as complementary technologies. UWB can be used for fine-grained localization with the disadvantage of a high energy consumption, while BLE can be used with a high energy efficiency to provide a rough estimate of the location of the tracked device. This chapter follows the work in publications [P4] and [P5].

Chapter 5 presents FlexTDOA, a TDOA localization system with a flexible TDMA scheduling scheme for anchor transmissions. The chapter covers the principles of the proposed TDMA scheme, the most important results achieved by the system in both LOS and NLOS propagation, and presents directions for improving the system in the future.

Finally, Chapter 6 summarizes the main contributions of this thesis and presents promising directions for future research.

1.4 Author's Contributions to the Publications

The thesis consists of the work presented in three journal papers [P2], [P5], [P6], two conference papers [P1], [P4], and one workshop paper [P3]. The Author of the thesis was the main author in the majority of the papers [P1]– [P4], shared an equal contribution with the second author of publication [P5], Viktoriia Shubina, and was the second author of publication [P6]. All the publications were the result of a tight collaboration with several researchers, whose input was essential throughout the projects.

The Author's main supervisors, Prof. Dragoş Niculescu and Prof. Simona

Lohan from University Politehnica of Bucharest (UPB) and Tampere University (TAU), respectively, contributed to all the publications by guiding the research direction, offering feedback during the data analysis and writing process, and generally sharing their ideas throughout the research projects.

In the papers in which the Author is the main author [P1]– [P4], the Author carried out the experiments for data collection, implemented the proposed localization algorithms, analyzed and interpreted the data, and was the main writer of the papers.

In [P1] and [P2], D.Sc. Silvan Wehrli implemented the backbone of the localization network on top of which the Author added new features, such as a new synchronization mechanism between anchors and new localization algorithms. In addition, he provided valuable insight into the UWB technology and feedback throughout the project. Dr. Michele Magno offered his supervision for publications [P1] and [P2], especially regarding the power and energy consumption measurements and their analysis.

The Author and Viktoriia Shubina share an equal contribution to publication [P5]. The publication is based on two measurement campaigns using BLE devices, which were performed in parallel at UPB and TAU. The measurements were carried out by the Author and Viktoriia Shubina, respectively. Therefore, each main author was responsible for their own data collection based on a commonly-agreed setup. The authors identified together eight key factors that lead to the instability of BLE RSS measurements: time, space, hardware, advertising channel, distance between the transmitter (TX) and the receiver (RX), device orientation, BLE–WiFi combo chipsets, and NLOS obstructions. When analyzing the data for each of the identified factors, the main authors divided their work equally: the Author was the principal responsible for the sections on the impact of the advertising channel, NLOS propagation, device orientation, and TX–RX distance on RSS measurements, while Viktoriia Shubina overlooked the analysis of the rest of the factors.

In publication [P6], the first author, Cristian Pătru, developed the hardware and software for the TDOA localization system under the close supervision of D.Sc. Iuliu Vasilescu. The Author identified the main gaps in the research on UWB TDOA localization systems and proposed a suitable research

direction. In addition, the Author designed the setup of the experiments, helped with the data collection, analyzed the data, and was the main writer of the publication. The Author contributed to at least 30% of the work for publication [P6]. Prof. Dragoş Niculescu and D.Sc. Daniel Rosner provided valuable feedback during the research and the writing process.

Besides the main publications included in the thesis [P1]– [P6], the Author was the main author of the work in [28], which proposes a ground-truth system for UWB localization based on the HTC Vive motion capture system. The HTC Vive is 10 times cheaper than professional-grade motion capture systems and still localizes objects with millimeter-level accuracy. The Author also contributed to a survey on the state-of-the-art (SOTA) and current challenges in wearables [83]. The Author was responsible for the sections related to the localization of wearable devices.

All the publications included in the thesis are available in an open-access format. In addition, the data sets used in publications [P2], [P3], and [P5] have also been published in an open-source format [56, 55, 57].

2 INTRODUCTION TO INDOOR LOCALIZATION WITH UWB DEVICES

In this chapter, we discuss the basic principles of ranging and localization using UWB communications, which provide a foundation for presenting the contributions of the thesis in the following chapters. Section 2.1 presents an introduction of UWB technology, the features that enable its accurate timing estimation, and the main applications of the technology. Section 2.2 discusses the principles of distance and time-difference of arrival measurements with UWB devices, with an emphasis on practical aspects. Finally, Section 2.3 introduces the framework for location estimation using distance and TDOA measurements.

2.1 UWB Technology

According to international regulations [12], UWB devices must have a bandwidth of at least 500 MHz, which means that they transmit pulses with a very high time resolution, on the order of 1 ns (Figure 2.1). This makes them particularly suitable for accurate time-based ranging and localization methods [131, 34].

The range of UWB communication varies from 30 m to over 100 m [2]. Therefore, they are most suited for short- to medium-range indoor applications, although previous works investigated their use in outdoor environments as well [137].

Another desirable feature of the UWB technology is that it can penetrate obstacles, so it is also suitable for through-the-wall communication in multi-room environments. However, its ranging accuracy degrades in NLOS propagation, which has led to extensive research into methods that mitigate NLOS errors [69, 134, 102]. We investigate the issue of NLOS propagation in UWB




Figure 2.1 UWB communications have a very large bandwidth, of over 500 MHz, which means that their pulses have a very high time resolution, of under 1 ns. Therefore, they can provide high accuracy in time-based ranging and localization methods.

communication in more detail in Chapter 4. In [P4], we propose a NLOS detection and mitigation method that does not need any manually-collected data, which is often a pre-requisite of such techniques.

The FiRa™ consortium lists four main industries for UWB applications: smart cities & mobility, smart building & industrial, smart retail, and smart home & consumer [14]. Figure 2.2 presents examples of applications in each area. Navigation is a key application of UWB with different requirements depending on the exact use case: for people, in autonomous driving, in drone delivery, asset tracking, or gaming. For instance, asset tracking and manufacturing applications might need highly accurate and robust localization, but they usually benefit from a controlled environment in which the sensors are deployed. For personal navigation, decimeter-level accuracy might suffice in most cases but the sparse and non-uniform deployment of sensors might pose issues. Autonomous driving and drone delivery localization applications need to deal with fast-moving objects and likely ensure cooperation between multiple sensing and localization technologies (e.g., GNSS, UWB, radar, inertial measurement unit). Gaming applications similarly need to cater to fast movements and also provide localization with low latency.

In secure building and residential access, as well as in payment systems, security is a critical component. Indeed, security is a central aspect of the latest amendment of the IEEE 802.15.4 standard [47] and the security of UWB PHYs has been analyzed in previous works [111, 112]. Key-less car access is already a large market for UWB chipsets and new studies aim to incorporate



Figure 2.2 Applications of UWB technology in four key industries: smart cities & mobility, smart building & industrial, smart retail, and smart home & consumer [14]. The figure lists examples of applications for each industry. We can identify several classes of applications: navigation, access, tracking, and sensing.



Figure 2.3 Practical localization scenario in everyday life. Several UWB anchors are deployed as part of the infrastructure. The user device (also called a tag) needs to be equipped with a UWB transceiver. The localization system provides the user either their position using global coordinates, local coordinates (relative to the building), or as a label depending on its proximity to surrounding landmarks.

sensing into the access functionality [53, 67].

In smart home applications, UWB devices deployed in home appliances can enable custom functionalities depending on the room occupancy or on the precise location of the user inside the room [43, 44, 133].

Although the localization infrastructure depends on the exact type of application, it generally consists of several devices with fixed and known locations which allow user localization in their local coordinate system. These devices are usually called *anchors* and deployed around the tracking area, ideally such that the user moves within the convex hull formed by the anchors. The anchors in an UWB localization system can be likened to satellites in GNSS, as they also provide a reference for positioning the user device. If the global coordinates of the anchors are known, the local coordinates of the tag can be transformed into global coordinates, allowing the user to seamlessly navigate outdoors as well as indoors. This setup is illustrated in Figure 2.3, which shows the anchors, the user device (also called the *tag*) and its local

and global coordinates. In other applications, instead of the local or global coordinates, the user might actually need a label such as the name of the room in which it is positioned. Therefore, localization applications can also incorporate other types of local cues.

2.2 Types of measurements

In this section, we discuss the most common time-based types of measurements acquired with UWB devices, namely distance and TDOA measurements. We focus on the practical aspects of performing these measurements using accurate timing information provided by UWB devices.

2.2.1 Distance Measurements

The distance between two devices can be estimated based on the time of flight (TOF) of the signal. If T_1 represents the transmission time of the signal measured by the sender and T_2 represents the arrival time of the signal at the receiver, the distance between the sender and the receiver is [106]:

$$d = (T_2 - T_1) \cdot c = T_p \cdot c, \quad (2.1)$$

where c is the speed of light and $T_p \triangleq T_2 - T_1$ is the propagation time of the signal. In reality, the clocks of the sender and the receiver run at slightly different speeds, so the devices would have to be clock-synchronized in order to estimate the TOF using Eq. (2.1). Because synchronizing the sender and the receiver to nanosecond-level (needed to achieve centimeter-level accuracy) is usually unfeasible in practice, more messages are exchanged in order to eliminate the clock synchronization requirements. This leads to ranging techniques such as single- or double-sided two-way ranging (SS-TWR and DS-TWR, respectively).

Single-Sided Two-Way Ranging. The SS-TWR uses two messages per distance estimate, as shown in Figure 2.4a. The propagation time can be computed as:

$$T_p = \frac{T_{round} - T_{proc}}{2}, \quad (2.2)$$




Figure 2.4 Message exchange in the single- and double-sided two-way ranging.

where T_{round} is the time spent in one message exchange and T_{proc} is the processing time on the responder side.

The main sources of errors in SS-TWR are T_{proc} , which ranges from hundreds of μ s to several ms [21], and the clock drift, which can be up to ± 20 ppm in systems compliant with the IEEE 802.15.4 standard [47].

Double-Sided Two-Way Ranging. Because in SS-TWR the errors increase proportionally to the processing time, one way to increase the ranging accuracy is to use an additional message denoted by FINAL in Figure 2.4b. The TOF can then be computed as:

$$T_p = \frac{T_{round1} - T_{proc1} + T_{round2} - T_{proc2}}{4}, \quad (2.3)$$

where T_{round1} is the time spent in the first message exchange (consisting of messages REQ and RESP) by the initiator, T_{round2} is the time spent in the second message exchange (consisting of messages RESP and FINAL) by the responder, and T_{proc1} and T_{proc2} are the processing times at the responder and initiator, respectively. The clock error in Eq. 2.3 is minimized when the processing times T_{proc1} and T_{proc2} are equal [76]. When this condition is not satisfied, an alternative formula for computing the TOF has been proposed in [76]:

$$T_p = \frac{T_{round1}T_{round2} - T_{proc1}T_{proc2}}{T_{round1} + T_{proc1} + T_{round2} + T_{proc2}}, \quad (2.4)$$

which minimizes the clock drift error even in the presence of asymmetrical processing times.




Figure 2.5 Synchronous DL TDOA. The anchors transmit the POLL messages simultaneously. Because the tag is closer to A_j than to A_i , the message POLL_j arrives the first. The TDOA is computed according to Eq. (2.6).

2.2.2 Time-Difference of Arrival Measurements

Whereas distance measurements are useful in themselves, since they can be used in proximity-based applications such as secure access or payment systems (Figure 2.2), TDOA measurements are generally used only in localization algorithms. In the general case, the TDOA between two anchors A_i and A_j and the tag T is:

$$t_{TA_iA_j} = \frac{1}{c}(d_{TA_j} - d_{TA_i}), \quad (2.5)$$

where d_{TA_k} is the distance between the tag and anchor A_k . This places the tag on a hyperbola with foci at A_i and A_j .

Conventionally, TDOA measurements required synchronization at the reference clocks [98]. However, in recent years, various *asynchronous* TDOA methods have been proposed [58, 88, 137]. We will first present the synchronous techniques to obtain TDOA measurements, after which we will examine the modifications needed to bypass the synchronization requirements.

Synchronous TDOA. If A_i and A_j were clock-synchronized, the TDOA could be obtained either on the downlink (DL) channel (from the anchors to tag) or on the uplink (UL) channel (from the tag to the anchors). In synchronous DL TDOA, presented in Figure 2.5, the anchors A_i and A_j can simultaneously transmit a message which will be received by the tag at times




Figure 2.6 Synchronous UL TDOA. The tag transmits the POLL message, which arrives at the anchors A_i and A_j at times T_{A_i} and T_{A_j} , respectively. The TDOA is computed according to Eq. (2.7).

T_i and T_j , respectively. In this case, the TDOA can be obtained as:

$$t_{T_{A_i}A_j} = T_j - T_i. \quad (2.6)$$

Note that in DL TDOA the tag is passive and does not need to transmit any messages. Therefore, DL TDOA methods can scale up to an unlimited number of users.

In UL TDOA, presented in Figure 2.6, the tag transmits a message which is received by the anchors A_i and A_j at times T_{A_i} and T_{A_j} , respectively. If the anchors are clock-synchronized, then T_{A_i} and T_{A_j} will have the same time reference. The TDOA can then be obtained as:

$$t_{T_{A_i}A_j} = T_{A_j} - T_{A_i}. \quad (2.7)$$

In UL TDOA, the anchors usually need to transmit the times at which they received the tag's broadcast to a central server, which will compute the tag's location. The advantage of UL TDOA is that the tag needs to transmit a single message, whereas in DL TDOA it needs to receive N messages, where N is the number of anchors in the system. Since the receive mode has a higher average power consumption than the transmit mode [21], the energy consumption of the tag is higher in DL TDOA than in UL TDOA.

Asynchronous TDOA. In asynchronous TDOA (A-TDOA) schemes, the clocks of anchors do not need to be synchronized. Multiple such variants




Figure 2.7 Asynchronous DL TDOA. One anchor called the reference anchor (in this case, A_i) broadcasts a message (POLL). Anchor A_j receives the message, processes it, and sends a response (RESP). The tag receives the POLL and RESP messages at times T_i and T_j , respectively. The TDOA is computed using Eq. 2.8. In the figure, $t_{A_i A_j}$ is the TOF between anchors A_i and A_j and t_{proc} is the processing time of the POLL message at anchor A_j .

have been proposed, but we will focus on asynchronous *DL* TDOA [58, 88, 137] which, similarly to its synchronous counterpart, uses only transmissions on the DL channel so that the tag can remain passive.

DL A-TDOA is illustrated in Figure 2.7, where A_i transmits the initial POLL message, which is received by both the Tag and anchor A_j . Upon the reception of the POLL, A_j processes the message and sends the response (RESP). The tag receives the POLL and RESP messages at times T_i and T_j , respectively. We denote the difference of the timestamps by $T_{ij} \triangleq T_j - T_i$. We can obtain the TDOA from Eq. (2.5) by subtracting the processing time at A_j (denoted by t_{proc}) and the TOF between A_i and A_j (denoted by $t_{A_i A_j}$) from the timestamp difference T_{ij} as follows:

$$t_{TA_i A_j} = T_j - T_i - t_{proc} - t_{A_i A_j}. \quad (2.8)$$

Because the anchors are usually placed at fixed and known locations, the TOF between the anchors ($t_{A_i A_j}$) is known. DL A-TDOA therefore allows a localization system to scale with an unlimited number of users without any clock synchronization between the anchors, which would otherwise increase the complexity of the system.

The drawback of using the asynchronous TDOA version instead of the synchronous one is that A-TDOA is affected by the relative clock skew be-

tween A_j and the tag. This is manifested when subtracting the processing time t_{proc} , which is measured by A_j and is on the order of hundreds of μs , from the timestamp difference T_{ij} measured by the tag. To correct these errors, we can estimate the relative clock offset between A_j and the tag and use this value to correct the processing time [20]. This is the method employed in publication [P6] to reduce clock errors in DL A-TDOA.

In a practical localization system, one of the anchors can be designated as the *reference* anchor (or the initiator) which sends the POLL message. Once the other anchors in the system receive the POLL message, they respond successively with the RESP messages. In practice, anchors are usually programmed to respond with different delays δ_i based on their index i , in order to avoid overlapping receptions at the tag [35].

2.3 Localization Methods

The position of the tag can be estimated based on distance or TDOA measurements between the tag and multiple anchors described in Section 2.2.

Let $\mathbf{x}_{A_i} = (x_{A_i}, y_{A_i}, z_{A_i})$ be the position of anchor A_i in a Cartesian coordinate system with axes x, y, z . Similarly, \mathbf{x}_T is the position of the tag in the same coordinate system. We will use a bold script to denote vectors, the subscript T to refer to the tag, and the subscript A_i to refer to the anchor with index $i \in \{1, \dots, N\}$, where N is the number of anchors in the system.

Range-based localization. A distance measurement between the tag and one anchor places the tag anywhere on a circle centered at the anchor with a radius equal to the measured distance [98]. Using three such measurements, we can localize the tag in a 2D coordinate system; similarly, we need at least four distance measurements to position the tag in a 3D coordinate system.

If we denote by d_i the true distance between the anchor A_i and the tag, then for range-based localization a system of equations:

$$d_{TA_i} = \sqrt{(x_{A_i} - x_T)^2 + (y_{A_i} - y_T)^2 + (z_{A_i} - z_T)^2} \quad (2.9)$$

$$\triangleq \|\mathbf{x}_T - \mathbf{x}_{A_i}\|, \quad (2.10)$$

needs to be solved for all M equations, where $i = 1, \dots, M$ and M is equal to the number of anchors in the system (denoted by N).

TDOA localization. A TDOA measurement between a tag and two anchors places the tag anywhere on the hyperbola with foci at the two anchors. As in the case of range-based localization, $M + 1$ equations are needed to localize the tag in an M -dimensional space. However, because two anchors determine one TDOA, the system needs at least $M + 2$ anchors for M -dimensional localization. In TDOA localization, we need to solve the system of equations:

$$d_{TA_i A_j} = d_{TA_j} - d_{TA_i} \quad (2.11)$$

$$= \|\mathbf{x}_T - \mathbf{x}_{A_j}\| - \|\mathbf{x}_T - \mathbf{x}_{A_i}\|, \quad (2.12)$$

for $i = 1, \dots, M$ and $M = N - 1$. In the noiseless case, any $N - 1$ TDOA measurements that form a minimum spanning subtree are sufficient for TDOA localization [114].

Location estimation. Because, in practice, both distance and TDOA measurements are noisy, we need to employ statistical techniques to solve the localization problem [98]. We consider the measurement model [98]:

$$z_i = f_i(\mathbf{x}_T) + n_i, \quad i = 1, \dots, M, \quad (2.13)$$

where n_i is the measurement noise and f_i is the measurement equation defined as:

$$f_i = \begin{cases} \|\mathbf{x}_T - \mathbf{x}_{A_i}\|, & \text{for range-based localization} \\ \|\mathbf{x}_T - \mathbf{x}_{A_i}\| - \|\mathbf{x}_T - \mathbf{x}_{A_0}\|, & \text{for TDOA localization.} \end{cases} \quad (2.14)$$

For TDOA, we denoted by A_0 the reference anchor which initiates the localization procedure.

In vector notations, we obtain the following measurement model:

$$\mathbf{y} = h(\mathbf{x}_T) + \mathbf{v}, \quad (2.15)$$

where $\mathbf{y} = (z_1, \dots, z_M)$ is the measurement vector, $\mathbf{v} = (v_1, \dots, v_M)$ is the error vector, and $h(\cdot)$ is the vector-valued measurement function. The goal of a localization algorithm is to find the position \mathbf{x}_T^* such that:

$$\mathbf{x}_T^* = \arg \min_{\mathbf{x}_T} \|\mathbf{y} - h(\mathbf{x}_T)\|, \quad (2.16)$$

where \mathbf{x}_T^* is the least-squares solution to the localization problem. Since the measurement equations are nonlinear and do not have a closed-form solution, the minimization can be done either via numerical search algorithms such as gradient descent or Gauss-Newton techniques [36] or by linearizing the equations using techniques such as the simple intersection method [9] or the Range-Bancroft method [113]. Algorithms that perform a local search require a good initialization to avoid converging to a local minimum of the loss function [36]. In practice, the algorithms can be initialized using the location found through one of the closed-form solutions obtained through linearization [113] or, in the case of continuous tracking, using the last estimated location. Another approach to solve the localization problem is to track the user's location over time using tracking filters such as the extended Kalman filter or the particle filter [98].

3 DIFFERENT UWB DEVICE ARCHITECTURES IN LOCALIZATION SYSTEMS

This chapter summarizes the main contribution of publications [P1], [P2], and [P3] on the impact of different UWB device architectures on the accuracy, energy consumption, and cross-platform compatibility of UWB localization systems. Section 3.1 provides an overview of the main types of UWB architectures, the devices that implement them, and previous comparisons of UWB devices in existing literature. Section 3.2 addresses the differences between the two main types of physical layer architectures in UWB devices in terms of energy consumption and ranging and localization accuracy. Section 3.3 presents challenges in obtaining platform-independent UWB localization systems. Finally, Section 3.4 provides a summary of the main contributions presented in this chapter and a discussion on their limitations.

3.1 State of the Art

The architecture of UWB devices is critical for their performance. The IEEE Standard for Low-Rate Wireless Networks 802.15.4 [47] specifies two types of UWB physical interfaces (PHYs), with low- and high-rate pulse repetition: LRP and HRP, respectively. Some devices not compliant with the standard can implement their own proprietary PHYs. Table 3.1 lists the most popular UWB device models currently on the market, the PHYs they implement, and references of studies that used these devices.

The device architecture and the accompanying software impact the energy efficiency, ranging and localization accuracy, maximum range, and channel estimation. The reception of pulses transmitted at a low rate (so correspond-

Table 3.1 The most popular UWB devices on the market, the physical interfaces they implement, and references for studies or articles mentioning the devices.

Device brand	Device model	PHY	References
Qorvo [*]	DW1000	HRP	[P1], [97, 50, 117, 100, 18]
Qorvo [*]	DW3000	HRP	[P3], [103, 61]
NXP	SR040 / SR150	HRP	[51, 61]
Apple	U1	HRP	[15, 61]
imec	ULP IR-UWB	HRP	[48]
3db Access [†]	3DB6830C	LRP	[P1], [P2], [P3]
Ubisense	Series 7000	LRP	[97], [105]
PulsON	P220, P400	LRP	[54, 71, 136, 68]
BeSpoon	B-UWB-MOD1	Proprietary	[97, 50]
TDSR	P452A	Proprietary	[P3]

^{*} Previously, Decawave.

[†] Also called “3db” in the following.

ing to the LRP PHY) can be performed using a *non-coherent* receiver, i.e., without phase information.

When using HRP, individual pulses have a lower energy than with LRP. Therefore, in the HRP PHY, the reception usually needs to be performed with a *coherent* receiver (so using phase information), in order to extract the weaker pulses from the noise. For optimal reception, the coherent receiver needs to estimate the multipath delays, their complex-valued channel coefficients, and the pulse shape distortion [4]. Also, coherent receivers have strict timing requirements in order to be able to track the phase of the signal. These requirements increase the power consumption of coherent demodulators compared to non-coherent ones [130]. Because non-coherent receivers can estimate the channel coefficients based on the envelope of the signal, this allows a more energy-efficient implementation of the LRP PHY compared to the HRP PHY.

Devices in the IoT can greatly benefit from location awareness, but these devices are frequently energy-constrained. It is therefore important to determine which UWB solutions are the most energy-efficient for the desired

applications. At the same time, the performance of localization systems depends on the ranging or timing accuracy of the devices. Most often, studies focused on the ranging or localization accuracy and overlooked the energy efficiency of the devices [97, 50, 127]. Prior to our work in [P1] and [P2], it was not clear whether there is a trade-off between energy efficiency and localization accuracy in UWB devices.

Although other works have provided a simulation-based comparison between coherent and non-coherent UWB receivers frequently used in HRP and LRP PHYs, respectively, we were the first to also perform power consumption measurements with devices representative for each of the PHYs in [P1] and [P2]. Other works that performed a measurement-based analysis of the ranging accuracy of UWB devices [117, 110, 100] have focused mostly on the Decawave¹ DW1000 IC [22], which implements the HRP PHY. LRP devices have been less studied and only in terms of their ranging accuracy [97, 105]. Only one other paper [97] compared HRP and LRP devices (developed by Decawave and Ubisense, respectively), but only in terms of ranging and localization performance, without analyzing their power and energy consumption. Therefore, a measurement-based comparison between LRP and HRP UWB devices including both the energy consumption and their ranging/localization accuracy has not been performed prior to our work in [P1] and [P2].

In [P3], we continued our work on comparing the impact of different device architectures on localization systems, but this time we focused on the *application* level. In order to accurately timestamp the receive times of incoming messages, UWB devices rely on channel estimation. Many of the most successful ranging error mitigation techniques rely on the resulting channel impulse response (CIR) of the signal to correct NLOS errors [138, 101, 29], since the CIR offers extensive information about the propagation path of the signal. However, prior to our work in [P3], the differences between CIRs acquired at the same physical locations with different devices had not been investigated.

¹Decawave was acquired by Qorvo in 2020 [92], but up until 2021 the devices were still referred to as “Decawave” in official documents, e.g., [21]. Since previous literature also referred to the devices under the “Decawave” name, we also kept this name and, for consistency reasons, will use it throughout the rest of the thesis.

As Table 3.1 shows, there are numerous commercial UWB solutions that must interoperate in localization scenarios. Numerous standardization bodies such as the FiRa™ Consortium [14], the Car Connectivity Consortium (CCC) [13], or Apple Nearby Interaction [49] now establish specifications and certifications that ensure the interoperability of different UWB solutions. It is, therefore, important to evaluate the impact of the various diagnostics used in error mitigation solutions across different UWB platforms, in order to avoid accuracy degradation in localization systems where users carry different models of UWB devices.

Other works are also starting to notice the difficulty in comparing the performance of different UWB solutions. In [103], the authors propose a testbed that allows the benchmarking of different UWB platforms. Whereas the work in [103] focuses on *network* performance indicators such as the packet reception rate (PRR) under different device configurations, we look at *application-level* performance indicators, such as the energy efficiency and ranging accuracy of the devices.

To sum up, our contributions from this chapter investigate the impact of the device architecture from two complementary angles:

1. The energy efficiency, maximum range, and ranging/localization accuracy, where we compared two devices implementing the LRP and HRP PHY. This work corresponds to the publications [P1], [P2].
2. The cross-platform compatibility of localization systems and error mitigation techniques, based on a measurement campaign performed with three devices, implementing an LRP, an HRP, and a proprietary interface. This work corresponds to the publication [P3].

3.2 LRP vs. HRP

In this section, we present our contributions from publications [P1] and [P2] on the comparison between the LRP and HRP PHYs. In Section 3.2.1 we present the architecture of both PHYs with a focus on the characteristics that impact our measurements. In Section 3.2.2 we present the devices used for the comparison [P1] and [P2]. We present only selected results from our comparison, regarding their power and energy consumption (Section 3.2.3),




Figure 3.1 Peak and average power level in LRP and HRP PHYs.

their ranging accuracy (Section 3.2.4), and their localization accuracy (Section 3.2.5). For the detailed analysis of the results we refer the reader to publications [P1] and [P2].

3.2.1 Architecture

UWB transmissions must satisfy two constraints [12]: (i) to have a maximum average power spectral density (PSD) averaged over 1 ms of -41.3 dBm/MHz, and (ii) to have a maximum peak power spectral density of 0 dBm/50MHz. Therefore, UWB devices can either transmit few pulses at high power levels, leading to the LRP PHY (Figure 3.1a) or many pulses at a low transmit power, leading to the HRP PHY (Figure 3.1b). Both technologies can achieve an equal *average* transmitted RF energy if employed optimally.

For the same path loss (or distance), the pulse energy received by the HRP PHY is lower than that of the LRP PHY because the former transmits pulses with lower energy, in order to maintain the same average transmitted energy. Therefore, the HRP PHY often needs more complex receivers to extract the weak pulses from the noise. A coherent receiver typically achieves this by correlating the received signal with a template one over many samples. Because the pulses transmitted by an LRP transceiver are stronger than those of HRP, the LRP receiver can be based on energy detection.

Coherent receivers need to estimate the phase shift induced by the channel in the detection process, which typically demands high computational resources and hardware complexity compared to a non-coherent receiver [4]. Precisely estimating the carrier phase is essential to recover the baseband pulse since any inaccuracy will introduce a signal power loss and crosstalk

interference in phase-shift keying (PSK) modulated signals [90]. A time shift of half the pulse period flips the phase of the signal, so coherent UWB systems usually tolerate rotations of at most $\pi/4$ of the signal phase. For instance, when using a carrier frequency of 8 GHz (UWB channel 11), this corresponds to a maximum time shift of 30 ps. These requirements increase the power consumption of coherent demodulators [130]. On the upside, coherent receivers have low sensitivity to inter-symbol interference and gain from the multipath diversity of the channel [130].

Non-coherent receivers estimate the channel coefficients using the envelope of the signal. Their synchronization constraints depend on the pulse envelope (and thus the pulse bandwidth), so they are not so stringent as those of a coherent receiver. For instance, for a pulse bandwidth of 500 MHz, a non-coherent receiver needs a timing resolution of 1 ns. As such, non-coherent receivers can be implemented in a more energy-efficient manner than coherent receivers. The disadvantages in using non-coherent receivers are that they have a higher bit error probability than coherent receivers [90] and that they cannot measure the angle of arrival (AoA) of incoming signals with closely-spaced antennas.

3.2.2 Devices Used for Measurements

Because the receiver implementation has implications on the resulting energy efficiency and ranging accuracy of the device, it is necessary to compare the two variants of UWB PHYs. Since theoretical comparisons have their limitations, we extend the analysis by comparing the energy efficiency and ranging accuracy of two devices, each implementing one type of UWB PHY: the 3db Access 3DB6830C (Release 2016) IC for the LRP PHY and the Decawave DW1000 (Release 2014) for the HRP PHY.

The comparisons in [P1] and [P2] include complementary functioning modes of the DW1000. In [P1], we used a robust configuration of the DW1000, suitable for long-range distance measurements, which attains a similar range to the 3db IC. In this case, the DW1000 chipset was integrated in the EVK1000 module. The devices used DS-TWR, which has a better accuracy than the SS-TWR but more message exchanges and therefore a higher power consumption. The long-range settings used were the lowest data rate

of 110 kb/s, a low pulse repetition frequency (PRF) of 16 MHz and a large preamble length (PLEN) of 2048 symbols. Low data rates increase the link budget and therefore also the range compared to high data rates, while the chosen preamble length and PRF minimize NLOS effects and make the system robust in the presence of multipath propagation.

In [P2], we used a short-range operating mode for the DW1000 which decreases the robustness compared to the long-range mode but increases the energy efficiency of the chip. Since the energy consumption analysis was an important part of our comparison, this operating mode places the DW1000 chip at an advantage compared to the 3db IC. In this case, the DW1000 chip was integrated in the DWM1001 module, which contains a Nordic Semiconductor nRF52832 BLE microprocessor mostly used for network communication and an STM LIS2DH12TR 3-axis motion detector besides the UWB chip. The DWM1001 module was part of the MDEK1001 kit for localization. The measurements used the default configuration of the kit, corresponding to operating mode 14, namely: a data rate of 6.8 Mb/s, a PRF of 64 MHz, and a preamble length of 128 symbols. In this implementation, the distances are computed using the SS-TWR, which is more energy efficient than the DS-TWR and therefore provides a more fair comparison to the 3db IC. The energy efficiency of the two operating modes of the DW1000 are also given by their packet length. Mode 14 (short-range) has a packet duration of 287 μ s, while a packet in mode 3 (long-range) has a $12\times$ higher packet duration of 3487 μ s. In comparison, the packet duration of a 3db message is approximately 400 μ s.

All devices operate at the center frequency of 6.5 GHz (UWB channel 5). The 3db IC has a total system 10 dB bandwidth of 620 MHz, but because the devices use binary frequency-shift keying (BFSK) modulation, the pulse bandwidth is 380 MHz. The DW1000 IC has a pulse bandwidth of ≈ 662 MHz (and a 3 dB bandwidth of 500 MHz). Because the 3db IC has a lower pulse bandwidth, this can decrease the ranging precision because it decreases the time resolution. However, the frequency diversity of the BFSK modulation partially compensates for this loss.

3.2.3 Power and Energy Consumption

We measured the current consumption of the devices and computed their power consumption according to the input voltage of each device. We measured the current consumption of all devices with a Keysight DC Power Analyzer. We isolated the most important modes, namely the idle, transmit (TX), and receive (RX) and computed their average current consumption. The input voltages of the 3db chip and the DW1000 chip are 1.25 V and 3.3 V, respectively.

As we mentioned, we used different hardware for the HRP PHY in [P1] and [P2]. In [P1], where we used the long-range mode of the DW1000, we could isolate the current consumption of the DW1000 chipset alone, which we were interested in. In [P2], however, we used the Decawave MDEK1001 board, which integrates the DWM1001 *module*. The module contains, besides the UWB chip, a BLE microprocessor and a motion detector. Unfortunately, the board allows for measuring the current consumption only of the entire module, so we could not isolate the current consumption of the UWB chip alone in those measurements. This is why we also present the power consumption of the DW1000 chipset reported in the device datasheet [21] for the two operating modes. We include our measurements for the sake of completeness, but we rely on the values reported in the datasheet of DW1000 when comparing it against the 3db chipset.

Table 3.2 shows, in order, the power consumption of: the 3db IC, the DW1000 IC in mode 3 reported in the datasheet, the DW1000 IC in mode 3 measured by us in [P1], the DW1000 IC in mode 14 reported in the datasheet, and the DWM1001 module in mode 14 measured by us in [P2]. Overall, the average power consumption of the 3db IC is at least 9 times lower than the one of the DW1000 IC in the short-range mode (reported in the datasheet), i.e., the energy-efficient mode. Moreover, the average power consumption of the DW1000 chip in the idle mode is about 1.45 times higher than the one of the 3db IC in the RX mode, which is the most power-consuming state. The results are in line with the theoretical analysis from Section 3.2.1 and show that the LRP PHY can be indeed implemented in a more power-efficient way than the HRP one.

Table 3.2 The average power consumption of 3db (LRP) and DW1000 (HRP) devices. We include the power consumption of the DW1000 chipset in two operation modes: 3 (long-range) and 14 (short-range). For mode 3, we include the power consumption of the DW1000 chipset from the datasheet and the measured power consumption of the isolated chipset when it is integrated in the EVK1000 module. For mode 14, we include the power consumption of the DW1000 chipset from the datasheet and the measured power consumption of the DWM1001 module (which includes, besides the DW1000 chipset, a BLE module and a motion detector).

Device	Mode	Average power consumption [mW]		
		TX	RX	Idle
3db Access [†]	—	20.7	40.7	6.6
DW1000, datasheet [‡]	3	165	267.3	59.4
DW1000 (EVK1000 module) [‡]	3	194.5	492.4	68.8
DW1000, datasheet [‡]	14	237.6	392.7	59.4
DWM1001 module [‡]	14	297.7	507.2	47.9

[†] Referenced to 1.25 V.

[‡] Referenced to 3.3 V.

The power consumption analysis is the basis for evaluating the energy consumption of an UWB localization system. A key challenge is minimizing the energy consumption of the *tag*, which is usually battery-powered. We consider that the tag initiates the message exchange and stays in the idle mode between consecutive rangings. We estimate the energy consumption only when the device is in the TX or RX mode. The packet length of the DW1000 chip in Mode 14 is 287 μ s and the one of the 3db chip is 400 μ s. Therefore, a DW1000 tag will consume 180 μ J per SS-TWR during transmission and reception, while a 3db tag will consume 28 μ J, so $6.4\times$ less energy than the DW1000 tag. When configured in the long-range mode, the packet duration of Decawave devices is 3487 μ s, so about $10\times$ larger² than that of 3db devices. This means that, in the long-range mode, the DW1000 chipset consumes at least $100\times$ more energy than the 3db chipset.

²The notation “ $n\times$ ” denotes an increase of the measured quantity by n times.

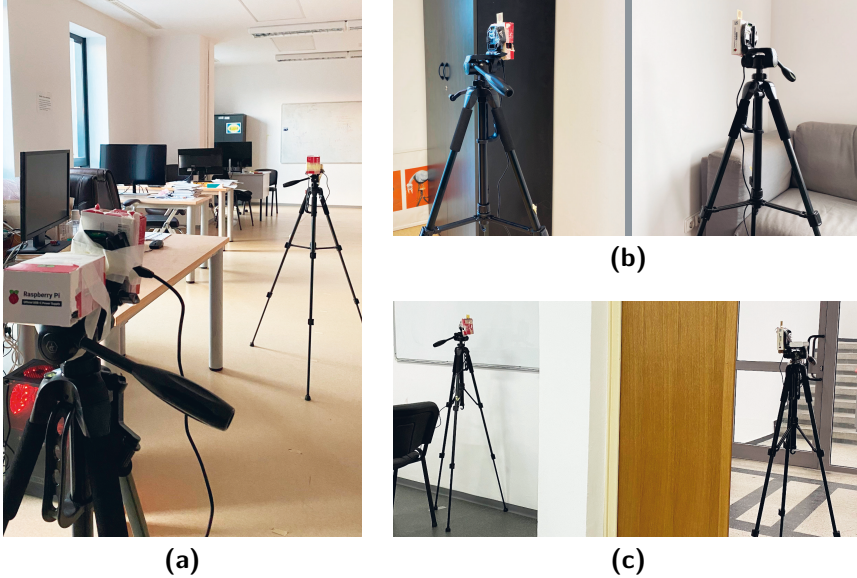


Figure 3.2 Setup of ranging measurements in (a) LOS, (b) NLOS with a drywall, and (c) NLOS with a concrete wall. The UWB devices are placed on tripods. The NLOS with human body shadowing setup is identical with the LOS one, except that a person is standing right in front of the transmitter (the device further away).

3.2.4 Ranging Accuracy

In this section, we compare the ranging accuracy of the 3db IC and the DW1000 IC in the energy-efficient configuration (mode 14) from [P2]. We performed measurements in which we placed pairs of devices from each brand *at the same locations*, in several LOS and NLOS setups illustrated in Figure 3.2. We considered four indoor settings: LOS inside a large office and NLOS caused by either a gypsum wall (12.5 cm thickness), a concrete wall (29 cm thickness), or a human body. We refer the reader to [P2] for more details on the experimental setup.

We define the ranging error as:

$$e_d = \hat{d} - d, \quad (3.1)$$

where \hat{d} is the measured and d is the true distance. We compare, on the one hand, the distribution of the aggregated errors for each type of LOS and

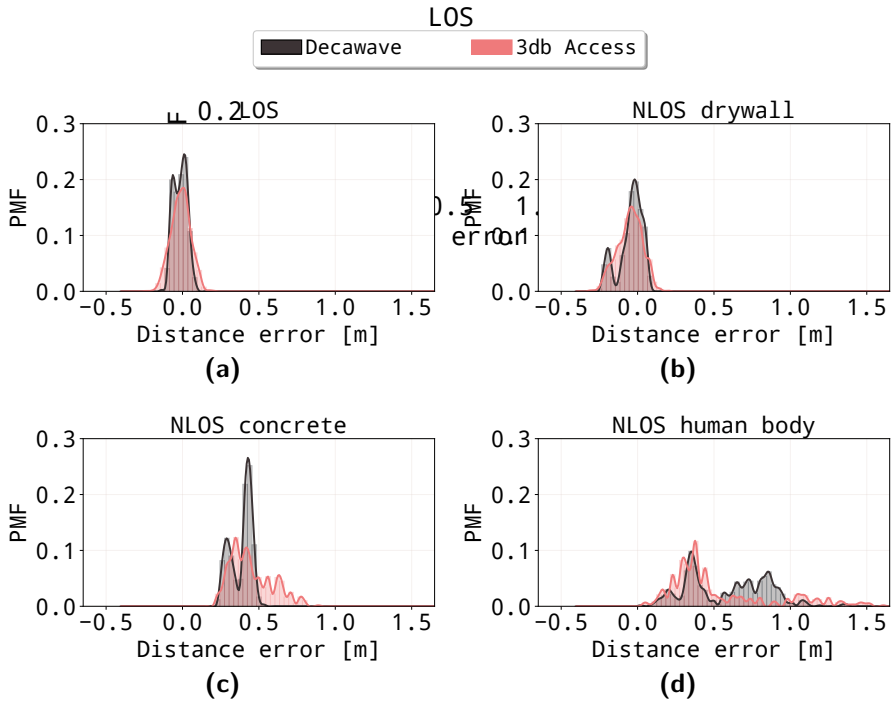


Figure 3.3 Comparison of the distributions of aggregated ranging errors of 3db Access and Decawave devices in LOS and NLOS settings where the obstruction is caused by drywall, concrete wall, or a human body.

NLOS scenario in Figure 3.3, and on the other hand the error distributions at each test point in each scenario in Figure 3.4³. Table 3.3 also presents the mean, standard deviation, and inter-quartile range (IQR) of the distance errors of 3db and DW1000 devices.

In all scenarios, both devices had mean errors within 2–5 cm of each other, with Decawave devices performing better in all scenarios except for the NLOS with human body shadowing one. In LOS and soft NLOS scenarios, the devices had a difference in the standard deviation and IQR of 1–2 cm, with Decawave devices yielding a higher accuracy in most cases. In the hard NLOS scenarios (with concrete wall and human body shadowing), Decawave devices had a lower spread than 3db devices by 1.23–2 \times , except for the IQR in the case with NLOS with human body, which was 1.59 \times higher than the one of 3db devices. At individual test points, Decawave devices had 2 \times lower spread than 3db devices. Compared to [P1], the performance of 3db devices

³We illustrate the distributions using Tukey’s boxplots.

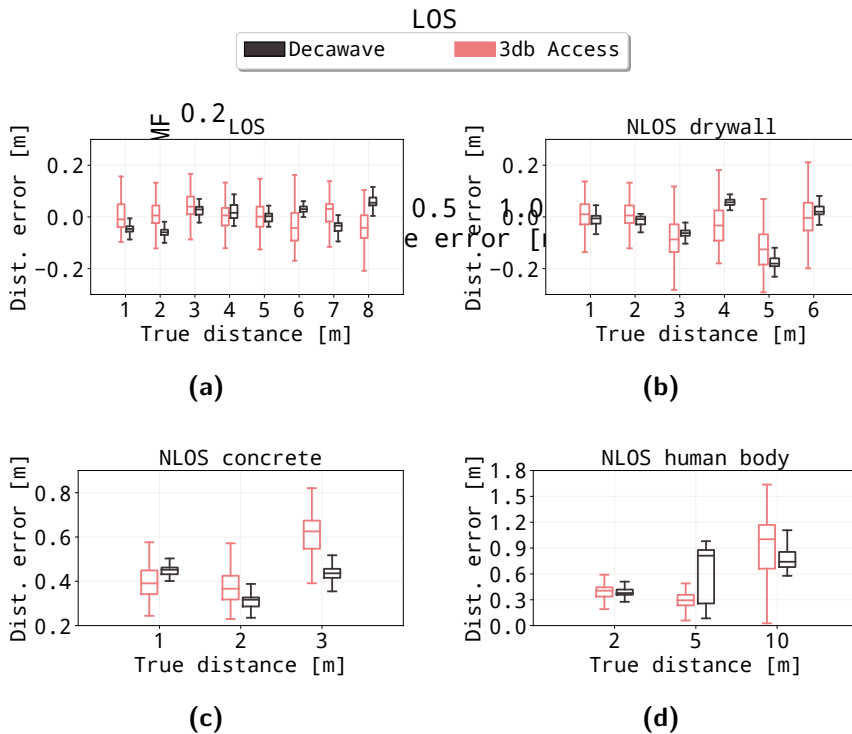


Figure 3.4 Comparison of the distributions of ranging errors at selected locations (for multiple true distances between the transmitter and the receiver) of 3db Access and Decawave devices in LOS and NLOS settings where the obstruction is caused by drywall, concrete wall, or a human body.

in [P2] has improved thanks to the refined calibration and to the correction of the firmware issues that previously caused large outliers in certain NLOS situations.

3.2.5 Localization Accuracy

In this section, we compare the localization performance of the DW1000 (short-range mode) and 3db Access devices using measurements acquired in the same localization area, shown in Figure 3.5. We performed measurements with each type of UWB IC in identical setups consisting of four anchors and one tag belonging to the same brand. In each case, we placed four anchors over an area of approximately 4.5×3.6 m at heights between 1.2–1.8 m. For localization, we used the regularized Gauss-Newton multilateration algorithm presented in [113]. More details about the setup can be found in [P2].

Table 3.3 Statistics of the distance measurement errors: mean, standard deviation (σ), and IQR.

Scenario	Device	Mean [m]	σ [m]	IQR [m]
LOS	3db Access	0.02	0.07	0.09
	DW1000	0.00	0.05	0.07
NLOS with drywall	3db Access	-0.04	0.08	0.12
	DW1000	-0.01	0.09	0.10
NLOS with concrete wall	3db Access	0.46	0.14	0.19
	DW1000	0.44	0.07	0.14
NLOS with human body	3db Access	0.55	0.32	0.29
	DW1000	0.60	0.26	0.46




Figure 3.5 Localization setup: the four anchors (A1 to A4) encompass an area of approximately 4.5×3.6 m and the tracking area is on the table. The ground truth was acquired using two HTC Vive base stations (BS1 and BS2) and a tracker that was colocated with the UWB tag, shown on the table.

To acquire the ground truth locations, we developed a system based on the HTC Vive motion capture system [28]. The work was performed and published during the Doctoral studies of the author, but not included among the main research publications of the thesis, as it is adjacent to the main topics presented in the dissertation. The system has an average accuracy of 5 mm, so it is suitable for acquiring ground-truth locations for centimeter-level localization technologies, such as those based on UWB communication. Its main advantage is its low cost, which is about $10 \times$ lower than the one of professional-grade motion capture systems, while still retaining a millimeter-level accuracy.

Table 3.4 Statistics of measurement-based localization errors.

Case	Device	Localization error			
		Mean [cm]		Standard deviation [cm]	
		2D	3D	2D	3D
LOS	3db Access	9.7	36.8	5.4	14.8
	Decawave	12.4	42.7	7.8	19.1
NLOS	3db Access	18.9	72.2	10.5	34.1
	Decawave	22.9	89.1	23.5	43.7

We performed two types of recordings: one in which the tag was in LOS with the anchors and one in which one of the anchors (hereby referred to as the NLOS anchor) was at all times in NLOS with the tag, obstructed by a human body.

We define the localization error as the Euclidean distance between the true and the estimated location, either in 2D or 3D. For the 3D case, the error is:

$$e = \sqrt{(x - \hat{x})^2 + (y - \hat{y})^2 + (z - \hat{z})^2}, \quad (3.2)$$

where (x, y, z) and $(\hat{x}, \hat{y}, \hat{z})$ are the Cartesian coordinates of the true and, respectively, estimated locations. For the 2D localization errors reported, the z -axis coordinate is not taken into account.

Table 3.4 presents the mean and standard deviation of the localization error of the 3db and DW1000 ICs in 2D and 3D. In LOS, the two devices have the mean and standard deviation of localization errors within at most 5.9 cm of each other, leading to a similar performance. In the NLOS scenario, the mean and standard deviation of the localization errors of the Decawave-based localization system are higher with 4–16.9 cm than 3db’s. In 3D, the average localization errors are significantly higher than in 2D, with about 30 cm in LOS and 60 cm in NLOS. The increase is due to the measurement noise and high GDOP on the z axis. In 2D localization, 90% of the LOS errors are under 20 cm, while in NLOS 75% of the localization errors are under 20 cm.

3.3 Challenges in Platform-Independent UWB Localization Systems

In this section, we present the work from publication [P3], where the goal was to identify the main challenges in developing platform-independent UWB localization systems, especially those that apply data-driven error mitigation techniques. For this purpose, we performed a measurement campaign with devices from three companies: Decawave (now Qorvo), TDSR, and 3db Access. We acquired measurements with pairs of devices from each brand at exactly the same locations, in order to assess the differences in their distance errors and CIRs under the same propagation conditions (i.e., the same physical location, room layout, furniture, room occupancy and, if the case, obstructions). We identified four main challenges pertaining to these differences that can hinder the cross-platform compatibility of error mitigation techniques frequently employed in UWB localization systems.

The accuracy of UWB localization systems can be improved in several ways. First, in case of multilateration localization algorithms, the errors of individual ranges (between each anchor and the tag) can be corrected before they are passed to the localization systems. Because high ranging errors frequently appear in NLOS propagation, a line of work has focused on detecting the NLOS condition and then either discarding the measurement [102], assigning it a lower weight in the localization algorithm [37], or correcting the distance measurement [69]. Another approach is to estimate the TOA using neural networks (NNs) trained for this task [27, 134]. NNs show better results than the conventional parameter-based TOA estimation which cannot easily adapt to changing propagation conditions [27]. Other works bypass the NLOS detection step and directly predict the distance error of the measurement, also using machine learning (ML) models [101, 29].

In recent years, ML-driven approaches for ranging and localization have taken precedence over statistics-based techniques [69]. Because the CIR offers rich information about the propagation path of the signal, most error-mitigation techniques rely on the CIR or on features extracted from the CIR to detect the NLOS condition. However, all the cited solutions have been trained on data from only one device brand. Given the breadth of commer-

cial UWB solutions on the market presented in Table 3.1, we cannot assume that all users accessing a localization system in a public space (e.g., mall, airport, university) will be equipped with the UWB hardware that the ML model was trained on.

Our measurement campaign was novel from several points of view. First, although measurements with different UWB devices at the same locations have been previously done in [P1], [P2], [97, 50, 68], these studies were focused on comparing the *performance* of the devices. In this work, we were interested not only in the *average* errors of each device, but also in the errors at selected locations, which gives insight into their different operating modes. Second, we were also the first to also compare the CIRs of different devices and to highlight how these differences can prevent the cross-platform compatibility of ranging and localization systems. Third, we provide one of the few open-source datasets [55] of this kind, which can be used to further study platform-independent error mitigation techniques. Fourth, these devices have not been previously compared. The 3db chipset was compared against the Qorvo DW1000 chipset in our previous works [P1], [P2], whereas here we use the new-generation DW3000 for the comparison. Also, to the best of our knowledge, the TDSR P452A has not been used in any other research works to this date. More details about the measurement setup and the configuration of the devices can be found in [P3].

In this section, we will discuss the main identified challenges hindering platform-independent UWB localization systems:

Section 3.3.1: The same environmental settings lead to different distance errors for different devices.

Section 3.3.2: The same (NLOS) environmental conditions can lead to a different TOA estimation in consecutive measurements even for the same device. Therefore, NLOS errors depend not only on the obstacle and environment, but also on the hardware and its TOA estimation.

Section 3.3.3: CIRs acquired by different platforms at the same locations have different lengths, shapes, and statistics.

Section 3.3.4: Different vendors provide different types of diagnostics (e.g., signal and noise power), sometimes using incompatible units of



(a) LOS



(b) NLOS with half wall



(c) NLOS with bar refrigerator



(d) NLOS with pillar

Figure 3.6 Examples of environments from the measurement campaign.

measurement.

3.3.1 Same Conditions, Different Distance Errors for Different Devices

The first factor that prevents the extension of error mitigation methods from one hardware to another is the variability in their distance errors even under the same propagation conditions.

We performed ranging measurements between pairs of devices belonging to each brand in several locations at Tampere University. We targeted LOS and NLOS scenarios, for the latter using as obstructions walls, pillars, furniture, a TV screen, and room divider panels. We performed measurements in nine different spaces which included meeting rooms, offices, corridors, or cafeterias. Figure 3.6 shows some of the locations in which we acquired measurements.

Figure 3.7 shows the cumulative distribution function (CDF) of distance errors obtained with the three devices in LOS and NLOS scenarios. We com-

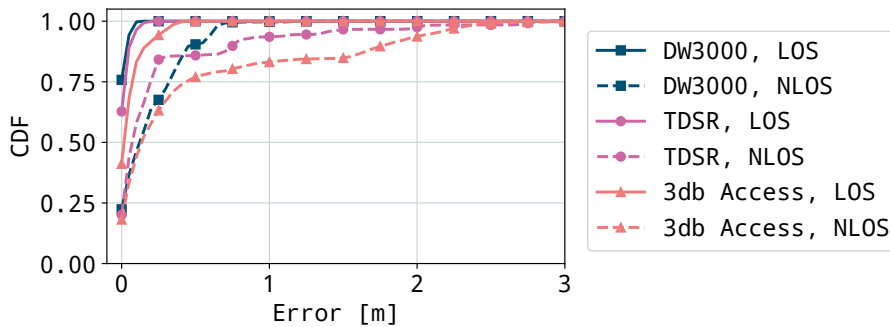


Figure 3.7 CDF of errors in LOS and NLOS for all devices.

pute ranging errors as in the previous chapter, using Eq. (3.1).

While in LOS most devices have a similar performance, as can be seen from the CDF of aggregated errors in LOS and NLOS for all devices from Figure 3.7, in NLOS the analysis is more complex.

We observe from Figure 3.7 that even NLOS errors can be very small. In fact, at least 50% of NLOS measurements are under 15 cm, which is usually regarded as an acceptable UWB ranging error. This means that, even if there were a NLOS detection model with 100% accuracy, discarding *all* NLOS measurements would imply discarding also measurements with acceptable errors. This can raise issues if the number of anchors in the area is already small.

Until now, “NLOS” was defined as the absence of a *visual* direct propagation path between two devices. However, in localization we are frequently interested only in the localization *errors*. It might be that NLOS does not, in fact, cause ranging or positioning errors because the device was able to lock on the correct TOA. Therefore, instead of detecting the LOS/NLOS situation, we could try to detect only the NLOS measurements that introduce significant ranging errors. One difficulty is setting a threshold beyond which an error is considered “significant” and should be discarded. This can be selected based on the requirements of a particular application.

However, this approach of detecting only NLOS scenarios which introduce high ranging error poses yet another issue. Figure 3.8 shows the distribution of distance errors obtained by the three devices with various obstacles. The distributions are illustrated as boxplots using Tukey’s definition. We notice that there is a high variability in the distance errors achieved by




Figure 3.8 Distribution of errors from selected locations with different obstacles.

different platforms, even under the same conditions. In some cases, we note even a high spread in the errors obtained by a *single* device at one location. Therefore, according to the definition of the NLOS condition as the situation which gives rise to “significant” localization errors, the same setup could be considered NLOS for some devices but LOS for others.

These examples show that a binary LOS/NLOS classification might not always be the best strategy for distance error mitigation. This is because, by adopting this strategy, we risk discarding useful measurements. Moreover, there is a high variability in the NLOS errors of different platforms. This compels us to look more closely into the inner working of the distance estimation algorithm and identify the cause of this variability.

3.3.2 Same Conditions, Different TOAs

We remind the reader that the distance between two devices is, in our case, computed based on the time of flight between the TX and the RX, obtained during a two-way message exchange. Therefore, any error in the timestamping of the TOA will also incur a distance error.

Even though the TOA estimation algorithms used by the devices are closed-source, it is fair to assume that most of them follow model-based approaches that require pre-set parameters such as the power ratio between two consecutive multipath components (MPCs), the peak-to-average power ratio, or the noise threshold beyond which to ignore any peaks [132]. We found that the most common reason for differences in ranging accuracy between differ-

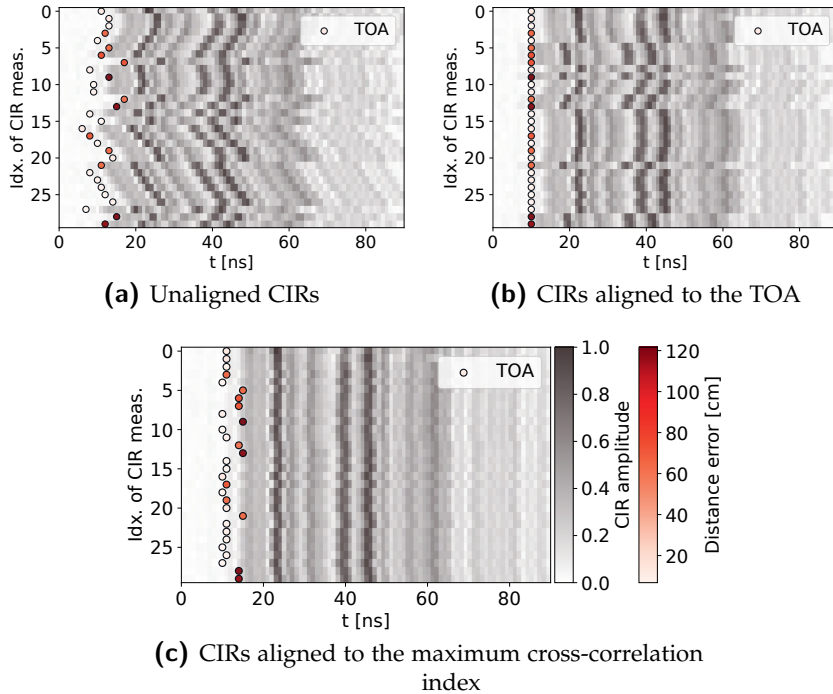


Figure 3.9 The plots show CIRs that are stacked vertically, with their amplitude encoded in the color depth. The measurements were acquired in a NLOS scenario. The estimated TOA is represented through circle markers and the color of the circle encodes the distance error of the measurements. Figure 3.9a shows the unaligned CIRs (i.e., as saved in the buffer by the device). Figure 3.9b shows the CIRs aligned to the TOA. Figure 3.9c shows the CIRs aligned to the maximum value of the cross-correlation function computed between the first CIR and all of the following [P3].

ent devices, on the one hand, and between different measurements acquired with the same device, on the other hand, is the TOA estimation.

To illustrate this, let us consider a series of CIRs acquired with one pair of devices from the same brand (DW3000) in a NLOS scenario. The devices are static during this experiment. Figure 3.9 shows a series of 30 CIRs stacked vertically with different alignments. The MPCs from each CIR have color-encoded amplitudes and the TOA is marked with a circle. The color of the circle encodes the distance error of the measurement corresponding to that CIR.

Figure 3.9a shows the “raw” CIR alignment, i.e., as the CIR was saved

in the internal buffer of the device. The documentation does not specify a preferred position of the TOA within the buffer. Under this representation, we cannot observe any pattern in CIRs that correspond to large errors.

In most works, the CIRs are cropped around the estimated TOA, as in Figure 3.9b. This provides a convenient representation for LOS/NLOS classification or for distance-error prediction, since CIRs are aligned to a common denominator and deviations from it (e.g., the maximum peak occurring earlier or later in the buffer, shorter or longer tails) can indicate different error magnitudes. Using this representation, we could believe that different CIR shapes are correlated with different magnitudes of ranging errors.

Instead, if we align the CIRs according to their maximum cross-correlation index, i.e., according to their similarity, we obtain the alignment from Figure 3.9c. Using this arrangement we see that, in fact, consecutive CIRs acquired in the same scenario are very stable. However, it is the estimated TOA that varies. The measurements with the lowest distance errors have the earliest TOA. However, because these early paths have amplitudes close to the noise threshold, they are ignored in some measurements and the TOA is identified as corresponding to a later reflection.

Therefore, the devices can experience different errors even at the same location because the SNR of the first path fluctuates or because of the TOA estimation algorithm implemented by the device.

3.3.3 Same Conditions, Different CIRs

The received signal can be modeled as [60]:

$$r(t) = s(t) * h(t) + n(t), \quad (3.3)$$

where $s(t)$ is the transmitted pulse which is convolved with the CIR $h(t)$. The signal $n(t)$ represents sensor or environmental noise (so not related to the propagation path), usually modeled as zero-mean white Gaussian noise.

The CIR can be decomposed into K multipath components (MPCs) with delays τ_k , $k = 1, \dots, K$ and amplitudes α_k and a stochastic process $\nu(t)$ which

results in diffuse multipath caused by scattering and diffraction:

$$h(t) = \sum_{k=1}^K \alpha_k \delta(t - \tau_k) + \nu(t). \quad (3.4)$$

The estimated CIR, denoted by $\hat{h}(t)$, is obtained by decorrelating the received signal $r(t)$ with the known reference pulse $s(t)$.

Because the measurements were acquired at exactly the same locations, we could expect to see the main MPCs at the same delays τ_k in the CIRs of all devices. However, we most note that the devices have different pulse bandwidths and/or center frequencies. Therefore, there will be different constructive/destructive interference patterns, usually reflected in $\nu(t)$. This means that the delays of the main MPCs τ_k can also be different, depending on how the reflections add up. The amplitudes of the MPCs α_k are affected by the different front-end circuits. The various signal processing components (e.g., low-noise amplifiers, mixers, automatic gain controls, analog filters, analog-to-digital converters) and digital processing steps can introduce different linear and/or non-linear distortions depending on the hardware. As such, we will, in fact, see different patterns in the CIRs acquired by different devices even if the channel and environmental conditions were kept the same.

To characterize CIRs acquired at the same location with different devices, we compute the average number of main peaks (or MPCs) and the average delay between the first and the last significant peaks. These measures indicate whether we can identify the same main MPCs in all CIRs and how long it takes until their amplitudes decay to a negligible level. We also compute the energy, mean excess delay (MED), and root-mean square (RMS) delay spread of the CIRs. These statistics have been previously used to discriminate LOS from NLOS measurements based on their CIRs [69].

We compute the number of significant MPCs in each CIR. We consider a peak significant if it exceeds 25% of the maximum amplitude of a CIR. We consider a minimum separation of 2 ns between peaks, to avoid detecting peaks belonging to the same path. The average number of significant peaks is denoted by N_p . We also compute the average delay between the first and the last significant peak, which reflects the decay time of MPCs in a CIR.

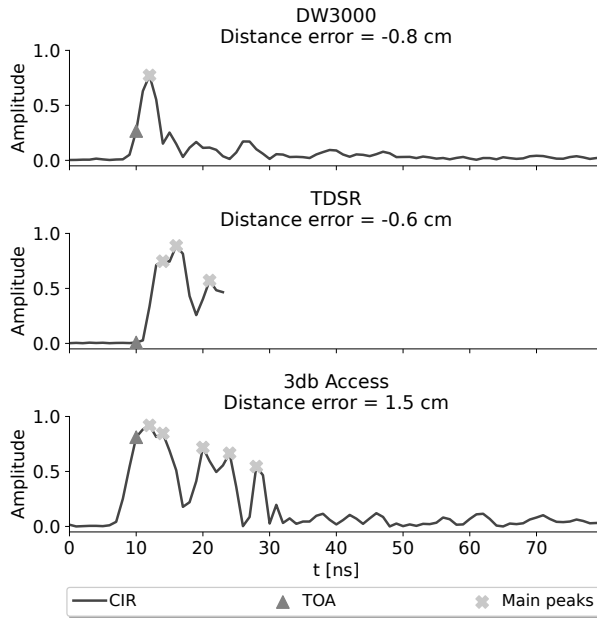


Figure 3.10 CIRs and their main peaks acquired with each device at the same location in a LOS scenario [P3].

Figure 3.10 shows the most significant peaks in a triplet of CIRs acquired at the same location with the three types of devices. Table 3.5 shows the CIR statistics averaged over all locations. 3db CIRs have more significant MPCs than the other two devices and their amplitudes take longer to decay, reflected in a higher delay μ_δ and energy. TDSR CIRs are the shortest, so they capture the fewest significant peaks and have the lowest energy, MED, and RMS delay spread.

In the literature, a low signal energy and high delay statistics have been associated with signal attenuation due to NLOS propagation [69]. However, if CIRs acquired in the *same* conditions with different brands of devices have, on average, different statistics, these can introduce issues in error mitigation methods customized for one device but applied on different brands of devices.

Table 3.5 Statistics of the CIRs: number of significant paths (N_p), delay between the first and the last significant path (μ_δ), energy (E , quantized), mean excess delay (τ_{MED}), and RMS delay spread (τ_{RMS}).

Device	N_p	μ_δ [ns]	E [-]	τ_{MED} [ns]	τ_{RMS} [ns]
DW3000	3	11.2	2.4	13.5	21.7
TDSR	2	5.3	1.6	8.2	3.0
3db	5	24.4	8.5	13.3	13.8

3.3.4 Different Devices, Different Diagnostics

The length of CIRs was between 256–1632 for the tested devices. For some applications, it might be too computationally expensive to process CIRs. Therefore, in some works, diagnostics such as power or noise figures are instead used for ranging or location error mitigation [101]. However, as we will show in this section, this approach can also pose issues for platform-independent localization systems. This is because not all devices provide the same diagnostics and often not in the same measurement units.

For instance, the DW3000 provides the power, maximum amplitude, and phase of arrival (POA) computed on the preamble for regular ranging and, additionally, the power and POA computed on the scrambled timestamp sequence for the secure ranging mode. The manufacturer (Decawave/Qorvo) provides formulas for converting the power to dBm [66].

TDSR devices provide the maximum value in the leading edge (LE) window of the received CIR and the noise level. However, the documentation of TDSR devices does not specify the measurement unit of these parameters [64]. Therefore, it is not clear if they can be compared, for instance, with the ones of the DW3000 chipset. TDSR devices also provide the noise amplitude and a coarse and a precise estimate of the range, of the distance error, and of the tag’s velocity.

The 3db chip provides the peak and the LE amplitudes (but not expressed in dBm) and other additional diagnostics [1].

Because the diagnostics provided by different platforms vary and because they are not expressed in the same unit of measurement, it is difficult to use this information in a platform-independent UWB localization system. There-

fore, UWB manufacturers should join forces to provide unified metrics that could be used to mitigate measurement errors even across different devices. Such a feature would also be beneficial in comparing the performance of different devices.

3.4 Summary

In this chapter, we presented the results on our study of UWB device architectures and their impact on the ranging and localization accuracy, power and energy comparison, and cross-platform compatibility of error-mitigation techniques.

We first compared the power energy consumption, distance measurement statistics, and localization performance of 3db Access and DW1000 ICs, representative of the two types of UWB physical interfaces, LRP and HRP, respectively. Publications [P1] and [P2] provided complementary comparisons: while [P1] was mostly focused on the power and energy consumption of the 3db IC and the DW1000 IC in the *long-range* mode, in [P2] we preferred a more energy-efficient operation mode for the DW1000, to provide a more fair comparison to the 3db chipset.

Both the 3db and the DW1000 IC have ranging and localization errors on the same order of magnitude: cm-level in LOS scenarios and dm-level in hard NLOS scenarios. The DW1000 chipset generally shows better performance in LOS and through-the-wall NLOS conditions, while 3db devices have slightly better performance in NLOS with human body shadowing. For a similar maximum range, Decawave devices have $125\times$ higher energy consumption than 3db Access devices, while in the short-range mode (which decreases the range by at least $8\times$) they have $6.4\times$ higher energy consumption than 3db Access. Therefore, devices implementing the LRP PHY might be more suitable for ultra-low power applications, while the HRP PHY might be a better choice for the highest ranging accuracy.

Second, we presented the main challenges in obtaining ranging or localization error-mitigation techniques that are compatible across different UWB devices. We derived our observations from a database of measurements acquired with three brands of UWB devices at exactly the same locations.

Applying error mitigation customized for one device to other hardware is hindered by the different errors yielded by the devices even under the same conditions, which is due to the different TOA estimation algorithm and hardware. The CIRs of different devices acquired under the same propagation conditions have different lengths and there is not a clear correspondence between the main MPCs identified in each of them. Since most error-mitigation methods in recent literature rely on the CIR [101, 29, 69], this could degrade the accuracy if the algorithms were deployed in locations in which users can have different UWB devices. We therefore brought to researchers' attention the issue of cross-platform compatibility in UWB localization systems and called for more research in this area.

Possible solutions to this issue include developing error mitigation solutions that can work across multiple platforms. The incorrect TOA estimation is the root problem for the different ranging errors of the devices we tested, even when the measurements were acquired under the same conditions. Future research should investigate whether it is possible to find a set of device-agnostic CIR features based on the estimated versus the correct TOA, which could help identify the correct distance errors even across multiple devices. Still, even if such features were found for some devices, they are unlikely to perform as well if applied on an unknown hardware that was not taken into account in the feature extraction process.

A different approach to the centralized solution suggested previously would be to implement error-correction methods at the edge. In this way, we could use models customized for the target device. However, in this case, more research is needed to develop models that do not depend on the environment in which the data used for training was acquired [7]. Measurement campaigns with only one device are already time-consuming; performing such measurement for every new location with every commercial device is simply unreasonable. A solution that estimates the TOA based on synthetic data was proposed in [27] and showed promising results. However, in order for the method to be used on-chip to correct clock drift errors, it must also be fast and lightweight. These aspects were not evaluated in [27], but can be a promising research topic for the future.

4 NLOS AND MULTIPATH PROPAGATION IN WIDEBAND AND NARROWBAND COMMUNICATIONS

NLOS propagation is arguably the most severe issue that affects the accuracy of localization systems. In this chapter, we will focus on the effect of NLOS and multipath propagation on two types of localization or ranging systems: based on UWB technology, which corresponds to publication [P4], and based on BLE communication, which corresponds to publication [P5].

Although the main topic of the dissertation has revolved around UWB localization systems, at the beginning of 2020, the world was shaken by the COVID-19 pandemic. This event (temporarily) shifted the research topic of many researchers, including myself, in order to address the pressing concerns at that time. What started as a project meant to fuse BLE and UWB measurements in order to increase the accuracy of contact tracing applications, soon turned into the study from [P5], which is focused solely on the challenges faced by RSS-based proximity detection applications using BLE. Nevertheless, we believe that the work still fits in the rest of the topics covered in the dissertation, for two main reasons.

First, the study on BLE complements our prior work by providing an overview of the challenges faced in localization systems based on *narrow-band* communications, as opposed to ultra-wideband. Second, many devices nowadays integrate both BLE and UWB chipsets (e.g., car keyfobs, the Apple AirTags). Although BLE modules are mostly used for ultra-low power communication, they can also be used for coarse energy-efficient ranging, followed by a refinement of the estimated location or distance using UWB devices [8]. Since BLE modules are already widespread in wearable devices, their integration with existing infrastructure is more seamless than in the case

of UWB. Therefore, we believe that the work presented in [P5] represents a useful benchmark for comparing and fusing BLE and UWB measurements in indoor localization applications.

The rest of the chapter is organized as follows. In Section 4.1, we present the state-of-the-art in error mitigation techniques for UWB and BLE communications. In Section 4.2, we present an error-mitigation method for UWB localization systems, corresponding to publication [P4]. The novelty of the proposed system is given by the fact that it does not need any manually-collected training data. In Section 4.3, we present the main sources of signal strength fluctuations in BLE proximity-detection algorithms and how to address them, corresponding to publication [P5]. Finally, in 4.4, we summarize our contributions related to the characteristics of NLOS errors in ranging/localization systems based on UWB and BLE communications.

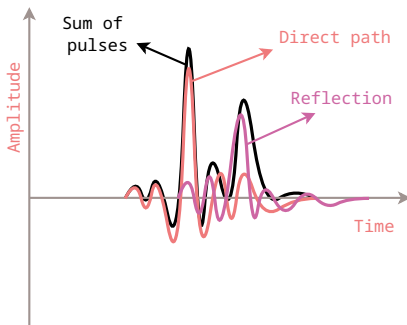
4.1 State of the Art

In this section, we discuss how NLOS and multipath propagation affects narrowband and wideband systems, exemplified by BLE and UWB communication protocols, respectively. Multipath propagation is the phenomenon in which the transmitted signal arrives at the receiver through multiple paths, due to the signal reflecting on surrounding objects, as shown in Figure 4.1a. Obstructions between the TX and the RX can attenuate and delay the RF signal traveling through the object. This is known as NLOS propagation.

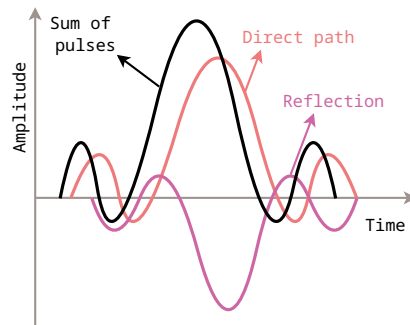
In UWB technology, because of its very large bandwidth (at least 500 MHz), pulses arriving through multiple paths can be resolved at the receiver (Figure 4.1b). Because UWB systems generally employ time-based methods for localization, i.e., they can precisely timestamp the moment at which the signal arrives at the receiver, they can therefore still recover the TOA of the first incoming path. Therefore, UWB signals are usually resilient to multipath propagation. In fact, UWB signals can *benefit* from the multipath diversity inherent in the channel [94]. However, when there is an obstruction between the TX and the RX, the signal traveling through the direct path might be attenuated close to the noise level. In this situation, a later path with higher amplitude might be identified by the TOA estimation algorithm as the first



(a)



(b) Ultra-wideband



(c) Narrowband

Figure 4.1 Illustrative example of the effect of multipath propagation on narrowband and ultra-wideband communication. In Figure 4.1a, the signal arrives from the transmitter (console near the TV) to the receiver (the hand-held device) through two main paths: the direct one and through the indirect path reflected on the wall. While in UWB communication the contribution of the signal arriving through the direct and reflected paths can be distinguished in the sum of pulses (Figure 4.1b), in narrowband communication the individual paths cannot be resolved (Figure 4.1c).

path, as we showed in Section 3.3.2.

In narrowband communication systems such as BLE, because of the poor time resolution, the signal arriving through multiple paths will not be resolvable at the receiver (Figure 4.1c). Such systems generally rely on the received

Table 4.1 Comparison of UWB and BLE accuracy in ranging and localization applications. We list the lowest value (or range) of the 95th percentile (denoted by P_{95}) of ranging/localization errors from the cited references.

Technology	Application	Method	References	P_{95}
UWB	Ranging	TWR	[P1], [P2], [18, 50, 97, 100, 117]	10 cm
	Localization	Multilateration	[P2], [P4], [50, 72, 78, 97, 127]	10 cm
		TDOA	[P6], [6, 35, 38, 58, 88, 118, 124, 137]	30 cm
BLE	Localization	Multilateration	[5, 77, 86, 87, 42]	2–5 m
		Fingerprinting	[3, 26, 82, 89]	1.5–3 m

signal strength (RSS) to estimate the distance between the TX and the RX or the location of an user. In proximity detection, multiple empirically-derived path-loss models have been proposed in the literature to estimate the distance between the target and a reference device based on the observed RSS [107, 65, 139, 41, 104]. For localization purposes, the most popular approaches are based on multilateration using the estimated ranges [5, 77, 86, 87, 42] or based on RSS fingerprinting [3, 26, 82, 89].

Because the narrowband signal that arrives through multiple paths can add up constructively or destructively at the receiver, multipath propagation causes fluctuations in the RSS. This, in turn, can cause fluctuations in the estimated distance or location solely because of the environment, even if the user is not moving. RSS fluctuations can therefore cause distance or location errors [42]. Similarly, obstacles between the TX and the RX can attenuate the signal and lead to ranging errors.

Table 4.1 provides a comparison of the ranging and localization accuracy obtained by UWB and BLE devices using various techniques such as multilateration, TDOA, and fingerprinting. Undoubtedly, UWB technology provides the highest accuracy on the order of centimeter, whereas BLE achieves at best only meter-level accuracy.

However, the ranging/localization accuracy is only one side of the picture. Because the devices tracked in positioning applications are generally battery-powered, their energy efficiency is also crucial. As we showed in [P2], the average power consumption of a UWB transceiver is between 20.7–40.7 mW for LRP devices and 165–267.3 mW for HRP devices in the short-range mode. Also, as we showed in Table 3.1, HRP devices are the most widespread on the market, especially in smartphones. Qorvo UWB chipsets are available in Google Pixel smartphones [93], Apple develops its own UWB U1 chipset [74], and NXP UWB devices are found in Samsung smartphones [81]. All the aforementioned chipsets implement HRP PHYs. We estimated that the energy consumption of the Qorvo DW1000 chipset for performing one SS-TWR measurement in the short-range mode is 180 μ J.

In comparison, the power consumption of a Nordic Semiconductor nRF52 BLE chipset is 22.7 mW [80] and its energy consumption for transmitting an advertising beacon is between 40–58 μ J (depending on the payload size) [115]. Therefore, BLE consumes around 3–4.5 \times less energy than a DW1000 chipset for obtaining a ranging measurement. The RX-to-TX power consumption ratio is also higher in UWB than in BLE, which makes UWB devices less suitable for discovery than BLE devices [8]. Therefore, there is a case for using BLE for localization applications for energy-constrained devices. Alternatively, a hybrid BLE–UWB approach has been proposed for tracking social interactions in [8]. Regardless of the approach, a comparison of the ranging capabilities of BLE and UWB is necessary to find the best strategy for the desired localization application.

In this section, we present the literature on UWB and BLE localization methods and NLOS mitigation techniques. We first focus on NLOS mitigation techniques in UWB localization and how they compare to our proposed methods from [P4] (Section 4.1.1). Then, we review the SOTA in studies related to the variability of BLE RSS measurements related to our work in [P5] (Section 4.1.2).

4.1.1 SOTA on UWB Error Mitigation Techniques

In UWB localization, there are some works that apply semi-definite programming to mitigate NLOS errors which do not need any measurement statis-

tics [122, 126]. However, such methods need more computational power than regular localization algorithms [122], which might not be available on low-power sensors. Other works propose the use of additional hardware, e.g., inertial measurement units [62]. In [39], only the measurement variance needs to be known to select one of two extended Kalman filters (EKF), one for each LOS/NLOS situation.

NLOS identification techniques that rely on the CIR have been proposed in [69, 134, 102] with over 90 % accuracy in correctly identifying LOS and NLOS errors. However, all data-aided NLOS mitigation methods need measurement campaigns to collect training data, which are rarely feasible in practice. Such databases usually include the true distance between two devices, the measured distance, the CIR of the signal, and a label which indicates whether the measurement was acquired in LOS or in NLOS. Therefore, such measurement campaigns are extremely time-consuming. Moreover, models trained on data from one location have difficulties in achieving the same accuracy when tested on measurements from other locations [7]. Therefore acquiring new data *from every new environment* in which a localization system is deployed is even more intractable.

For this reason, in [P4], we devise a detection and mitigation method for measurements acquired in NLOS, which does not need any training data. The proposed framework learns in real-time which measurements were acquired in NLOS based on the localization errors they incur. The framework then trains a model to recognize such measurements in the future based on their CIRs. Our method relies on distance *residuals* of individual anchors, i.e., the difference between the distance measured by an anchor and the distance between the location of the tag estimated using *all anchors* and the said anchor. In other words, the residual distance of an anchor is the deviation of the measured distance from the location computed using the measurements from all anchors. Because usually only a few anchors will experience a bad link to the tag, the correct location of the tag might still be recovered if there are enough anchors with good links to the tag. In cases in which there is more than one anchor in NLOS with the tag, the distance residuals need to be computed over all combinations of anchor subsets in order to identify the ones that cause large deviations.

To the best of our knowledge, there are two works [109, 11] that identify and mitigate NLOS measurements through anchor residuals. In [109], the authors used anchor residuals to train a classifier, whereas we use the CIR features of those measurements to train the model. In [109], the authors simulated a ranging NLOS error $e_{\text{NLOS}} \sim \mathcal{U}(0.75, 3.5)$, i.e., sampled from a uniform distribution in the interval 0.75–3.5 m. The LOS errors were normally-distributed with zero mean. However, in our measurement campaign, NLOS errors were spread in a much smaller range, of 0.25–0.8 m. In this case, it is more difficult to distinguish NLOS measurements based only on the residuals, which is why we perform the mitigation in two steps: first by identifying NLOS anchors based on their residuals, and then by training a model to recognize anchors with large errors/residuals based on their CIR.

In [11], the authors used residual analysis to identify NLOS errors and a combination of a fuzzy logic algorithm and Kalman filters to mitigate the NLOS errors. The method was evaluated through simulation for different types of NLOS error distributions (Gaussian, uniform, and exponential) and through measurements using UWB devices. In comparison with this work, in [P4] we also provide a self-learning NLOS detection method which can be useful beyond reducing the localization accuracy, for instance to detect obstacles, to create building maps, or to estimate the crowd density in a room. In [11], the authors focused more on NLOS error mitigation than on NLOS identification and they did not mention the accuracy of the classification method alone.

There are also a number of works using unsupervised or semi-supervised ML models. In [24], the authors propose an unsupervised NLOS identification method. However, the proposed solution can classify data only in bulk (so not online), whereas our method could also be used online.

In [85] and [129], the authors proposed NLOS identification methods based on pre-trained models. In [85], a convolutional neural network (CNN) is trained using data from one environment and updated with data from another environment. The method is evaluated in two office environments with similar setups. It is not clear how well the model could be transferred between two very different environments, e.g., a mall and an industrial hall. This is also an open research question for our method. However, the differ-

ence is that we do not rely on a pre-trained model, but can train such a model online. In [129], a pre-trained model is improved by retraining using unlabeled samples. This approach could also be used to improve the accuracy of our model.

4.1.2 SOTA on BLE RSS Variability

BLE was developed as an energy-efficient communication protocol for wireless personal area networks. In recent years, multiple works have investigated the capabilities of BLE communication for ranging and localization purposes [87, 26]. The RSS variability of WiFi signals has been extensively studied in [19, 40, 135, 128, 99]. However, no such comprehensive work yet existed for BLE signals. Previous work focused on only one source of RSS fluctuations or observed the combined effect of multiple error sources.

In contrast to prior work, in [P5], we isolated the main factors that lead to RSS variability, such as time, space, hardware, advertising channels, the distance between the TX and the RX, orientation, and obstructions, and we evaluated the impact of each factor on the observed RSS. We also documented a new source of errors: the influence of combo WiFi-BLE chips on the RSS of BLE signals.

Our measurement campaign includes the same type of measurements acquired in parallel at two sites: at University Politehnica of Bucharest (UPB) and at Tampere University (TAU). The Author was responsible for the measurement campaign conducted at UPB and the analysis of the data concerning the impact of the BLE channels, multipath propagation, device orientation, and transmit power on the variability of BLE RSS measurements.

In the following, we present the SOTA on each of the error sources investigated in [P5] and our main contributions regarding them.

BLE Channels. BLE has 40 channels of 2 MHz bandwidth each, indexed as shown in Figure 4.2. Three channels (37, 38, and 39) are reserved for advertising and the rest are for data transfer. The advertising channels are spread over the 2.4 GHz band at center frequencies of 2.402, 2.426, and 2.48 GHz to avoid interference with other devices operating in the same band. In [79], channel 39 was found the most reliable, since it is further away from the cen-




Figure 4.2 BLE and WiFi channels.

ter frequency of WiFi channels, also illustrated in Figure 4.2. In contrast, the other advertising channels partially overlap with other WiFi channels.

Variations between the RSS of advertising beacons sent on different channels at the same location was up to 15 dB in [33] and up to 6 dB in [26]. This can happen because the RSS varies inversely with the squared carrier frequency and because antennas do not always have a flat response over the entire bandwidth. Distance estimates can be improved by modifying the pathloss equation according to the channel on which the beacon was transmitted [33, 73], but knowledge of the channel is often not available by default at the receiver. Moreover, smartphones usually hop between different advertising channels. In [33], the authors exploited some patterns with which smartphones switch between different advertising channels to recover information about the center frequency.

Multipath propagation. The channel-dependent multipath fading of BLE signals was studied in [26]. The effect of multipath fading can be reduced by averaging multiple RSS values in a window. However, this solution is, in practice, hindered by the low rate with which devices send advertising beacons, in order to save battery power. Fluctuations in the RSS at the same TX–RX distance were found to be as high as 6 dB in [75] and even 20 dB [59] or 25 dB [32] over short time periods. In [59], the authors noted that the average RSS does not always decrease with the distance according to pathloss models [65].

Device orientation. Changes in the device orientations can result in RSS fluctuations of up to 30 dB [70] or 3 dB [96] even at the same distance and location. The manner in which people hold their mobile devices (e.g., in hand,




Figure 4.3 Overview of the main steps in the proposed NLOS error mitigation technique.

inside a pocket or a bag, etc.) influences the relative orientation between TX and RX antennas.

Transmit power and hardware. The RSS depends on the transmission power of the signal and the receiver’s RF front-end and antenna gain. Because BLE chipsets from different manufacturers have different settings and designs, the measured RSS can vary depending on hardware even under the same environmental conditions. In [84], transmitters from different vendors that transmitted at *different* powers had almost the same RSS. In [10], the authors show that the transmission power affects the localization accuracy and propose a ML-based method to learn the TX power of beacons in order to maximize the localization accuracy.

4.2 Self-learning NLOS Error Detection and Mitigation for UWB Localization Systems

4.2.1 Overview of the Proposed Framework

In this work, we propose a method for detecting and mitigating NLOS errors in UWB localization systems. The proposed method does not require manually-collected training data nor makes assumptions about the channel statistics.

Figure 4.3 shows the overview of our approach. When a tag is first deployed in an area, it starts the *initial phase* of the algorithm, in which it collects distance measurements and CIRs from the anchors in the area. In this first step, we use *residual analysis* to label the measurements as LOS, NLOS, or




Figure 4.4 Using noise-less distance measurements, a geometric localization algorithm would estimate the tag’s location at the intersection point between the circles determined by the anchor–tag distances (Figure 4.4a). However, in practice, the measurements are noisy so the circles do not have a single intersection point anymore (Figure 4.4b). The distance residual (r_i) of an anchor A_i is the difference between the measured distance (denoted by d_i) and the Euclidean distance between the estimated location of the tag and the location of the anchor (denoted by d'_i).

ambiguous. If there are more anchors than the minimum necessary for 2D or 3D localization, we can compute locations using each subset of anchors. This is a reasonable assumption, as typically tens of anchors can be deployed over a certain area for centimeter-level accuracy.

Since NLOS anchors¹ introduce higher location errors than LOS anchors, they also have higher distance residuals. A distance residual is defined as the difference between the measured distance and the Euclidean distance between the anchor and the estimated position of the tag. Distance residuals are illustrated in Figure 4.4b, where all anchors have noisy measurements. The location of the tag is found in the intersection *area* of the circles determined by the anchor–tag distances. In a noise-less case, like in Figure 4.4a, the tag’s location is at the intersection point of the three circles and the residuals are equal to zero.

The residual-based labeling step has an accuracy between 70–80% but it cannot be used on all sets of measurements, since for small NLOS errors the separation between LOS and NLOS residuals is not clear. This is why,

¹“NLOS anchor” is shorthand for “anchor in NLOS with the tag.”

besides the LOS or NLOS label, the method can also assign to a measurement the “ambiguous” label, when it cannot say with high confidence if the measurement was acquired in LOS or NLOS. Moreover, the residual-based labeling step requires computing the location using all anchor combinations, which scales with $\mathcal{O}(2^N)$ for N anchors. Therefore, for a large number of anchors, computing the residuals over all anchor subsets becomes computationally intensive.

Therefore, we introduce the second step (named “Model Training” in Figure 4.3), in which we train a Random Forest (RF) classifier which takes as input the CIR features of the measurements and the labels predicted in the first step. The model can recover the correct class boundary even with noisy labels, reaching a higher classification accuracy than the residual-based labeling, of over 90%. After the RF model is trained, it can directly classify all distance measurements as LOS or NLOS and the residual-based labeling can be skipped. Classifying samples using RF has a constant complexity regardless of the number of anchors.

We further propose a location-correction method based on identified LOS/NLOS measurements which does not discard the NLOS measurements. We evaluate the accuracy and localization error of the proposed method through a simulation of a localization network but based on a database of real distance measurements. The simulation therefore resembles a realistic localization scenario.

More details about the implementation of the residual-based labeling and model training steps can be found in [P4]. In the following, we will focus on the results obtained by the proposed framework.

4.2.2 Simulation Framework

We simulate a localization scenario based on a database of real UWB measurements collected in [P2] to evaluate the feasibility and performance of the proposed method for NLOS error detection and correction. Some of the locations in which we acquired measurements are illustrated in Figure 3.2.

Figure 4.5a shows a typical localization application in which four anchors are deployed over the perimeter of a localization area. However, because there are two rooms, some of the anchors might be in NLOS with the user’s

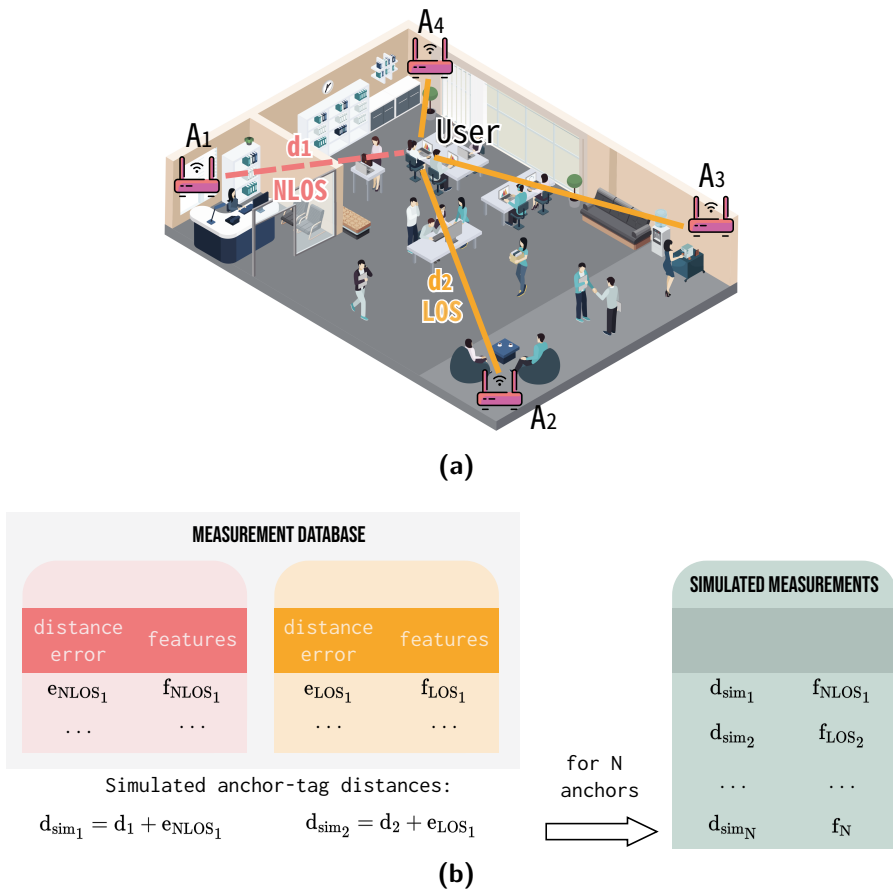


Figure 4.5 (a) Simulated localization setup and (b) data generation process based on the measurement database.

device. We consider between 5 to 9 anchors distributed uniformly on the perimeter of the area. We consider a grid of approximately 1700 true locations of the tag spaced 20 cm apart within the area delimited by the anchors. For each true location of the tag, we choose a percentage of anchors (either 0%, 30%, or 50%) to be in NLOS with the tag. The anchors which are in NLOS with the tag are chosen at random among all anchors.

Figure 4.5b explains the data generation process. We start from a measurement database collected in [P2] which contains distance errors and their corresponding CIRs. We simulate the distance measurements between each anchor and the tag by adding to the true distance a distance error selected from the measurement database. The distance error is selected depending on




Figure 4.6 Label generation process using residual analysis.




Figure 4.7 Model training and testing based on generated labels.

whether the anchor is in LOS or in NLOS with the tag. For instance, if we consider that the anchor is in NLOS with the tag, we will select one of the measurements acquired in NLOS from our database. We also store the CIR corresponding to the selected measurement.

From the CIR, we extract CIR features typically used in NLOS detection: the energy of the received signal, the maximum amplitude of the signal, the RMS delay spread, the mean excess delay, the kurtosis, and the difference between the TOA and the time at which the CIR has the maximum amplitude.

Next comes the label generation process, shown in Figure 4.6. We feed the N simulated distances from each anchor to a localization engine. The engine computes the 2D location and average anchor residuals over all anchor subsets. The residual analysis block takes as input the average anchor residuals and predicts the labels of the measurement between each anchor and the tag. The label can be either LOS, NLOS, or ambiguous (in case the system cannot determine the LOS/NLOS label with high confidence).

The model training and testing part is detailed in Figure 4.7. After the labels are predicted for N anchors and M locations, they are split into a training set and a test set. The training set contains the measurements predicted as LOS/NLOS, while the test set contains the ambiguous measurements. The training set is used to train a Random Forest classifier, which takes as input the CIR features and the labels predicted in the residual analysis step. Once

Table 4.2 Performance of unsupervised residual-based labeling in UWB localization [P4].

N	Q [%]	Accuracy [%]			Classified instances [%]
		LOS	NLOS	Balanced	
5	30	68.2	98.7	83.4	37.6
	50	66.8	82.1	74.4	40.2
6	30	68.3	98.9	83.6	29.3
	50	72.0	76.5	74.2	39.6
7	30	64.7	84.5	74.6	34.6
	50	62.4	79.5	70.9	37.6
8	30	74.4	81.1	77.7	32.8
	50	69.9	73.3	71.6	38.1
9	30	70.2	79.2	74.7	36.8
	50	66.0	72.5	69.2	35.2

the model is trained, we can use it to directly classify all measurements, without going through the residual analysis step. Finally, we mitigate the NLOS measurements and reduce the localization error.

4.2.3 Results

Residual-based labeling. Table 4.2 shows the accuracy in correctly identifying LOS measurements, NLOS measurements, and the balanced accuracy between both. It also shows the percentage of classified instances for the optimum KDE shaping parameter h used in the residual-based labeling step. More details about how the optimum value was found can be found in [P4]. The classification accuracy is higher for 30% NLOS anchors than for 50%. This is expected, since when more anchors are in NLOS with the tag, the location estimate is more skewed and the residuals of all anchors (not only of NLOS anchors) are larger. In this case, LOS and NLOS anchor residuals are harder to distinguish.

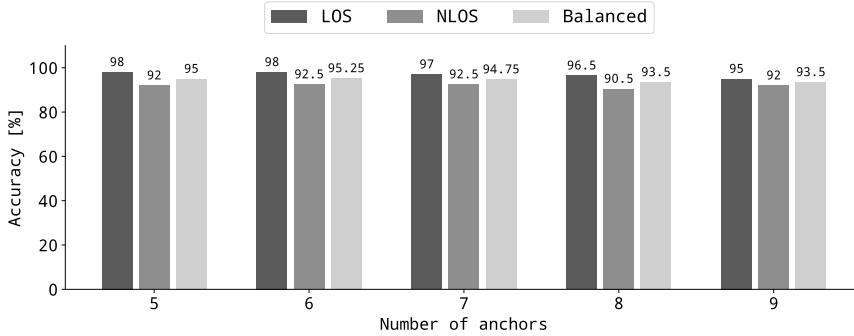


Figure 4.8 The accuracy of the trained random forest (RF) classifier using data labeled with the residual method.

Supervised classification. We now evaluate the accuracy of the RF classification applied on measurements labeled with residual analysis. The training set consists of labeled measurements and their features. Please refer to [P4] for more details about the simulation setup.

Figure 4.8 shows the LOS, NLOS, and balanced accuracy of the RF model. The accuracy in identifying NLOS measurements exceeds 90% in all cases, while the accuracy in identifying LOS measurements is higher than 95%. Compared with only the residual-based labeling, we gain 10–20% accuracy. Although it might seem surprising that the RF classification accuracy is that high even though the labeling accuracy is, in some instances, as low as 63%, our scenario fits into the research area of learning in the presence of noisy labels. Research has shown that certain ML models can recover to a certain extent the correct class boundaries even in the presence of outliers or anomalies in the training data [30].

NLOS mitigation. We also devised a strategy to mitigate NLOS measurements after their detection. The detailed algorithm can be found in [P4] and, in the following, we will only present a summary of it.

If there are enough LOS anchors to compute one location and the set of anchors is not degenerate (i.e., the anchors are not colinear), we could use only the LOS anchors to compute the location, since they generally have the smallest distance errors. However, we noticed that we obtain better location estimates if we correct NLOS measurements first and then use them for local-

Table 4.3 Localization error before and after the proposed NLOS mitigation method [P4].

N	Method	Q = 30% NLOS		Q = 50% NLOS	
		Mean [m]	Std. dev. [m]	Mean [m]	Std. dev. [m]
5	No mitigation	0.14	0.27	0.20	0.39
	Proposed	0.07	0.06	0.11	0.14
6	No mitigation	0.11	0.08	0.19	0.31
	Proposed	0.06	0.04	0.10	0.17
7	No mitigation	0.14	0.30	0.21	0.64
	Proposed	0.07	0.05	0.08	0.07
8	No mitigation	0.14	0.36	0.22	0.67
	Proposed	0.06	0.05	0.08	0.06
9	No mitigation	0.13	0.35	0.20	0.64
	Proposed	0.05	0.04	0.07	0.14

ization, since the LOS anchors might not have an optimum placement, e.g., if the tag is located outside the convex hull of the anchors. In these areas, multilateration algorithms usually have larger errors than inside the convex hull [98].

For correction, we first estimate the intermediate location using only the LOS anchors. We compute the residuals of the NLOS anchors based on the intermediate location. Then, we subtract the residuals from the measured distances of NLOS anchors. We estimate the final location using the distance measurements of LOS anchors *and* the corrected distances of NLOS anchors. Our algorithm also targets cases in which there are not enough LOS anchors to compute one estimate or the LOS anchors are colinear.

Table 4.3 compares the localization errors obtained with the proposed mitigation algorithm (denoted by “proposed”) with those obtained when using all anchor–tag distances, without mitigation. The localization errors of the proposed method have $1.8\text{--}2.8\times$ smaller mean and $1.8\text{--}11.6\times$ smaller stan-

standard deviation after NLOS mitigation. On average, the algorithm reduces the mean and standard deviation of localization errors by 2.2 and 5.8 times, respectively. Therefore, the proposed method can successfully mitigate localization errors caused by NLOS propagation.

4.3 Signal Fluctuations in BLE Communication

This section describes a part of the sources of BLE RSS fluctuations identified in [P5], focusing on those where the Author's contribution was the largest. As we mentioned, the observations are based on a measurement campaign that took place, in parallel, at UPB and TAU. Figure 4.9 shows examples of locations at which we acquired measurements, in LOS and NLOS, at both universities.

In total, we documented eight sources of variability in BLE RSS data: time, space (or environment), hardware, advertising channel used, distance, device orientation, the co-existence of WiFi and BLE on the same chipset, and obstructions between the TX and the RX.


In this section, we will discuss only the variations in time, caused by hardware, with orientation, and in LOS/NLOS propagation, while the rest can be found in the publication [P5].

4.3.1 Variations in Time

We investigated the stability of the BLE RSS over time, when the devices were left unmoved at one location. Figure 4.10 shows the RSS distribution in increasing time window sizes. The median RSS decreases with 3 dB from the time window of 2 min to the time window of 10 h. However, the mean RSS (diamond-shaped marker) is relatively stable in all window sizes and varies with less than 1 dB. The results are meaningful to determine the recording time needed to obtain a stable value of the mean RSS. In this case, a recording time of 2 min is enough to capture the characteristic RSS in a location.



(a) LOS at TAU



(b) LOS at UPB



(c) NLOS with wall at TAU




(d) NLOS with wall at UPB



(e) NLOS with plexiglass at UPB



(f) NLOS with human body at UPB



(g) NLOS with door at UPB



(h) NLOS with wall and whiteboard at UPB

Figure 4.9 The pictures of the receiver and transmitter in LOS and NLOS with wall acquired at UPB and TAU (Figure 4.9a to 4.9d). Figure 4.9e to 4.9h show NLOS scenarios with a plexiglass panel, human body, a door, and a wall with a whiteboard at UPB.

4.3.2 Variations Caused by Hardware

In all the measurements, we used as the evaluation hardware Raspberry Pi 4 Model B boards equipped with the same model of BLE chip. However,




Figure 4.10 The evolution of the BLE RSS distribution over different periods, up to 10 hours. The measurements were acquired in LOS, on channel 37, at a distance of 1 m between the devices.




Figure 4.11 The impact of hardware on the RSS in recordings acquired with different device pairs placed at a distance of 2 m at exactly the same location. The median RSS varies with 5 dB, even though the boards contain the same BLE chipset model.

we noted that, when performing measurements with different boards at the same locations, the mean RSS values differed depending on the pair of boards used. In order to evaluate the magnitude of the fluctuations, we performed measurements with different device pairs at the same locations, at a distance of 2 m, configured to operate only on channel 37. More details about the devices used in each pair can be found in [P5].

Figure 4.11 shows the RSS distribution of each pair of devices. In device pairs 1, 2, and 3 we used the same transmitter, while in device pairs 3, 4, and 5 we used the same receiver. The median RSS varies with up to 5 dB even between devices from the same model. This experiment shows the difficulty of calibrating the transmitter and receiver efficiency according to the device model, since even devices that share the same device model have RSS variations of several dB.

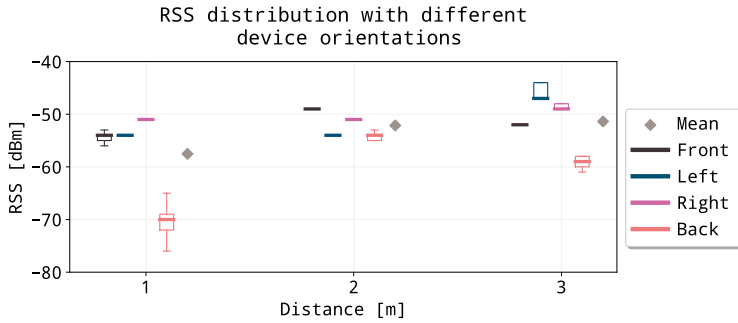


Figure 4.12 The impact of orientation on the BLE RSS. In the back, left, and right poses the receiver was rotated clock-wise with, respectively, 180° , 90° , and 270° with respect to the front orientation.

4.3.3 Variations with Orientation

We evaluated the effect of the relative orientation between the TX and the RX on the BLE RSS. We analyzed four poses, where the pose of the transmitter is fixed and the receiver is rotated clock-wise with 90° , 180° , and 270° with respect to the “front” orientation, resulting in the “left,” “right,” and “back” poses, respectively. The devices were placed on tripods at distances of 1, 2, and 3 m, and the receiver was rotated around its center axis.

Figure 4.12 shows the RSS distribution obtained in each of the scenarios. We notice that the RSS fluctuates with the pose for a particular distance but that there is no orientation which results in a higher RSS at *all* distances. The “back” pose has a lower median RSS than the others, most likely because in this pose the metallic USB and Ethernet ports of the board are placed in the direct path of the signal and can attenuate the signal. We also noticed inconsistencies with the distance, in which the average RSS at a distance of 1 m is lower than the RSS at 2 m and 3 m.

4.3.4 Variations in LOS and NLOS Propagation

We investigated the effect of different types of obstructions on the BLE RSS. We acquired measurements at both locations (UPB and TAU) in LOS and NLOS with wall shadowing. At UPB, we also tried other types of obstructions: human body, wall and whiteboard, door, and plexiglass panel. More

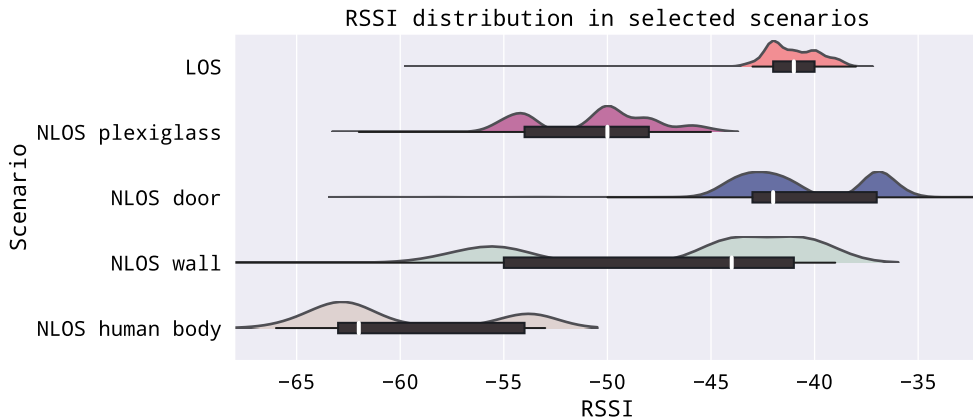


Figure 4.13 Comparison of selected RSS distributions acquired on all advertising channels at 2 m, in LOS, NLOS with a plexiglass panel, NLOS with a door, NLOS with a wall and a whiteboard, and NLOS with human body shadowing. Underneath each distribution there is the boxplot representation, where the black box is the IQR that contains 50% of the values and the white marker denotes the median RSSI.

information about the setup and the procedure for selecting the measurements to avoid duration biases can be found in [P5].

Figure 4.13 shows the distribution of RSS measurements in selected LOS and NLOS scenarios at a distance of 2 m between the devices. Measurements acquired in hard NLOS scenarios (with wall or human body between the devices) show higher spread than in LOS or in NLOS with a shallow obstruction (the plexiglass panel or the door). The median RSS in the NLOS scenario with a door is -42 dBm, which is only 1 dB smaller than the median RSS in LOS for the same distance. Such inconsistencies can also be caused by the fact that we used only a single recording to plot the statistics. However, we showed in [P5] that single recordings can deviate from statistics computed on aggregated data with more than 10 dB.

The inconsistency can be also caused by the fact that most of the NLOS statistics were computed based on a single recording. As we show in [P5], single recordings can deviate from statistics computed on aggregated data with more than 10 dB. Therefore, we also compared LOS and NLOS setups when aggregating data from multiple recordings from campaigns conducted independently at TAU and UPB. The results are shown in Figure 4.14.

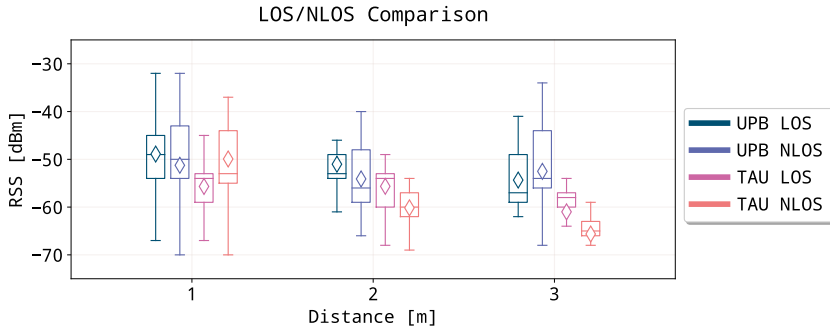


Figure 4.14 Comparison of BLE RSS distributions based on data acquired in LOS and NLOS with a wall at UPB and TAU at distances of 1, 2, and 3 m.

First, although we would expect the mean RSS to be higher in LOS than in NLOS at the same distance, this is not always the case. At 1 m, the mean RSS values in LOS and NLOS are almost equal at both sites, while at 3 m, at UPB, the mean RSS in NLOS is *higher* than the one in LOS. The standard deviation of RSS measurements acquired in NLOS is generally higher than in LOS. This is expected, since obstructions can cause the diffraction and reflection of the signal, which can introduce higher RSS fluctuations.

In the literature, multiple solutions have been proposed to deal with localization or ranging errors caused by RSS instabilities in NLOS propagation [73, 17, 52]. Most of the solutions, however, are based on large datasets for training ML models. However, as we showed in our measurement campaign, the data collected at different sites might show different statistics, so solutions customized for one environment alone might not work in another setting.

4.4 Summary

In this chapter, we investigated the effects of multipath propagation on the ranging and localization accuracy of narrowband and wideband systems, exemplified through the BLE and UWB technologies, respectively. The system bandwidth has important implications on the type of localization methods that can be employed and their accuracy. BLE is advantageous owing to its very low energy consumption and widespread availability. However, most

proximity detection applications based on BLE use RSS measurements and hence their accuracy is in the range of meters. UWB, on the other hand, allows centimeter-level ranging and localization accuracy but has higher power and energy consumption than BLE.

In [P4], we proposed a method to detect and mitigate NOS errors in UWB localization systems, which does not need any manually-collected training databases. The novelty lies in using residual analysis to automatically label measurements as LOS or NLOS. Therefore, we eliminated the need for performing an expensive and time-consuming measurement campaign. Instead, the measurements can be collected and processed online and the model can be trained on the user device. The complete framework reduces the mean localization error by $2.2\times$ and its standard deviation by $5.8\times$ compared to the case in which no error mitigation method is applied.

In [P5], we identified and characterized eight main sources of variability in BLE RSS measurements, basing our observations on a measurement campaign performed in tandem at two locations: University Politehnica of Bucharest, Romania, and Tampere University, Finland. Having the experiments replicated at two locations enabled us to identify trends in the data without being biased by local artifacts. Besides documenting the most important sources of variability of BLE RSS measurements, we also identified a new factor, namely the co-existence of the WiFi and BLE functionalities on the same chipset. The measurements were published in an open-source format [57], to aid future research focused on proximity detection with BLE signals.

5 TDOA LOCALIZATION WITH FLEXIBLE ANCHOR ALLOCATION

Ultra-wideband (UWB) localization systems currently face high robustness and scalability demands, as they operate in challenging multipath environments and need to satisfy an ever-growing number of users. Time-difference of arrival (TDOA) localization in which the user device remains passive and listens to beacons sent by the anchors can scale to an unlimited number of users. This technique is also called *downlink* (DL) TODA, because the communication occurs only on the downlink, from the anchors to the tag.

The drawback of DL TDOA is that the anchors need to synchronize their clocks. To bypass this requirement, a common approach is to designate a reference anchor which periodically sends synchronization beacons. The rest of the anchors respond to the beacon, while the tag passively listens to every message and estimates its location based on their TDOAs. This method is illustrated in Figure 5.1a and will be called the “classic” DL TDOA. However, the classic TDOA is not robust to obstructions between the devices, especially between the reference anchor and the tag.

In publication [P6], we propose FlexTDOA, a localization system with a flexible TDMA scheduling scheme, in which the reference anchor, the number of responders, and the order of responders can change every time slot. The TDMA scheme of FlexTDOA is illustrated in Figure 5.1b, where the first anchor is the initiator in the first time slot, the second anchor is the initiator in the second time slot, etc.

We compare FlexTDOA against the classic TDOA and TWR-based localization in a deployment that consists of ten anchors and one tag. We evaluate scenarios with and without obstructions. Results show that FlexTDOA, the proposed method, has the highest localization accuracy in most of the evaluated scenarios. FlexTDOA reduces the localization error with up to 38%

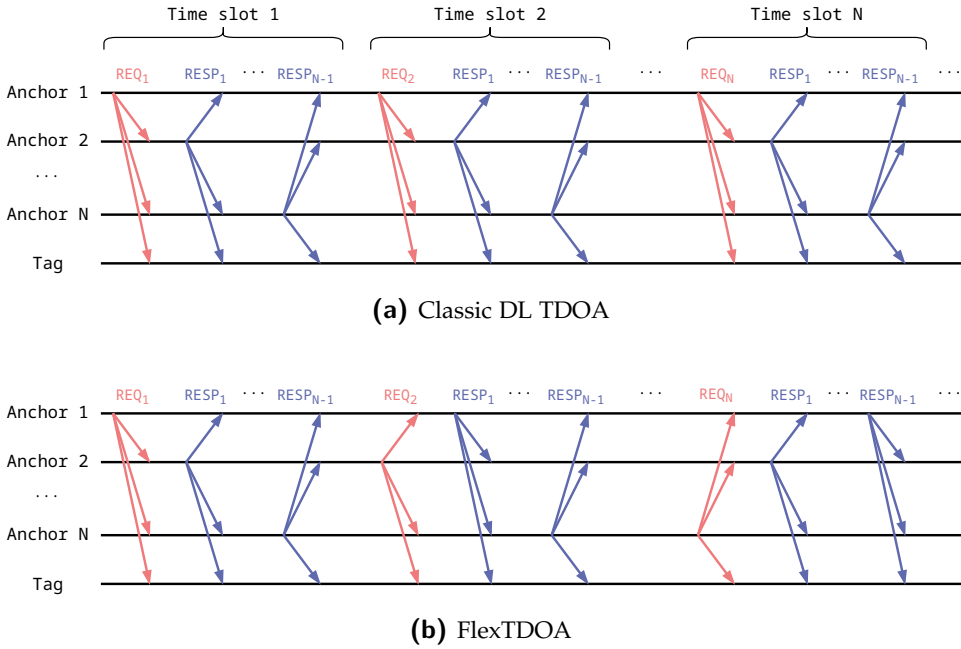


Figure 5.1 The scheduling of anchor transmissions in (a) classic DL TDOA and (b) FlexTDOA. In the classic TDOA, there is a single reference anchor which broadcasts messages. The rest of the anchors respond to this initial message (usually in a pre-defined order). In FlexTDOA, we propose a flexible scheduling scheme in which all the anchors can become reference anchors and the order of responses can also change.

compared to the classic TDOA. We also evaluate the impact of multiple parameters such as the number of responses, their order, or the number of anchors on the localization accuracy.

This chapter is organized as follows. Section 5.1 presents the SOTA in TDOA localization and how our work differentiates from previous solutions. Section 5.2 introduces the TDMA scheduling algorithm which is the core of FlexTDOA. Section 5.4 shows the most important results achieved by the proposed localization system. Finally, Section 5.5 summarizes the results and proposes directions for future research.

5.1 State of the Art

Next, we review the most important state-of-the-art on TDOA localization. We first discuss the theoretical differences between TOA (or range-based) and TDOA localization in Section 5.1.1. Section 5.1.2 presents the most relevant UWB localization systems that implement *DL* TDOA, since they are comparable to FlexTDOA in terms of multi-user scalability. We also present a comparison against uplink (UL) TDOA localization systems which were optimized for multi-user scalability. Section 5.1.3 discusses the carrier frequency offset (CFO) estimation technique used in the proposed method to correct TDOA measurements.

5.1.1 TOA and TDOA Comparisons

In [121], TOA and TDOA methods with correlated and uncorrelated noise are compared from a theoretical standpoint. When the tag transmits an uplink message which arrives at multiple anchors, the measurements observed by the anchors are *correlated*. If the tag is passive and only the anchors exchange messages, if enough time passes between two consecutive transmissions, the measurements at the tag are *uncorrelated*. It is shown that the circular (TOA) method with an unknown clock offset and the hyperbolic (TDOA) method with correlated noise are equivalent. It is also shown here that both hyperbolic methods, i.e., with correlated and uncorrelated noise, achieve similar performance. Therefore, there is no loss in localization accuracy when using DL TDOA instead of UL TDOA.

5.1.2 TDOA Localization

Although TDOA localization systems have been extensively studied in the literature, most previous works have focused on the classic TDOA scheme in which the reference anchor and the order of responses are fixed. In this work, we evaluate the impact of the channel diversity brought by using more pairs of anchors than in the standard configuration. We also evaluate the impact of several parameters (e.g., number of responses and of anchors) on the localization accuracy. We compare the FlexTDOA scheme against the classic

TDOA and TWR localization, which are the most widely-used techniques in UWB localization. Moreover, most previous TDOA systems have been evaluated only in LOS conditions, whereas we also covered the performance of all three localization algorithms in a NLOS scenario as well.

In [58], the authors propose a DL TDOA localization system with a clock synchronization protocol based on a reference anchor. It is mentioned, however, that the system does not scale to large networks. With eight anchors, the system obtained a 2D localization root-mean squared error (RMSE) of 14 cm and a 3D RMSE of 28 cm. Using a comparable setup of seven anchors, FlexTDOA obtained a 2D RMSE of 16.2 cm and a 3D RMSE of 23.5 cm (averaged over all considered locations). Therefore, the performance of FlexTDOA is similar to the system proposed in [58]. In [38], a similar approach to [58] is proposed. Here, the pairwise clock error is tracked using a Kalman filter that can handle variations in the reception period. In the proposed method, each new anchor joining the network has to synchronize itself to all the other anchors already in the network. The authors only evaluate the case where each anchor has an allocated (fixed) time slot for transmission. In comparison, we also evaluate more complex schemes in which the transmission order of the anchors can change. This flexibility is facilitated by the fact that we do not track pairwise clock errors and can hence easily change the order of transmissions.

In [88], the authors propose a DL TDOA scheme in which the anchors respond only to the previous message, as opposed to responding to a synchronization beacon, as in our case. The mean and maximum localization errors obtained with a configuration of four anchors were 31 cm and 81 cm, respectively.

Although the works in [35, 16] are said to perform concurrent ranging, they essentially implement the classic DL TDOA scheme. However, in these works, the novelty is brought by processing all the responses within a single reception period. This is done by leveraging multipath information from the CIR. Because of several hardware limitations of the DW1000 chipset used in these works, the 90% error obtained by a setup with five anchors in [35] was 55.8 cm, so about $2\times$ larger than the one obtained by FlexTDOA in a similar setup in which the tag is placed in the center of the room. However, the

limitations were solved in the new-generation DW3000 chipset used in our work, so in the future it would be interesting to integrate concurrent ranging in our system in order to decrease the energy consumption of the tag.

In [118], a UWB-based TDOA localization system called ATLAS is introduced. Even though the system also uses a reference anchor, it is not clear if the tag is completely passive. Because the paper mentions that only tags on a white list are processed at a localization server, this implies that the localization system must be aware of the presence of the tags. This is usually achieved using uplink transmissions. Therefore, ATLAS uses a centralized processing to compensate errors in the TDOAs. In comparison, in FlexTDOA, the localization is offloaded to the user device. This preserves the user's location privacy. Iterations of ATLAS have been introduced in [120, 119]. However, in these works, it is clear that the tag is active, which poses scalability issues.

A recent localization system called VULoc has been recently proposed in [137]. The system implements DL TDOA. However, compared to our method, in VULoc, the reference anchors send an extra message after receiving all the responses. This means that the tag has to listen to an additional message compared to FlexTDOA. The most significant difference between our system and VULoc is that in FlexTDOA we also propose a highly configurable and flexible TDMA scheme for scheduling anchor transmissions. In comparison, in [137], the authors mention that a scheduling protocol is not needed since the tags are passive. However, we argue that, especially in large-scale deployments, there is a need to easily add or remove anchors from the localization system and to have a protocol for scheduling their transmissions. This allowed us to easily evaluate the performance of FlexTDOA with up to ten anchors, whereas the performance of VULoc was evaluated only with five anchors. In [137], the impact of channel diversity is also not evaluated. In a setting similar to ours, VULoc had a median error of 15.5 cm and a 90% error of 23.6 cm. In a similar experiment in which we placed the tag in the center of the room, FlexTDOA yielded a median error of 15.4 cm and a 90% error of 22.2 cm. Therefore, the performance of the two systems are similar with a small anchor deployment.

Although it does not implement a DL TDOA scheme (so the tag is *active*), the work in [124] proposes a scalable UL TDOA localization scheme called

TALLA. The tags broadcast uplink messages, which are received by wirelessly-synchronized anchors. A server maintains the synchronization necessary to correct clock drift errors in TDOA localization. The server can synchronize to any reference anchor that is part of the system. This approach provides flexibility in large-scale deployments, but poses more privacy risks, as the network can record the location of all the tags. In comparison, in FlexTDOA, the network has no information of the users' locations since the tags localize themselves. The system is evaluated in an experiment using 12 anchors in an area of 100×60 m. Simulations were performed for more anchors. The effect of the synchronization rate and the number of anchors (4 or 8) were evaluated experimentally. The authors found only small improvements with more anchors, which confirms our own observations.

One important contribution of [P6] is that we evaluated the performance of the system in NLOS *experimentally*. Most previous works considered only LOS scenarios. VULoc [137] was also evaluated in NLOS conditions. The authors also proposed an anchor selection algorithm based on a confidence parameter chosen empirically. Whereas in [137] the goal of the NLOS evaluation was to demonstrate the efficacy of the anchor selection mechanism, we also quantified the effect brought by the added channel diversity when using the TDOAs between all anchor-tag trios. In [45], the authors derived theoretical error bounds for TDOA localization under NLOS propagation. In [141], an anchor-selection method for NLOS scenarios was proposed and evaluated using simulations. In [140], a sensor-placement strategy for cluttered multipath environments was proposed. The method was evaluated using experiments. In [91], an UL TDOA localization system that accounts for NLOS conditions was proposed and evaluated in experiments. In [123], a method for selection anchor pairs in an UL TDOA localization system was developed. The method takes into account measurement errors incurred by NLOS propagation.

5.1.3 Clock offset correction

In the proposed TDOA scheme, we avoid tracking clock parameters like in previous works [58, 38] using, for instance, Kalman filters. Instead, we use the CFO estimation feature of the DW3000 chipset that allows us to compen-

sate for the relative clock offset between a transmitter and a receiver directly at the receiver. The method has been described in [20], where the authors derived the systematic error for SS-TWR, A-TDOA, and SS-TWR with A-TDOA extension. The authors evaluated the error experimentally but only for TWR schemes. A similar CFO correction is applied in a TDOA scheme in [108]. However, in [108], the TDOA scheme is based on the alternative DS-TWR method. This means that the tag is active, so the system uses UL TDOA localization scheme. In [23], a CFO-assisted synchronization algorithm for TDOA localization has been proposed. Here, the CFO is used to correct the TOA at each receiver with respect to the reference node. Although FlexTDOA is different from the TDOA scheme proposed in [23], both works correct errors in a similar way. We also evaluate the feasibility of the CFO-based correction experimentally using a commercial UWB system. In [23], the method was evaluated in simulations and using software-defined radios.

5.2 Improved TDMA Scheduling for TDOA Localization

In the “classic” TDOA approach, there is a single designated reference anchor which broadcasts the synchronization message. The rest of the anchors respond to the broadcast in a pre-defined order. However, since the TDOAs are computed with respect to the message broadcast by the synchronization node, this method is not robust in case the path between the tag and the synchronization anchor is obstructed. In this case, the synchronization message will arrive at the tag with a delay δ . All the TDOAs derived from the subsequent responses will be biased with $-\delta$ ¹. This scenario motivated us to implement the proposed FlexTDOA scheme. In the proposed scheme, all the anchors in the system can become initiators and the order of responses can also change.

The TDMA scheme we proposed is shown in Figure 5.2. The same TDMA scheme can be configured for either TWR or TDOA localization. At this point, we do not distinguish between anchors and tags; instead, we consider

¹The change of sign happens because, in TDOA, the TOF between the initiator and the tag is subtracted from the TOF between the responder and the tag.

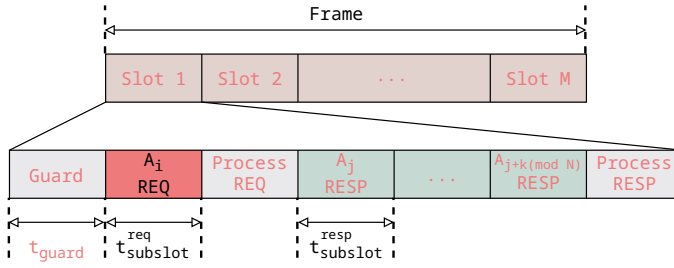


Figure 5.2 TDMA scheduling with changing initiator and responders.

all of them equally-participating nodes. We will make the distinction according to the implemented localization method.

The TDMA scheme consist of *time slots*. One time slot consists of a broadcast message, called *request*, sent by an initiating node, and K responses from other nodes. We denote the total number of nodes by N and the number of responses is always smaller than the number of nodes, i.e., $K < N$.

In a time slot, each transmission of a node takes place in a *subslot*. The duration of a subslot is equal to $t_{\text{subslot}}^{\text{req}} = 2 \text{ ms}$ if the transmission consists of a request message or equal to $t_{\text{subslot}}^{\text{resp}} = 250 \mu\text{s}$ if it is a response message.

A time slot starts with a guard time of $250 \mu\text{s}$. This is followed by the request of the initiator. The request includes, in the payload, the number of nodes that should respond, their ID, and the order in which they respond. All the other nodes in the system process this request. If a responder was scheduled to respond in the subslot with index $k \in \{1, \dots, K\}$, then the node will wait a time equal to $(k - 1) \times t_{\text{subslot}}^{\text{resp}}$ and then respond. At the end of the time slot, the initiator processes the responses.

For TWR-based localization, the tag will be configured as the initiator and the anchors will play the role of responders in the TDMA scheme.

We can implement four variants of TDOA localization with the proposed TDMA scheme:

- *Fixed initiator, fixed responders (FI-FR)*, or the classic TDOA. The initiator (or the reference anchor), as well as the order in which the anchors respond, are both fixed. The anchors respond in ascending order of their index.
- *Fixed initiator, changing responders (FI-CR)*. The initiator is fixed, but the



Figure 5.3 Office setup with ten anchors (A_0 to A_9) and one tag.

order of the responding anchors changes in a round-robin (RR) manner.

- *Changing initiator, fixed responders (CI-FR)*. The initiator changes every time slot in a RR manner, but the rest of the anchors respond in a fixed order, according to their index.
- *Changing initiator, changing responders (CI-CR)*. Both the initiator and the order of responses change every time slot in a RR manner. This is the scheduling proposed in FlexTDOA. This is the method illustrated in Figure 5.2, where a succession of time slots cycling through all the possible initiator anchors in the system is called a *frame*.

5.3 Evaluation Setup

We evaluate the localization systems in the office shown in Figure 5.3. The tag is mounted on a rail of approximately 150 cm which moves the tag and outputs its ground-truth location with millimeter-level precision. Figure 5.4 shows the anchor placement more clearly. The figure also shows four placements of the rail on which the tag was mounted, at which we evaluated the algorithms.

For the evaluations presented in the thesis, we implemented a localization algorithms which estimates the user’s location based on a series of consecutive measurements using squared error minimization. We denote this algorithm by AlgMin. AlgMin does not track the user’s location and it does

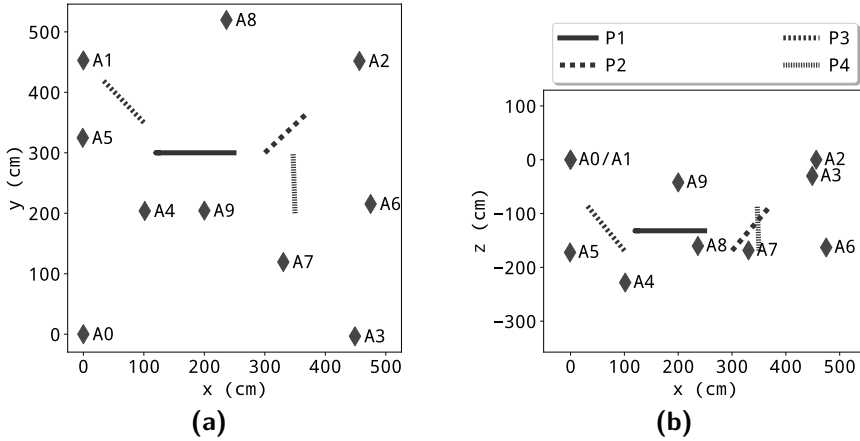


Figure 5.4 Setup of the anchors and the tag in the (a) xy and (b) xz planes. The anchors are denoted, according to their index, by A_0 to A_9 . The tracks P1 to P4 represent the positions of the linear actuator (on which the tag is moving) at which we evaluate the localization methods.

not smooth the estimated locations. Therefore, it is suitable to evaluate the impact of various parameters on the localization accuracy (e.g., the number of responses or anchors). The algorithm can operate with either TWR or TDOA data. We refer the reader to [P6] for more information related to the algorithm implementation.

We also implemented a second algorithm based on an Extended Kalman Filter, called AlgEKF, which updates the location for every incoming measurement. This approach is beneficial because we do not need to wait for the minimum number of measurements (four in the case of TWR localization and five for TDOA localization) in order to update the tag's location. However, AlgEKF smooths the location estimates and minimizes the impact of noisy measurements, which makes the comparison between the localization methods more difficult. Therefore, we use AlgEKF only when we compare setups that generate different numbers of equations per time slot, as in the case where we evaluate the impact of the number of responses for different tag speeds. We did not include this evaluation in the body of the thesis, but we direct the reader to [P6] for more information on the algorithm and its corresponding evaluation.

CFO correction is implemented for all localization techniques (based on

TDOA and TWR alike).

Throughout the rest of the chapter, localization errors are computed as the Euclidean error between the true location of the tag and the estimated location of the tag in a 3D space, as in Eq. (3.2).

5.4 Results

In [P6], we evaluated the impact of multiple parameters on the accuracy of TDOA measurements or of the localization accuracy: the order of response in a time slot, the number of responses per time slot for different location update rates, the number of anchors, the chosen TDOA variant (with fixed/changing initiator and/or responders), and NLOS propagation.

In this section, we will review the results only for selected parameters. In Section 5.4.1, we discuss the effect of the response order on the accuracy of TDOA measurements and its implications for localization. In Section 5.4.2, we evaluate the impact of each TDOA scheduling variant on the localization accuracy. In Section 5.4.3, we analyze the classic TDOA, the proposed FlexTDOA, as well as TWR-based localization under NLOS conditions.

5.4.1 Optimal Number of Responses per Time Slot

In the proposed TDMA scheme, we can vary the number of responses in a time slot between $K \in \{1, \dots, N - 1\}$, where N is the total number of anchors in the system. However, with more responses, there is a longer delay between the request and the last response. Because of the relative clock offset between the responder and the tag, the timing error increases proportionally with the delay between the request and the response.

To evaluate the magnitude of the TDOA measurement error, we record TDOA measurements during an experiment with $N = 10$ anchors, $K = 9$ responses, during which the tag does not move. Knowing the ground truth locations of the anchors and of the tag, we can compute the *true* TDOA between the tag and each two anchors. We denote the TDOA *measurement* between the tag and anchors A_i and A_j as $\tilde{d}_{TA_iA_j}$. Similarly, the *true* TDOA is

denoted by $d_{TA_iA_j}$. We then define the TDOA error as:

$$e_{TDOA} = \tilde{d}_{TA_iA_j} - d_{TA_iA_j}. \quad (5.1)$$

The TDOA error is computed for each anchor and for each response order. For instance, for a fixed initiating anchor A_i , the errors for the first response order are computed based on the measurements from all anchors A_j for $j \in \{0, \dots, 9\}$ with $A_i \neq A_j$, where A_j was the first anchor to respond in a time slot.

Figure 5.5a shows the distribution of TDOA measurement errors according to the order of the response (or subslot). The whisker spread of TDOA errors increases with 14.7 cm between the last and the first in the list of responses. The results are as we would expect according to the theory. The error caused by the relative clock skew between two devices increases with a longer waiting time between the initiator's message and the response. The measurements are therefore corrupted by noise.

Over many measurements, the mean error converges to approximately zero for every order of response. On the other hand, the noise in each measurement increases with the order of response. Therefore, for a relatively static tag, averaging the TDOAs over more responses could increase the accuracy. However, for tags moving at high speeds, where averaging is not possible because the location changes at a fast rate, it is preferable to use a smaller number of responses to minimize the clock drift error. In [P6], we further explore this aspect by evaluating the optimal number of responses for different simulated tag speeds.

Since the TDOA measurement error increases with the order of response, we evaluate how the *localization* accuracy changes with the number of responses in a time slot in FlexTDOA.

For this experiment, we keep the same number of anchors as before ($N = 10$) but we vary the number of responses $K \in \{1, \dots, 9\}$. The tag is moving along the P1 position of the rail shown in Figure 5.4. Figure 5.5b shows the TDOA localization error with an increasing number of responses. As predicted, the localization error is the lowest for the least amount of responses. However, the increase in the mean and IQR of the localization error for a higher number of responses is almost negligible: less than 3 cm with nine vs. one responses.




Figure 5.5 (a) TDOA *measurement* error variation with the order of response in a time slot. The TDOA measurement error increases with the index of response because of the higher clock drift error in the processing time. (b) TDOA *localization* error variation with the number of responses in a time slot. Similarly, the TDOA localization error increases with the number of responses because of the higher clock drift error present in later responses.

Although we would be tempted to say that one response per time slot is the optimal configuration, because it results in the lowest TDOA measurement error, there are more factors to consider. First, for a fixed period of time, the number of TDOA measurements decreases inversely proportionally to the number of responses per time slot. This is because the initiator’s request adds an overhead. This means that, over the same time period, we have less measurements for localizing tags with a smaller number of responses.

Second, the energy consumed by the tag to receive a certain number of TDOA measurements increases inversely proportionally to the number of responses, because of the same overhead. Publication [P6] includes a more in-depth analysis of the optimal number of responses given either a fixed time or a fixed energy budget for different localization update rates.

5.4.2 Channel Diversity

In this section, we evaluate to what extent the localization accuracy is improved by changing the initiating anchor and/or the order of responses in the four variants of TDOA scheduling: FI-FR (also called the classic TDOA), FI-CR, CI-FR, CI-CR (implemented in FlexTDOA). We perform three experiments for each scheme, for $(N, K) \in \{(5, 4), (7, 6), (10, 9)\}$. For these experiments, the tag is moving only along position P1 (Figure 5.4).




Figure 5.6 Comparison between FI-FR (classic TDOA), FI-CR, CI-FR, and CI-CR (FlexTDOA) for different numbers of anchors (N) and of responders (K).

Figure 5.6 shows the error distributions obtained with each of the setups. For five and seven anchors, the schemes with a fixed initiator (FI) have the highest localization errors. However, the trend is reversed for ten anchors. One explanation for this is that, when using ten anchors, we placed some of the anchors on the ground or on the tables. We noticed that these anchors have slightly higher ranging errors than the ones placed, for instance, on the ceiling. When these anchors with non-ideal conditions become initiators, they drive the localization errors higher. Therefore, in the setup with ten anchors, the localization errors obtained by the FI schemes are slightly lower than for CI schemes, because the initiating anchor in the former is placed at an ideal location.

We also notice that changing the responders does not significantly affect the localization accuracy. Instead, it is the *initiator* who plays a crucial role in the location estimation. This is because all the TDOAs in a time slot are computed relative to the initiator's time frame. If an error occurs when timestamping the initiator's message, then the error will affect every TDOA in that time slot. Because of this, starting from now, we will consider only the FI-FR and CI-CR schemes. We alternatively call them the classic TDOA and FlexTDOA, respectively. FlexTDOA is the proposed method for improving the classic TDOA localization.

We further evaluate to what extent the channel diversity improves the localization accuracy in LOS conditions when comparing the classic TDOA and FlexTDOA. We evaluate the localization errors for $(N, K) \in \{(5, 4), (7, 6), (10, 9)\}$ at three positions of the rail on which the tag moves, denoted by P1, P2, and P3 in Figure 5.4.




Figure 5.7 Distributions of the localization errors of classic TDOA (FI-FR) and FlexTDOA (CI-CR), in LOS, aggregated over all the evaluated positions (P1, P2, P3).

We let each experiment for each location, for every method, and combination of (N, K) run for approximately one minute and a half. This gives us around 7000 location estimates computed using AlgMin. We provide the aggregated statistics for the three locations of the rail (P1, P2, P3), which amount to approximately 21,000 location estimates for each method and (N, K) combination.

Figure 5.7 shows the distributions of the localization errors for the classic TDOA and FlexTDOA, aggregated over locations P1, P2, and P3. Because the measurements were acquired in LOS, where the channel between the anchors and the tag is already “ideal,” we do not expect to see a large improvement brought by the channel diversity in the case of FlexTDOA. Indeed, FlexTDOA yields similar errors to the FlexTDOA.

The LOS evaluation sets a baseline for comparison and allows us to verify that FlexTDOA does not *degrade* the accuracy in LOS conditions. A degradation could be possible if, for instance, not all anchors had equally good channels to the tag. As we have previously seen in Figure 5.6, the channel between the *initiating* anchor and the tag is the most important factor for obtaining a high accuracy. Therefore, if anchors with poor channels to the tag (for instance, caused by strong nearby reflectors) become initiators in FlexTDOA, we could expect higher localization errors from FlexTDOA than from the classic TDOA.



Figure 5.8 Setups used to acquire NLOS measurements. The setup in Figure 5.8a corresponds to position P1 from Figure 5.4. It includes one aluminum panel placed as a blocker between anchor A_1 and the tag. Figure 5.8a corresponds to the position P4 from Figure 5.4. It includes two aluminum panels that block the initiating anchor A_1 and various anchors depending on the tag’s position on the rail.

5.4.3 NLOS Propagation

Obstacles between the nodes of a localization system are an integral part of any real-life deployment of a localization system. In this section, we evaluate the performance of the three localization approaches considered so far, based on TWR, the classic TDOA, and FlexTDOA, in NLOS conditions.

We acquired measurements at positions P1 and P4 from Figure 5.4. At P1, we placed a panel covered in aluminum foil between the anchor A_1 and the tag, shown in Figure 5.8a. At P4, we placed two aluminum panels, shown in Figure 5.8b. The panel on the left blocks anchors A_0 and A_9 completely and anchors A_4 , A_7 , A_8 partially, depending on the position of the tag on the rail. The rightmost panel blocks A_1 completely (which is the initiator in the classic TDOA scheme) and A_5 partially.

Figure 5.9 shows the distribution of 3D localization errors aggregated over both locations. The proposed TDOA scheme, FlexTDOA, yields a higher accuracy than the classic TDOA in all NLOS scenarios. With five anchors, the improvement is modest, with a reduction of 5–7 cm in the median and P_{95} errors. However, for seven and ten anchors, FlexTDOA achieves a marked reduction in the P_{95} localization error compared to the classic TDOA, by 19% and 38%, respectively.

As predicted, TWR has the lowest errors in all evaluated scenarios. This




Figure 5.9 Localization errors in NLOS obtained with TWR, classic TDOA (FI-FR), and FlexTDOA (CI-CR). The errors are aggregated over both NLOS positions of the tag (P1 and P4).

is because, in TWR-based localization, obstructions between an anchor and the tag affect only the distance between them. Consequently, the error might affect the location estimate only to a small degree if the rest of the distance measurements have high accuracy.

In contrast, in TDOA localization, an obstruction between the reference anchor and the tag affects *all* the TDOAs from that time slot. Therefore, the location estimate is more prone to biases. This is why FlexTDOA has an advantage over the classic TDOA. If we change the initiating anchor, we minimize the risk that one obstructed anchor will lead to high localization errors at all times. The increased channel diversity brought by using all anchor-tag combinations improves the robustness of the system.

We must remember that, even though TWR-based localization yields the highest localization accuracy, it scales poorly with the number of tags. If we use $N = 9$ anchors, the time slot needs to have a duration of at least 10 ms (according to the time slot structure presented in Section 5.2). Therefore, at most 100 tags could be localized every second using TWR-based localization. Actually, the number could be even smaller given that a certain air time would be needed to synchronize the access of the tags to the channel on the uplink. On the other hand, TDOA localization can scale to an unlimited number of tags. In the same example, each tag could individually obtain 100 locations per second regardless of how many other users are in the area.

5.5 Summary

In [P6], we presented a new flexible scheduling scheme for TDOA localization called FlexTDOA that fully exploits the channel diversity in the environment. We compared FlexTDOA against the classic TDOA implementation and against range-based localization in a deployment of ten anchors and one tag in an office environment. FlexTDOA achieves the lowest localization errors in most scenarios, with and without obstructions. We also evaluated the impact of several factors on the TWR/TDOA and localization accuracy: the order of response, the number of responses, the update rate (or speed) of the tag, and the number of anchors. Our evaluation can be used to choose the optimal configuration of a UWB localization system in future deployments.

FlexTDOA has higher localization accuracy than the classic TDOA in most scenarios, most notably in the presence of obstructions. In LOS, the improvement brought by FlexTDOA in the median localization accuracy compared to the classic TDOA is small, around 2–3 cm. This is not surprising, because the initiator in the classic TDOA already had a good propagation link to the tag. The robustness added by the increased diversity of FlexTDOA is more evident in NLOS. When using obstructions, FlexTDOA achieves up to 38% smaller P_{95} localization errors than the classic TDOA. Overall, FlexTDOA has a median localization error of 13–17 cm in LOS and 15–22 cm in NLOS (the error depends on how many anchors are used).

While the flexible TDOA scheduling scheme decreased the localization error in most scenarios, the channel diversity can be detrimental to the accuracy in cases in which anchors that are obstructed from the tag or from other anchors become initiators. Therefore, in the future, we will propose an anchor selection method that prevents such anchors from becoming initiators. Such a method should adjust to the channel conditions in real time, so that anchors can be both added to *and* removed from the blacklist, and should be flexible enough to adapt to multiple environments (and not, for instance, based on hard-coded thresholds).

In a real deployment, the proposed system needs to be scaled up to a multi-room or building environment. This setup will raise more issues that need to be addressed. First, communication range rarely extends beyond a

room. Therefore, the TDMA scheme needs to be improved so that initiators interrogate only other anchors in their radio range (or, preferably, in LOS). Also, the same time slots should be reused between locations that have no direct radio connections. Second, the tag should select and use in the localization process only the “best” TDOAs. A quality metric should be devised such that TDOAs resulting from poor anchor–tag links are discarded.

A pressing issue in the large-scale adoption of UWB localization systems is the initial calibration and deployment of the anchors. In most anchor-based localization systems, the anchors’ location is assumed to be known with millimeter accuracy, in order to minimize localization errors. In FlexTDOA, because the anchors were in LOS of each other, we could use the pair-wise distances between anchors measured via UWB. We averaged the distances over several minutes such that the mean error of the distance measurements was under 1 cm. Moreover, we used a laser range finder to validate that the measured distances were accurate. However, in a multi-room environment, some of the anchors will be in NLOS of each other. In this scenario, the anchor localization will need to be adapted to a network with densely-connected groups (formed by anchors placed in the same room) but with sparse connections between different groups (where the groups are formed by anchors placed in different rooms).

6 CONCLUSIONS

The thesis explored issues related to the robustness, energy efficiency, accuracy, scalability, and cross-platform compatibility of UWB localization systems. We compared the energy efficiency and accuracy of UWB PHYs in [P1] and [P2]. In [P3], we focused on the cross-platform compatibility of error-mitigation techniques in UWB localization systems. The robustness aspect was covered in all publications, since we investigated the performance of the localization systems in NLOS scenarios in all cases. In addition, in [P4], we proposed a solution for detecting and mitigating NLOS errors in UWB localization systems without manually-collected data. In [P5], we investigated the robustness of BLE RSS measurements under multipath and NLOS propagation. The study can be used as a benchmark to compare BLE and UWB ranging/localization systems or to integrate them in a hybrid UWB-BLE positioning system. In [P6], we proposed a scalable and robust TDOA localization system based on UWB devices. The accuracy aspect was an integral part of all the publications related to UWB, as we investigated the ranging and/or localization accuracy of the proposed systems under multiple conditions.

6.1 Main Results

We summarize the main results and findings of the thesis in view of the research questions (RQs) proposed in Chapter 1.

RQ1: *How do UWB physical interfaces compare in terms of energy efficiency, maximum range, and accuracy?*

The analysis on the energy efficiency and accuracy of HRP and LRP UWB PHYs from [P1] and [P2] revealed that the LRP PHY can enable localization applications with $6.4\times$ less energy consumption than HRP devices in

the short-range mode or with $100\times$ less energy consumption than HRP devices that achieve a similar range. Therefore, the LRP PHY is suitable for ultra-low power devices and wearables. On the other hand, the accuracy and precision of the LRP devices we investigated was slightly lower than the one of HRP devices. For instance, the distance measurements of HRP devices had a standard deviation of up to 2 times lower than the one of LRP devices. Therefore, the HRP PHY might be preferable in applications where ranging/localization accuracy is critical. The studies [P1] and [P2] serve as a benchmark for comparing UWB PHYs and assessing which devices are more suitable for particular applications.

RQ2: *What implications do different UWB hardware architectures have on the cross-platform compatibility of localization systems?*

In [P3], we brought to researchers' attention the issue of cross-device compatibility in error-mitigation techniques for UWB localization systems. To this end, we identified four main challenges standing in the way of platform-independent UWB localization systems.

First, we observed that measurements acquired with different devices under identical environmental conditions have different error statistics. Second, we showed that NLOS errors depend not only on the obstacle and environment, but also on the hardware and its TOA estimation algorithm. Therefore, even measurements acquired with the same device can fluctuate. Third, we showed that CIRs acquired by different platforms at the same locations have different effective lengths, shapes, and statistics. Therefore, error-mitigation techniques based on CIRs are highly device-dependent. Finally, we documented the fact that different vendors provide different types of measurement diagnostics, using incompatible or unspecified units of measurements. Therefore, error-mitigation techniques based on such diagnostics are also device-dependent.

RQ3: *How can we mitigate the impact of NLOS propagation on UWB localization accuracy without manually-collected training data?*

In [P4], we proposed a method to improve the robustness of UWB localization systems in NLOS propagation without the need to collect training data sets. The method arose from a real difficulty in the Author's re-

search: the extremely laborious and time-consuming process of collecting measurements for developing and testing localization algorithms. Therefore, we sought to bypass this process by proposing a method that could learn from measurements acquired in real-time. We did that by observing that NLOS anchors leave artifacts in the localization process that can be used to identify them and mitigate their errors. The proposed method reduced the mean localization error by $2.2\times$ and the standard deviation of location errors by $5.8\times$ compared to the baseline case with no error mitigation.

One drawback of the proposed method is that it needs at least one more anchor than the minimum necessary for localization in order to identify the NLOS anchor(s). However, UWB devices are not expensive: their price is in the range of tens of dollars and will decrease further as the technology becomes more widely deployed. Therefore, having more than the minimum number of anchors in a particular setting should not be an issue. Another limitation is that we did not evaluate the method in a real localization scenario. Therefore, future work should investigate how many measurement samples are necessary to achieve the same improvement in a practical deployment.

RQ4: *How does UWB compare to complementary low-power technologies, such as BLE, in terms of ranging performance?*

Based on the research from [P5], we showed that BLE RSS measurements are highly vulnerable to fluctuations in indoor environments. The variability can be caused by the environment (i.e., the multipath propagation conditions), obstructions, hardware, the advertising channel used, or the co-existence of Wi-Fi and BLE technologies on the same chipset. Therefore, BLE can typically achieve a ranging accuracy on the order of meters, while UWB has a ranging accuracy of several centimeters. However, BLE can complement UWB in localization applications by providing advertising and communication capabilities with an energy consumption up to 3–4.5 times lower than UWB. In addition, since BLE is already widely available in many wearables and sensors, it can facilitate the integration of a joint BLE–UWB localization solution in existing systems.

RQ5: *How can we achieve scalable and robust UWB localization?*

In publication [P6], we proposed a localization system based on UWB devices that implements a flexible TDMA scheme for TDOA localization. We chose to implement a TDOA variant in which the tag is passive, which allows our system to scale with an unlimited number of users, similar to GNSS. Previous to our work, most TDOA localization systems considered only rigid scheduling schemes, with a designated reference anchor and a fixed order of responders. However, we discovered that this approach is not robust in NLOS propagation, especially when the link between the tag and the reference anchor is obstructed. Therefore, we proposed a flexible TDOA scheme in which every anchor can become a reference and the order of responses can also change. We evaluated multiple system parameters and the impact brought by the improved channel diversity of the proposed TDMA scheme.

We compared the proposed TDOA system to the classic TDOA approach and range-based localization. The improved robustness of the proposed method was evident in the experiments performed in NLOS conditions. The proposed TDOA method outperformed the classic TDOA in all experiments, reducing the 95% localization error with up to 16 cm. The proposed TDOA method even outperformed range-based localization with five and seven anchors, as it decreased the localization error with up to 10 cm. With ten anchors, range-based localization slightly outperformed the proposed TDOA method, the two algorithms achieving a 95% error of 29.4 cm and 32.4 cm, respectively. However, unlike TDOA, range-based localization does not scale with an unlimited number of tags. Therefore, we showcased a new TDOA localization system with unlimited scalability and improved robustness in most scenarios compared to state-of-the-art localization algorithms.

One possible drawback of the proposed TDOA method is that it might not achieve the best localization accuracy in case several anchors are obstructed or have poor links to the tag. In this case, the localization accuracy might be lower than, for instance, in the case in which there is a single initiating anchor with a good link to the tag. This is because, in the proposed method, all anchors become initiators and the resulting localization accuracy is averaged over all possible anchor combinations. One solution is to devise a measure-

ment selection procedure in which TDOAs with a poor quality factor are not used in the localization process. This could be achieved, for instance, by tracking the standard deviation of the TDOA measurements over time.

6.2 Future Research Directions

While UWB technology has been extensively researched during the past 20 years, localization systems based on UWB communication have not yet reached their maturity. As such, there are still pressing issues in UWB localization systems that need to be addressed.

The adoption of UWB devices for building-wide public deployments is still impeded by certain factors. The first is related to the presence of UWB chipsets in user devices. Although UWB chipsets have been introduced in smartphones since 2019 [74, 81, 93], mobile applications using UWB have not followed the hardware adoption at the same rate. One reason is that the link budget of UWB communication still has to be improved. This is an enhancement currently targeted by the IEEE 802.15.4 Task Group 4ab [46]. In addition, the energy (in)efficiency of UWB devices could prevent their adoption in small wearables. As we saw in Chapter 4, UWB HRP devices consume $3\text{--}4.5\times$ more energy per ranging measurement than BLE devices. In comparison, UWB LRP devices have an energy consumption of the same order of magnitude as BLE devices but are not as widely-adopted as their HRP counterparts.

Second, the deployment of anchors on top of existing infrastructure is difficult. Anchors are frequently assumed to have fixed and known locations in order to locate the tag with maximum accuracy. In a single room, it is relatively easy to obtain the anchors' locations with millimeter-level accuracy using common tools such as a laser range finder. In [P6], we relied on UWB distances averaged over several minutes to compute the anchors' locations. We validated that the anchors' locations computed in this way were within several millimeters of the true locations. However, this procedure was facilitated by the fact that the anchors were in LOS of each other. In a multi-room deployment, the anchors would be placed, most likely, in room-wide clusters and anchors from different clusters will be in NLOS. Therefore, future

research should investigate how to simplify the deployment of infrastructure necessary for UWB localization.

Third, and perhaps the most important, UWB localization systems have not been extensively evaluated in multi-room deployments besides work done in [31]. In such scenarios, scalable TDMA schemes that cover anchor and tag transmissions need to be implemented. We believe that the FlexTDOA system from [P6] provides a solid base for extending location services to building-wide areas. In the future, the FlexTDOA system will also be evaluated in multi-room scenarios. The localization accuracy in transition areas between clusters of anchors also needs to be evaluated, as location estimates usually degrade outside the convex hull formed by the anchors [98].

The robustness of UWB localization systems in NLOS propagation is still far from being fully addressed. As we highlighted in [P3], ML techniques can show quick and easy improvements of the localization accuracy, but the results often do not extend past the environment in which the model was trained and tested [7]. Work in adjacent fields has shown that the performance of models often generalizes better when training on a combination of synthetic and real datasets instead of relying on real measurements alone [95, 125]. In [P4], we proposed a framework that learns based on data acquired in real-time, but more work is needed to see how much data is needed to obtain the same results in real deployments and how often the models would need to be updated with new data.

Most existing ML approaches for improving the resilience of localization systems in NLOS propagation suffer from cross-device compatibility issues, as shown in [P3]. Therefore, more work is needed to ensure that robustness-enhancing methods can yield similar results across different devices.

Finally, integrated sensing and communication (ISAC) is a hot topic in WiFi, 5G, and 6G communications [63]. Recent works are investigating ICAS for UWB devices as well [25] and integrating the concept in real-life applications, e.g., detecting car occupancy using UWB infrastructure [67]. As UWB devices become more present in our day-to-day lives, the possibilities for integrated sensing and communication—or, rather, integrated sensing and localization—should also expand to enable surrounding devices to better react to our presence and movement.

REFERENCES

- [1] 3db Access AG. *3DB6830D User Guide. Version 4.0*. Dec. 2018.
- [2] *APS017 Application Note. Maximising Range in DW1000 Based Systems. Version 1.1*. www.qorvo.com/products/p/DW1000. Accessed on 29-11-2022.
- [3] Fernando J Aranda, Felipe Parralejo, Fernando J Álvarez, and José A Paredes. "Performance analysis of fingerprinting indoor positioning methods with BLE". In: *Expert Systems with Applications* 202 (2022), p. 117095.
- [4] Huseyin Arslan, Zhi Ning Chen, and Maria-Gabriella Di Benedetto. *Ultra wideband wireless communication*. John Wiley & Sons, 2006.
- [5] Hyoin Bae, Jaesung Oh, KangKyu Lee, and Jun-Ho Oh. "Low-cost indoor positioning system using BLE (bluetooth low energy) based sensor fusion with constrained extended Kalman Filter". In: *2016 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE. 2016, pp. 939–945.
- [6] Luca Barbieri, Mattia Brambilla, Andrea Trabattoni, Stefano Mervic, and Monica Nicoli. "UWB localization in a smart factory: Augmentation methods and experimental assessment". In: *IEEE Transactions on Instrumentation and Measurement* 70 (2021), pp. 1–18.
- [7] Valentín Barral, Carlos J Escudero, José A García-Naya, and Pedro Suárez-Casal. "Environmental cross-validation of NLOS machine learning classification/mitigation with low-cost UWB positioning systems". In: *Sensors* 19.24 (2019), p. 5438.
- [8] Andreas Biri, Neal Jackson, Lothar Thiele, Pat Pannuto, and Prabal Dutta. "SociTrack: Infrastructure-free interaction tracking through mo-

- ble sensor networks". In: *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 2020, pp. 1–14.
- [9] James J Caffery. "A new approach to the geometry of TOA location". In: *Vehicular Technology Conference Fall 2000. IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference (Cat. No. 00CH37152)*. Vol. 4. IEEE. 2000, pp. 1943–1949.
- [10] Manuel Castillo-Cara, Jesús Lovón-Melgarejo, Gusseppe Bravo-Rocca, Luis Orozco-Barbosa, and Ismael García-Varea. "An empirical study of the transmission power setting for bluetooth-based indoor localization mechanisms". In: *Sensors* 17.6 (2017), p. 1318.
- [11] Long Cheng, Yifan Li, Yan Wang, Yangyang Bi, Liang Feng, and Mingkun Xue. "A triple-filter NLOS localization algorithm based on fuzzy c-means for wireless sensor networks". In: *Sensors* 19.5 (2019), p. 1215.
- [12] Federal Communications Commission et al. "Revision of part 15 of the commission's rules regarding ultra-wideband transmission systems". In: *First Report and Order, FCC 02-48* (2002).
- [13] Car Connectivity Consortium. *Digital Key Release 3 v1.1.0 Specification*. 2022. URL: <https://carconnectivity.org/press-release/car-connectivity-consortium-makes-digital-key-release-3-v1-1-specification-available-to-the-public/>.
- [14] FiRa Consortium. 2022. URL: firaconsortium.org.
- [15] Dieter Coppens, Eli De Poorter, Adnan Shahid, Sam Lemey, Ben Van Herbruggen, and Chris Marshall. "An Overview of UWB Standards and Organizations (IEEE 802.15. 4, FiRa, Apple): Interoperability Aspects and Future Research Directions". In: *IEEE Access* (2022).
- [16] Pablo Corbalán, Gian Pietro Picco, and Sameera Palipana. "Chorus: UWB concurrent transmissions for GPS-like passive localization of countless targets". In: *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE. 2019, pp. 133–144.

- [17] Gabriel De Blasio, Alexis Quesada-Arencibia, Carmelo R García, Jezabel Miriam Molina-Gil, and Cándido Caballero-Gil. "Study on an indoor positioning system for harsh environments based on Wi-Fi and bluetooth low energy". In: *Sensors* 17.6 (2017), p. 1299.
- [18] Mickael Delamare, Remi Bouteau, Xavier Savatier, and Nicolas Iriart. "Static and dynamic evaluation of an UWB localization system for industrial applications". In: *Sci* 1.3 (2019), p. 62.
- [19] F. Della Rosa, T. Paakki, J. Nurmi, M. Pelosi, and G. D. Rosa. "Hand-grip impact on range-based cooperative positioning". In: *2014 11th International Symposium on Wireless Communications Systems (ISWCS)*. 2014, pp. 728–732. doi: 10.1109/ISWCS.2014.6933449.
- [20] Igor Dotlic, Andrew Connell, and Michael McLaughlin. "Ranging methods utilizing carrier frequency offset estimation". In: *2018 15th Workshop on Positioning, Navigation and Communications (WPNC)*. IEEE. 2018, pp. 1–6.
- [21] *DW1000 Datasheet. Version 2.09*. Decawave Ltd.
- [22] *DW1000 User Manual. Version 2.11*. Decawave Ltd.
- [23] Zohreh Ebadi, Cédric Hannotier, Heidi Steendam, François Horlin, and François Quitin. "An over-the-air CFO-assisted synchronization algorithm for TDOA-based localization systems". In: *2020 IEEE 92nd Vehicular Technology Conference (VTC2020-Fall)*. IEEE. 2020, pp. 1–5.
- [24] Jiancun Fan and Ahsan Saleem Awan. "Non-line-of-sight identification based on unsupervised machine learning in ultra wideband systems". In: *IEEE Access* 7 (2019), pp. 32464–32471.
- [25] Zhongyuan Fang, Wensong Wang, Jipeng Wang, Bingqiang Liu, Kai Tang, Liheng Lou, Chun-Huat Heng, Chao Wang, and Yuanjin Zheng. "Integrated Wideband Chip-Scale RF Transceivers for Radar Sensing and UWB Communications: A Survey". In: *IEEE Circuits and Systems Magazine* 22.1 (2022), pp. 40–76.
- [26] Ramsey Faragher and Robert Harle. "Location fingerprinting with bluetooth low energy beacons". In: *IEEE journal on Selected Areas in Communications* 33.11 (2015), pp. 2418–2428.

- [27] Tobias Feigl, Ernst Eberlein, Sebastian Kram, and Christopher Mutschler. “Robust ToA-Estimation using Convolutional Neural Networks on Randomized Channel Models”. In: *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2020, pp. 1–8.
- [28] Laura Flueratoru, Elena Simona Lohan, Jari Nurmi, and Dragos Niculescu. “HTC Vive as a Ground-Truth System for Anchor-Based Indoor Localization”. In: *2020 12th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. IEEE. 2020, pp. 214–221.
- [29] Jaron Fontaine, Matteo Ridolfi, Ben Van Herbruggen, Adnan Shahid, and Eli De Poorter. “Edge inference for UWB ranging error correction using autoencoders”. In: *IEEE Access* 8 (2020), pp. 139143–139155.
- [30] Benoît Frénay and Michel Verleysen. “Classification in the presence of label noise: a survey”. In: *IEEE transactions on neural networks and learning systems* 25.5 (2013), pp. 845–869.
- [31] Johannes Friedrich, Janis Tiemann, and Christian Wietfeld. “Accurate multi-Zone UWB TDOA localization utilizing cascaded wireless clock synchronization”. In: *2021 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2021, pp. 1–8.
- [32] G. De Blasio, A. Quesada-Arencibia, C. R. García, J. C. Rodríguez-Rodríguez, and R. Moreno-Díaz. “A Protocol-Channel-Based Indoor Positioning Performance Study for Bluetooth Low Energy”. In: *IEEE Access* 6 (2018), pp. 33440–33450. doi: 10.1109/ACCESS.2018.2837497.
- [33] Christian Gentner, Daniel Günther, and Philipp H. Kindt. *Identifying the BLE Advertising Channel for Reliable Distance Estimation on Smartphones*. 2020. arXiv: 2006.09099 [cs.NI].
- [34] Sinan Gezici and H Vincent Poor. “Position estimation via ultra-wide-band signals”. In: *Proceedings of the IEEE* 97.2 (2009), pp. 386–403.
- [35] Bernhard Großwindhager, Michael Stocker, Michael Rath, Carlo Alberto Boano, and Kay Römer. “SnapLoc: An ultra-fast UWB-based indoor localization system for an unlimited number of tags”. In: *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*. 2019, pp. 61–72.

- [36] Fredrik Gustafsson and Fredrik Gunnarsson. "Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements". In: *IEEE Signal processing magazine* 22.4 (2005), pp. 41–53.
- [37] İsmail Güvenç, Chia-Chin Chong, Fujio Watanabe, and Hiroshi Inamura. "NLOS identification and weighted least-squares localization for UWB systems using multipath channel statistics". In: *EURASIP Journal on Advances in Signal Processing* 2008.1 (2007), p. 271984.
- [38] Michael Hamer and Raffaello D'Andrea. "Self-calibrating ultra-wideband network supporting multi-robot localization". In: *IEEE Access* 6 (2018), pp. 22292–22304.
- [39] Ulrich Hammes and Abdelhak M Zoubir. "Robust MT tracking based on M-estimation and interacting multiple model algorithm". In: *IEEE Transactions on Signal Processing* 59.7 (2011), pp. 3398–3409.
- [40] P. Handayani, L. Mubarokah, and G. Hendrantoro. "Pathloss and shadowing characteristics in indoor environment at 2.4 GHz band". In: *2015 International Seminar on Intelligent Technology and Its Applications (ISITIA)*. 2015, pp. 423–428. doi: 10.1109/ISITIA.2015.7220018.
- [41] Z. He, Y. Li, L. Pei, R. Chen, and N. El-Sheimy. "Calibrating Multi-Channel RSS Observations for Localization Using Gaussian Process". In: *IEEE Wireless Communications Letters* 8.4 (2019), pp. 1116–1119. doi: 10.1109/LWC.2019.2908397.
- [42] Robert Heyn, Marc Kuhn, Henry Schulten, Gregor Dumphart, Janick Zwyssig, Florian Trosch, and Armin Wittneben. "User tracking for access control with bluetooth low energy". In: *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*. IEEE. 2019, pp. 1–7.
- [43] Robert Heyn and Armin Wittneben. "Comprehensive Measurement-Based Evaluation of Posture Detection from Ultra Low Power UWB Signals". In: *2021 IEEE 32nd Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE. 2021, pp. 1518–1524.

- [44] Robert Heyn and Armin Wittneben. “WBAN Node Topologies for Reliable Posture Detection from On-Body UWB RSS Measurements”. In: *ICC 2022-IEEE International Conference on Communications*. IEEE. 2022, pp. 5298–5303.
- [45] Jiyan Huang and Qun Wan. “Analysis of TDOA and TDOA/SS based geolocation techniques in a non-line-of-sight environment”. In: *Journal of Communications and Networks* 14.5 (2012), pp. 533–539.
- [46] *IEEE 802.15 WSN™ Task Group 4ab (TG4ab). 802.15.4 UWB Next Generation*. <https://www.ieee802.org/15/pub/TG4ab.html>. Accessed on 29-11-2022.
- [47] “IEEE Standard for Low-Rate Wireless Networks”. In: *IEEE Std 802.15.4-2020 (Revision of IEEE Std 802.15.4-2015)* (2020), pp. 1–800. DOI: 10.1109/IEEESTD.2020.9144691.
- [48] *Imec Showcases World’s First Sub-5mW, IEEE 802.15.4z Ultra-Wideband Transmitter Chip*. <https://www.imec-int.com/en/press/imec-showcases-worlds-first-sub-5mw-ieee-802154z-ultra-wideband-transmitter-chip>. Accessed on 29 August 2022.
- [49] Apple Inc. *Apple Nearby Interaction protocol specification*. 2020.
- [50] Antonio Ramón Jiménez and Fernando Seco. “Comparing Decawave and Bespoon UWB location systems: Indoor/outdoor performance analysis”. In: *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2016, pp. 1–8.
- [51] Radovan Juran, Petr Mlynek, Martin Stusek, Pavel Masek, Michal Mikulasek, and Aleksandr Ometov. “Hands-On Experience with UWB: Angle of Arrival Accuracy Evaluation in Channel 9”. In: *2022 14th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. IEEE. 2022, pp. 45–49.
- [52] Arief Affendi Juri, Tughrul Arslan, and Fengzhou Wang. “Obstruction-aware bluetooth low energy indoor positioning”. In: *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*. 2016, pp. 2254–2261.

- [53] Avinash Kalyanaraman, Yunze Zeng, Sushanta Rakshit, and Vivek Jain. “Caraokey: Car states sensing via the ultra-wideband keyless infrastructure”. In: *2020 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE. 2020, pp. 1–9.
- [54] Ashith Kumar, Zhuo Li, Qilian Liang, Baoju Zhang, and Xiaorong Wu. “Experimental study of through-wall human detection using ultra wide-band radar sensors”. In: *Measurement* 47 (2014), pp. 869–879.
- [55] Flueratoru Laura, Lohan Elena Simona, and Niculescu Dragoş. *Dataset: Ultra-Wideband Ranging Measurements Acquired With Three Different Platforms (Qorvo, TDSR, 3db Access)*. DOI: 10.5281/zenodo.6984698.
- [56] Flueratoru Laura, Lohan Elena Simona, and Niculescu Dragoş. *UWB Ranging and Localization Dataset for “High-Accuracy Ranging and Localization with Ultra-Wideband Communication for Energy-Constrained Devices”*. 2021. DOI: 10.5281/zenodo.4686379.
- [57] Flueratoru Laura, Shubina Viktoriia, Lohan Elena Simona, and Niculescu Dragoş. *Open-Access Data for Received Signal Strength Measurements with BLE Signals for Contact Tracing and Proximity Detection*. DOI: 10.5281/zenodo.4643668.
- [58] Anton Ledergerber, Michael Hamer, and Raffaello D’Andrea. “A robot self-localization system using one-way ultra-wideband communication”. In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2015, pp. 3131–3137.
- [59] Douglas J. Leith and Stephen Farrell. *Coronavirus Contact Tracing: Evaluating The Potential Of Using Bluetooth Received Signal Strength For Proximity Detection*. 2020. arXiv: 2006.06822 [eess.SP].
- [60] Erik Leitinger, Paul Meissner, Christoph Rüdissler, Gregor Dumphart, and Klaus Witrisal. “Evaluation of position-related information in multipath components for indoor positioning”. In: *IEEE JSAC* 33.11 (2015), pp. 2313–2328.
- [61] Patrick Leu, Giovanni Camurati, Alexander Heinrich, Marc Roeschlin, Claudio Anliker, Matthias Hollick, Srdjan Capkun, and Jiska Classen. “Ghost Peak: Practical Distance Reduction Attacks Against {HRP}{UWB}

- Ranging". In: *31st USENIX Security Symposium (USENIX Security 22)*. 2022, pp. 1343–1359.
- [62] Shenghong Li, Mark Hedley, Iain B Collings, and David Humphrey. "Joint trajectory and ranging offset estimation for accurate tracking in NLOS environments". In: *IEEE Transactions on Aerospace and Electronic Systems* 56.1 (2019), pp. 3–14.
- [63] An Liu, Zhe Huang, Min Li, Yubo Wan, Wenrui Li, Tony Xiao Han, Chenchen Liu, Rui Du, Danny Kai Pin Tan, Jianmin Lu, et al. "A survey on fundamental limits of integrated sensing and communication". In: *IEEE Communications Surveys & Tutorials* 24.2 (2022), pp. 994–1034.
- [64] TDSR LLC. *TDSR RangeNet Application Programming Interface (API) Specification. Version 320-0313I*. Oct. 2020.
- [65] E. S. Lohan, J. Talvitie, P. Figueiredo e Silva, H. Nurminen, S. Ali-Löytty, and R. Piché. "Received signal strength models for WLAN and BLE-based indoor positioning in multi-floor buildings". In: *2015 International Conference on Localization and GNSS (ICL-GNSS)*. 2015, pp. 1–6. doi: 10.1109/ICL-GNSS.2015.7217154.
- [66] Decawave Ltd. *DW3000 User Manual. Version 1.1*. May 2021.
- [67] Yongsen Ma, Yunze Zeng, and Vivek Jain. "CarOSense: Car occupancy sensing with the ultra-wideband keyless infrastructure". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4.3 (2020), pp. 1–28.
- [68] G MacGougan, K O'Keefe, and R Klukas. "UWB ranging and ranging measurement accuracy". In: *Measurement Science and Technology* 20.9 (2009), pp. 0957–0233.
- [69] Stefano Marano, Wesley M Gifford, Henk Wymeersch, and Moe Z Win. "NLOS identification and mitigation for localization based on UWB experimental data". In: *IEEE JSAC* 28.7 (2010), pp. 1026–1035.
- [70] Nasreen Mohsin, Shahram Payandeh, Derek Ho, and Jean Pierre Gelinass. "Study of Activity Tracking through Bluetooth Low Energy-Based Network". In: *Journal of Sensors* 2019 (Feb. 2019), p. 6876925. ISSN: 1687-725X. doi: 10.1155/2019/6876925.

- [71] Stefania Monica and Gianluigi Ferrari. "An experimental model for UWB distance measurements and its application to localization problems". In: *2014 IEEE International Conference on Ultra-WideBand (ICUWB)*. IEEE. 2014, pp. 297–302.
- [72] Philipp Müller, Henk Wymeersch, and Robert Piché. "UWB positioning with generalized Gaussian mixture filters". In: *IEEE Transactions on Mobile Computing* 13.10 (2014), pp. 2406–2414.
- [73] Sharareh Naghdi and Kyle O’Keefe. "Detecting and correcting for human obstacles in BLE trilateration using artificial intelligence". In: *Sensors* 20.5 (2020), p. 1350.
- [74] *Nearby Interaction with UWB*. <https://developer.apple.com/nearby-interaction/>. Accessed on 29-11-2022.
- [75] J. Neburka, Z. Tlamsa, V. Benes, L. Polak, O. Kaller, L. Bolecek, J. Sebesta, and T. Kratochvil. "Study of the performance of RSSI based Bluetooth Smart indoor positioning". In: *2016 26th International Conference Radioelektronika (RADIOELEKTRONIKA)*. 2016, pp. 121–125. DOI: 10.1109/RADIOELEK.2016.7477344.
- [76] Dries Neiryneck, Eric Luk, and Michael McLaughlin. "An alternative double-sided two-way ranging method". In: *2016 13th workshop on positioning, navigation and communications (WPNC)*. IEEE. 2016, pp. 1–4.
- [77] Quang Huy Nguyen, Princy Johnson, Trung Thanh Nguyen, and Martin Randles. "Optimized indoor positioning for static mode smart devices using BLE". In: *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. IEEE. 2017, pp. 1–6.
- [78] Vlad Niculescu, Daniele Palossi, Michele Magno, and Luca Benini. "Energy-efficient, Precise UWB-based 3-D Localization of Sensor Nodes with a Nano-UAV". In: *IEEE Internet of Things Journal* (2022).
- [79] Ali Nikoukar, Mansour Abboud, Borna Samadi, Mesut Güneş, and Behnam Dezfouli. "Empirical analysis and modeling of Bluetooth low-energy (BLE) advertisement channels". In: *2018 17th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*. IEEE. 2018, pp. 1–6.

- [80] *nRF52840: Multiprotocol Bluetooth 5.3 SoC supporting Bluetooth Low Energy, Bluetooth mesh, NFC, Thread and Zigbee*. <https://www.nordicsemi.com/Products/nRF52840>. Accessed on 29-11-2022.
- [81] *NXP's Trimension™ Ultra-Wideband Technology Helps Samsung Users Easily Find Their Misplaced Belongings*. [https://www.nxp.com/docs/en/supporting-information/NXPs-Trimension™-Ultra-Wideband-Technology-Helps-Samsung-Users-Easily-Find-Their-M.pdf](https://www.nxp.com/docs/en/supporting-information/NXPs-TrimensionTM-Ultra-Wideband-Technology-Helps-Samsung-Users-Easily-Find-Their-M.pdf). Accessed on 29-11-2022.
- [82] Serban Georgica Obreja and Alexandru Vulpe. "Evaluation of an indoor localization solution based on bluetooth low energy beacons". In: *2020 13th International Conference on Communications (COMM)*. IEEE. 2020, pp. 227–231.
- [83] Aleksandr Ometov, Viktoriia Shubina, Lucie Klus, Justyna Skibińska, Salwa Saafi, Pavel Pascacio, Laura Flueratoru, Darwin Quezada Gai-bor, Nadezhda Chukhno, Olga Chukhno, et al. "A survey on wearable technology: History, state-of-the-art and current challenges". In: *Computer Networks* 193 (2021), p. 108074.
- [84] Jeongyeup Paek, JeongGil Ko, and Hyungsik Shin. "A measurement study of BLE iBeacon and geometric adjustment scheme for indoor location-based mobile applications". In: *Mobile Information Systems* 2016 (2016).
- [85] JiWoong Park, SungChan Nam, HongBeom Choi, YoungEun Ko, and Young-Bae Ko. "Improving Deep Learning-Based UWB LOS/NLOS Identification with Transfer Learning: An Empirical Approach". In: *Electronics* 9.10 (2020), p. 1714.
- [86] Pavel Pascacio, Joaquín Torres-Sospedra, Sven Casteleyn, and Elena Simona Lohan. "A Collaborative Approach Using Neural Networks for BLE-RSS Lateration-Based Indoor Positioning". In: *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE. 2022, pp. 01–09.
- [87] Pavel Pascacio, Joaquín Torres-Sospedra, and Sven Casteleyn. "A Lateration Method based on Effective Combinatorial Beacon Selection for Bluetooth Low Energy Indoor Positioning". In: *2021 17th International*

- Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. IEEE. 2021, pp. 397–402.
- [88] Mathias Pelka and Horst Hellbrück. “S-TDoA—Sequential time difference of arrival—A scalable and synchronization free approach for Positioning”. In: *2016 IEEE Wireless Communications and Networking Conference*. IEEE. 2016, pp. 1–6.
- [89] Yiran Peng, Wentao Fan, Xin Dong, and Xing Zhang. “An iterative weighted KNN (IW-KNN) based indoor localization method in bluetooth low energy (BLE) environment”. In: *2016 Intl IEEE Conferences on Ubiquitous Intelligence & Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress (UIC/ATC/ScalCom/CB-DCom/IoP/SmartWorld)*. IEEE. 2016, pp. 794–800.
- [90] John G Proakis and Masoud Salehi. *Digital communications*. Vol. 4. McGraw-hill New York, 2001.
- [91] Amanda Prorok, Phillip Tomé, and Alcherio Martinoli. “Accommodation of NLOS for ultra-wideband TDOA localization in single-and multi-robot systems”. In: *2011 international conference on indoor positioning and indoor navigation*. IEEE. 2011, pp. 1–9.
- [92] Qorvo® *Completes Acquisition of Decawave*. <https://www.qorvo.com/newsroom/news/2020/qorvo-completes-acquisition-of-decawave>. Accessed: 25 November 2022.
- [93] Qorvo® *Delivers Ultra-Wideband in Google®Pixel 6 Pro*. <https://www.qorvo.com/newsroom/news/2021/qorvo-delivers-ultra-wideband-in-google-pixel-6-pro>. Accessed on 29-11-2022.
- [94] Tony QS Quek and Moe Z Win. “Analysis of UWB transmitted-reference communication systems in dense multipath channels”. In: *IEEE Journal on Selected Areas in Communications* 23.9 (2005), pp. 1863–1874.
- [95] Darwin Quezada-Gaibor, Joaquín Torres-Sospedra, Jari Nurmi, Yevgeni Koucheryavy, and Joaquín Huerta. “SURIMI: Supervised Radio Map Augmentation with Deep Learning and a Generative Adversarial Network for Fingerprint-based Indoor Positioning”. In: *2022 IEEE*

- 12th International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2022, pp. 1–8.
- [96] U. M. Qureshi, Z. Umair, Y. Duan, and G. P. Hancke. “Analysis of Bluetooth Low Energy (BLE) Based Indoor Localization System with Multiple Transmission Power Levels”. In: *2018 IEEE 27th International Symposium on Industrial Electronics (ISIE)*. 2018, pp. 1302–1307. DOI: 10.1109/ISIE.2018.8433787.
- [97] Antonio Ramón Jiménez Ruiz and Fernando Seco Granja. “Comparing Ubisense, Bespoon, and Decawave UWB location systems: Indoor performance analysis”. In: *IEEE Transactions on instrumentation and Measurement* 66.8 (2017), pp. 2106–2117.
- [98] Zafer Sahinoglu, Sinan Gezici, and Ismail Guvenc. “Ultra-wideband Positioning Systems: Theoretical Limits, Ranging Algorithms, and Protocols”. In: *Cambridge, New York* (2008).
- [99] M. Sasaki, T. Nakahira, K. Wakao, and T. Moriyama. “Human Blockage Loss Characteristics of 5 GHz Wi-Fi Band in A Crowded Stadium”. In: *IEEE Antennas and Wireless Propagation Letters* (2021), pp. 1–1. DOI: 10.1109/LAWP.2021.3069004.
- [100] Anthony Schenck, Edwin Walsh, Jonas Reijnders, Ted Ooijevaar, Risang Yudanto, Erik Hostens, Walter Daems, and Jan Steckel. “Information Theoretic Framework for the Optimization of UWB Localization Systems”. In: *2018 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. ISSN: 2471-917X. Sept. 2018, pp. 1–8. DOI: 10.1109/IPIN.2018.8533802.
- [101] Lorenz Schmid, David Salido-Monzú, and Andreas Wieser. “Accuracy assessment and learned error mitigation of UWB ToF ranging”. In: *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2019, pp. 1–8.
- [102] Jens Schroeder, Stefan Galler, Kyandoghene Kyamakya, and Klaus Jobmann. “NLOS detection algorithms for ultra-wideband localization”. In: *2007 4th Workshop on Positioning, Navigation and Communication*. IEEE. 2007, pp. 159–166.

- [103] Maximilian Schuh, Hannah Brunner, Michael Stocker, Markus Schuß, Carlo Alberto Boano, and Kay Römer. “First Steps in Benchmarking the Performance of Heterogeneous Ultra-Wideband Platforms”. In: *2022 Workshop on Benchmarking Cyber-Physical Systems and Internet of Things (CPS-IoTBench)*. IEEE. 2022, pp. 34–39.
- [104] Henry Schulten, Marc Kuhn, Robert Heyn, Gregor Dumphart, Florian Trosch, and Armin Wittneben. “On the crucial impact of antennas and diversity on BLE RSSI-based indoor localization”. In: *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*. IEEE. 2019, pp. 1–6.
- [105] Arash Shahi, Afrooz Aryan, Jeffrey S West, Carl T Haas, and Ralph CG Haas. “Deterioration of UWB positioning during construction”. In: *Automation in Construction* 24 (2012), pp. 72–80.
- [106] Xuemin Shen, Mohsen Guizani, Robert Caiming Qiu, and Tho Le-Ngoc. *Ultra-wideband wireless communications and networks*. John Wiley & Sons, 2007.
- [107] S. Shrestha, J. Talvitie, and E. S. Lohan. “Deconvolution-based indoor localization with WLAN signals and unknown access point locations”. In: *2013 International Conference on Localization and GNSS (ICL-GNSS)*. 2013, pp. 1–6. doi: 10.1109/ICL-GNSS.2013.6577256.
- [108] Juri Sidorenko, Volker Schatz, Norbert Scherer-Negenborn, Michael Arens, and Urs Hugentobler. “Error corrections for ultrawideband ranging”. In: *IEEE Transactions on Instrumentation and Measurement* 69.11 (2020), pp. 9037–9047.
- [109] Bruno J Silva and Gerhard P Hancke. “Non-Line-of-Sight Identification Without Channel Statistics”. In: *IECON 2020 The 46th Annual Conference of the IEEE Industrial Electronics Society*. IEEE. 2020, pp. 4489–4493.
- [110] Bruno J. Silva and Gerhard Petrus Hancke. “Ranging Error Mitigation for Through-the-Wall Non-Line-of-Sight Conditions”. In: *IEEE Transactions on Industrial Informatics* (2020). Conference Name: IEEE Transactions on Industrial Informatics, pp. 1–1. ISSN: 1941-0050. doi: 10.1109/TII.2020.2969886.

- [111] Mridula Singh, Patrick Leu, and Srdjan Capkun. “UWB with pulse reordering: Securing ranging against relay and physical-layer attacks”. In: *Cryptology ePrint Archive* (2017).
- [112] Mridula Singh, Marc Roeschlin, Ezzat Zalzal, Patrick Leu, and Srdjan Čapkun. “Security analysis of IEEE 802.15. 4z/HRP UWB time-of-flight distance measurement”. In: *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 2021, pp. 227–237.
- [113] Niilo Sirola. “Closed-form algorithms in mobile positioning: Myths and misconceptions”. In: *2010 7th Workshop on Positioning, Navigation and Communication*. IEEE. 2010, pp. 38–44.
- [114] Julius Smith and Jonathan Abel. “Closed-form least-squares source location estimation from range-difference measurements”. In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 35.12 (1987), pp. 1661–1669.
- [115] Yang Song. *Investigation on the energy consumption of wireless RPM sensor*. 2020.
- [116] Michael Stocker, Bernhard Großwindhager, Carlo Alberto Boano, and Kay Römer. “Towards secure and scalable UWB-based positioning systems”. In: *2020 IEEE 17th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*. IEEE. 2020, pp. 247–255.
- [117] Qinglin Tian, Kevin I-Kai Wang, and Zoran Salcic. “Human Body Shadowing Effect on UWB-Based Ranging System for Pedestrian Tracking”. In: *IEEE Transactions on Instrumentation and Measurement* 68.10 (Oct. 2019). Conference Name: IEEE Transactions on Instrumentation and Measurement, pp. 4028–4037. ISSN: 1557-9662. DOI: 10.1109/TIM.2018.2884605.
- [118] Janis Tiemann, Fabian Eckermann, and Christian Wietfeld. “Atlas-an open-source tdoa-based ultra-wideband localization system”. In: *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2016, pp. 1–6.

- [119] Janis Tiemann, Yehya Elmasry, Lucas Koring, and Christian Wietfeld. "ATLAS FaST: Fast and simple scheduled TDOA for reliable ultra-wideband localization". In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE. 2019, pp. 2554–2560.
- [120] Janis Tiemann and Christian Wietfeld. "Scalable and precise multi-UAV indoor navigation using TDOA-based UWB localization". In: *2017 international conference on indoor positioning and indoor navigation (IPIN)*. IEEE. 2017, pp. 1–7.
- [121] Andreu Urruela, Josep Sala, and Jaume Riba. "Average performance analysis of circular and hyperbolic geolocation". In: *IEEE Transactions on Vehicular Technology* 55.1 (2006), pp. 52–66.
- [122] Reza Monir Vaghefi, Javier Schloemann, and R Michael Buehrer. "NLOS mitigation in TOA-based localization using semidefinite programming". In: *2013 10th Workshop on Positioning, Navigation and Communication (WPNC)*. IEEE. 2013, pp. 1–6.
- [123] Ben Van Herbruggen, Jaron Fontaine, and Eli De Poorter. "Anchor pair selection for error correction in Time Difference of Arrival (TDoA) Ultra Wideband (UWB) positioning systems". In: *2021 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2021, pp. 1–8.
- [124] Davide Vecchia, Pablo Corbalán, Timofei Istomin, and Gian Pietro Picco. "TALLA: Large-scale TDoA localization with ultra-wideband radios". In: *2019 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE. 2019, pp. 1–8.
- [125] Anran Wang, Maruchi Kim, Hao Zhang, and Shyamnath Gollakota. "Hybrid neural networks for on-device directional hearing". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 10. 2022, pp. 11421–11430.
- [126] Gang Wang, Weichen Zhu, and Nirwan Ansari. "Robust TDOA-based localization for IoT via joint source position and NLOS error estimation". In: *IEEE Internet of Things Journal* 6.5 (2019), pp. 8529–8541.

- [127] Jia Wang, Asad Khalid Raja, and Zhibo Pang. "Prototyping and experimental comparison of IR-UWB based high precision localization technologies". In: *UIC-ATC-ScalCom*. IEEE. 2015, pp. 1187–1192.
- [128] Q. Wang, D. W. Matolak, and B. Ai. "Shadowing Characterization for 5-GHz Vehicle-to-Vehicle Channels". In: *IEEE Transactions on Vehicular Technology* 67.3 (2018), pp. 1855–1866. doi: 10.1109/TVT.2017.2764267.
- [129] Tianyu Wang, Keke Hu, Zhihang Li, Kangbo Lin, Jian Wang, and Yuan Shen. "A Semi-Supervised Learning Approach for UWB Ranging Error Mitigation". In: *IEEE Wireless Communications Letters* (2020).
- [130] Martin Weisenhorn and Walter Hirt. "Robust noncoherent receiver exploiting UWB channel properties". In: *2004 International Workshop on Ultra Wideband Systems Joint with Conference on Ultra Wideband Systems and Technologies. Joint UWBST & IWUWBS 2004 (IEEE Cat. No. 04EX812)*. IEEE. 2004, pp. 156–160.
- [131] Moe Z Win, Davide Dardari, Andreas F Molisch, Werner Wiesbeck, and W Jinyun Zhang. "History and applications of UWB". In: Institute of Electrical and Electronics Engineers. 2009.
- [132] IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs). *Application of IEEE Std 802.15.4*. Dec. 2014.
- [133] Klaus Witrisal, Paul Meissner, Erik Leitinger, Yuan Shen, Carl Gustafson, Fredrik Tufvesson, Katsuyuki Haneda, Davide Dardari, Andreas F Molisch, Andrea Conti, et al. "High-accuracy localization for assisted living: 5G systems will turn multipath channels from foe to friend". In: *IEEE Signal Processing Magazine* 33.2 (2016), pp. 59–70.
- [134] Henk Wymeersch, Stefano Marandò, Wesley M Gifford, and Moe Z Win. "A machine learning approach to ranging error mitigation for UWB localization". In: *IEEE transactions on communications* 60.6 (2012), pp. 1719–1728.
- [135] C. Xiang, P. Yang, C. Tian, L. Zhang, H. Lin, F. Xiao, M. Zhang, and Y. Liu. "CARM: Crowd-Sensing Accurate Outdoor RSS Maps with Error-Prone Smartphone Measurements". In: *IEEE Transactions on Mobile Computing* 15.11 (2016), pp. 2669–2681. doi: 10.1109/TMC.2015.2508814.

- [136] Wang Yan, Zhao Jing, and Zhang Nailong. "The designing of indoor localization system based on self-organized WSN using PulsON UWB sensors". In: *2015 2nd International Conference on Information Science and Control Engineering*. IEEE. 2015, pp. 965–969.
- [137] Jing Yang, BaiShun Dong, and Jiliang Wang. "VULoc: Accurate UWB Localization for Countless Targets without Synchronization". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6.3 (2022), pp. 1–25.
- [138] Kegen Yu, Kai Wen, Yingbing Li, Shuai Zhang, and Kefei Zhang. "A novel NLOS mitigation algorithm for UWB localization in harsh indoor environments". In: *IEEE Trans. on Vehicular Technology* 68.1 (2018), pp. 686–699.
- [139] A. Zanella. "Best Practice in RSS Measurements and Ranging". In: *IEEE Communications Surveys Tutorials* 18.4 (2016), pp. 2662–2686. doi: 10.1109/COMST.2016.2553452.
- [140] Wenda Zhao, Abhishek Goudar, and Angela P Schoellig. "Finding the Right Place: Sensor Placement for UWB Time Difference of Arrival Localization in Cluttered Indoor Environments". In: *IEEE Robotics and Automation Letters* 7.3 (2022), pp. 6075–6082.
- [141] Yue Zhao, Zan Li, Benjian Hao, and Jia Shi. "Sensor selection for TDOA-based localization in wireless sensor networks with non-line-of-sight condition". In: *IEEE Transactions on Vehicular Technology* 68.10 (2019), pp. 9935–9950.

PUBLICATIONS

PUBLICATION 1

**On the energy consumption and ranging accuracy of ultra-wideband
physical interfaces**

L. Flueratoru, S. Wehrli, M. Magno, and D. Niculescu

In: *IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2020, pp. 1–7

DOI: 10.1109/GLOBECOM42002.2020.9347984

Publication reprinted with the permission of the copyright holders.

On the Energy Consumption and Ranging Accuracy of Ultra-Wideband Physical Interfaces

Laura Fluera^{*†}, Silvan Wehrli[‡], Michele Magno[§] and Dragoş Niculescu^{*}

^{*} Computer Science Department, University Politehnica of Bucharest, Romania

[†] Electrical Engineering Unit, Tampere University, Finland

[‡] 3db Access AG, Zürich, Switzerland

[§] Department of Information Technology and Electrical Engineering, ETH Zürich, Switzerland

Abstract—Ultra-wideband (UWB) communication is attracting increased interest for its high-accuracy distance measurements. However, the typical current consumption of tens to hundreds of mA during transmission and reception might make the technology prohibitive to battery-powered devices in the Internet of Things. The IEEE 802.15.4 standard specifies two UWB physical layer interfaces (PHYs), with low- and high-rate pulse repetition (LRP and HRP, respectively). While the LRP PHY allows a more energy-efficient implementation of the UWB transceiver than its HRP counterpart, the question is whether some ranging quality is lost in exchange. We evaluate the trade-off between power and energy consumption, on the one hand, and distance measurement accuracy and precision, on the other hand, using UWB devices developed by Decawave (HRP) and 3db Access (LRP). We find that the distance measurement errors of 3db Access devices have at most 12 cm higher bias and standard deviation in line-of-sight propagation and 2–3 times higher spread in non-line-of-sight scenarios than those of Decawave devices. However, 3db Access chips consume 10 times less power and 125 times less energy per distance measurement than Decawave ones. Since the LRP PHY has an ultra-low energy consumption, it should be preferred over the HRP PHY when energy efficiency is critical, with a small penalty in the ranging performance.

Index Terms—Ultra-Wideband (UWB), Distance Measurement, Accuracy, Energy Efficiency

I. INTRODUCTION

Ultra-wideband (UWB) radio frequency (RF) signals have a high time resolution which enables the precise timestamping of their reception. As a result, they can provide time-of-flight (ToF) measurements with sub-nanosecond accuracy which can be further converted into distance and location information with cm- or dm-level accuracy [1], [2]. Location awareness can augment the capabilities of devices in the Internet of Things (IoT) and is often used in wireless sensor networks, industrial processes, or health-related applications. UWB devices usually consume tens to hundreds of mA, making them fit for localization tasks on energy-constrained devices. Their popularity has therefore risen and they have been recently included in smartphones [3], facilitating their large-scale deployment on the consumer market. Moreover, enhancements to the ranging capabilities of UWB devices are

currently being developed by the IEEE 802.15.4z Enhanced Impulse Radio (EIR) Task Group [4]. An overview of the enhancements to the IEEE 802.15.4 standard proposed by the EIR Task Group 4z can be found in [5].

The IEEE Standard for Low-Rate Wireless Networks 802.15.4 [6] specifies two UWB physical interfaces (PHYs), that use that use high- and low-rate pulse repetition (HRP and LRP, respectively). Transmitting pulses at low rates means that a single pulse can have the highest energy under UWB regulations [7]. The high instantaneous pulse amplitude enables the implementation of LRP PHYs with non-coherent receivers which are more energy-efficient than coherent ones. If we increase the pulse rate, we must decrease the energy per pulse, causing link budget losses which can be compensated by coherent pulse integration [8].

The LRP PHY, therefore, can be implemented with a more energy-efficient transceiver design than the HRP PHY, suitable for low-complexity active RFID tags. Coherent and non-coherent architectures have been compared from a theoretical standpoint in [9] and the authors found that the latter typically have an SNR loss of at least 5 dB but better multipath, phase jitter, and synchronization characteristics than the former.

The effect of LRP and HRP PHYs on distance measurement quality has not yet been evaluated. Previous work has focused mostly on the ranging accuracy of UWB devices [10]–[12] — most often, the Decawave DW1000 IC [13], which implements the HRP PHY — or on integrating UWB devices in localization systems [1], [2]. Comparisons of UWB devices have evaluated only the distance measurement accuracy, without regards to the power and energy consumption [14]–[16]. Moreover, only one of them has included an LRP PHY device (developed by Ubisense) [14].

To the best of our knowledge, we are the first to compare the distance measurements and power and energy consumption of LRP and HRP devices. To quantify the trade-off between energy consumption and ranging quality, we perform measurements using two commercially-available UWB devices: the Decawave DW1000 IC (HRP) and the 3db Access 3DB6830C IC [17] (LRP)¹.

The main contributions of this paper are the following:

The authors gratefully acknowledge funding from European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints, <http://www.a-wear.eu/>).

¹We will refer to the 3db 3DB6830C (Release 2016) and the Decawave DW1000 (Release 2014) as the 3db and Decawave ICs, respectively.

- We measure the average power consumption of the chosen devices in the receive, transmit, and idle modes and compute their energy consumption per ranging.
- We compare the channel impulse responses (CIRs) of 3db and Decawave devices acquired in identical settings, which are essential to understand the distance estimation.
- We compare the range of 3db and Decawave devices.
- We conducted an extensive distance measurement campaign with 3db devices in indoor multipath environments. We compare our ranging results with already-published results on Decawave devices [10]–[12], [15]. We classify distance measurement errors based on whether they were acquired in line-of-sight (LOS) propagation, with no obstruction between the transmitter and the receiver, or non-line-of-sight (NLOS) propagation, when an object blocks the direct path.

II. DEVICE CHARACTERISTICS

Decawave devices are compliant with the HRP PHY defined in the IEEE 802.15.4a amendment, now part of the main standard [6]. They are the most widely-used UWB devices, which is why we chose them to represent the HRP PHY class. The 3db IC is compliant with the LRP PHY specified in the upcoming IEEE 802.15.4z amendment [4]. The chip is already used for secure keyless car access but it has not been evaluated in high-accuracy applications so far.

In this section, we study how Decawave and 3db devices differ in terms of pulse repetition frequency, receiver architecture, and ranging methods in Sections II-A, II-B, and II-C, respectively. We also highlight how their characteristics affect the power consumption and distance measurement accuracy.

A. High- and Low-Rate Pulse Repetition

UWB transmissions have to satisfy two constraints imposed by international regulations [18]: a maximum *average* power spectral density (PSD) of -41.3 dBm/MHz (averaged over 1 ms) and a maximum *peak* power spectral density of 0 dBm/50MHz. UWB devices can therefore transmit over a fixed period of time either few pulses at high power levels or a large number of pulses with lower transmit power. The first situation falls under the LRP specification and is employed by 3db devices, while the latter is known as HRP and is used by Decawave. If optimally employed, both of these technologies benefit from an *equal* average transmitted RF energy.

Since the HRP PHY transmits individual pulses with lower energy than the LRP, the received pulse energy is also lower for the same path loss (same distance). Therefore, the HRP PHY needs more sophisticated techniques to extract weaker pulses from the receiver noise, typically performed with correlations over a large number of samples.

B. Device Architecture

Owing to the LRP PHY, 3db devices can be implemented with a *non-coherent* receiver based on energy detection (ED) for signals modulated with binary frequency-shift keying

(BFSK). This signaling scheme allows the circuitry implementation of the receiver to be more energy efficient than the *coherent* Decawave receiver.

Coherent receivers have low sensitivity to inter-symbol and co-user interference and benefit from the multipath diversity of the UWB channel [19]. At the same time, the receiver architecture demands high computational resources and hardware complexity [20]. For optimal reception, the coherent receiver needs to estimate the multipath delays, their channel coefficients, and the pulse shape distortion [20]. Precisely estimating the carrier phase is crucial for recovering the baseband pulse, since inaccuracies will result in signal power loss and crosstalk interference in PSK-modulated signals [21]. For a carrier frequency of 8 GHz, a time shift of half of the pulse period flips the phase of the signal, so coherent UWB systems generally tolerate rotations only within $\pi/4$ of the signal phase (around 30 ps). These requirements increase the power consumption of coherent demodulators [19].

Non-coherent receivers estimate channel coefficients based on the envelope rather than the phase and amplitude of the received signal, relaxing synchronization constraints. The timing requirements of a non-coherent receiver are dependent only on the pulse envelope, which is related to the pulse bandwidth. For instance, if the pulse bandwidth is 500 MHz, the non-coherent receiver needs to operate with a timing resolution of 1 ns and does not need sophisticated RF carrier synchronization. Therefore, non-coherent receivers can be more energy-efficient albeit with a higher bit error probability in comparison with the coherent architecture [21].

C. Ranging Methods

A popular application for UWB devices is indoor localization. Owing to the high time resolution of UWB signals, time-based localization techniques are the most suitable for UWB devices [22]. In this paper, we chose to compare the *ranging* accuracy and precision of the UWB devices instead of the *localization* ones for several reasons. First, many popular localization algorithms (e.g. trilateration) work directly with distance estimates between the tracked device (tag) and the reference devices (anchors), so our results can be used to compute the expected localization accuracy of those algorithms. Second, localization results are heavily influenced by factors unrelated to the devices themselves, such as the anchor placement, the location of the tag², or the localization algorithm. Since we are interested in comparing the devices themselves, it is easier to avoid these effects by evaluating the ranging performance instead of the localization one. Third, there are important applications of UWB devices which do not involve localization, such as keyless car access and, in the future, possibly contact tracing, so our results can be used to evaluate which physical interface is more suitable for them.

The distance between two devices can be estimated based on the time of flight (ToF) of the signal. If we know the

²Localization errors are larger near the anchors and lower in the center of the tracking area [23].

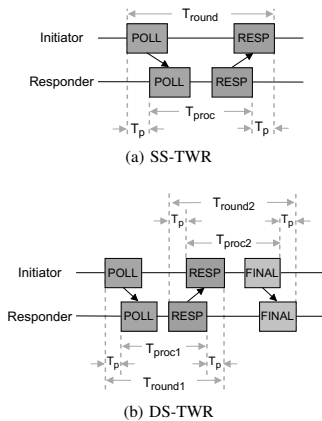


Fig. 1: Message exchange in the single- and double-sided two-way ranging which are the ranging methods of choice for the LRP and HRP PHYs, respectively.

transmission time (T_1) of the signal measured by the sender and the arrival time (T_2) at the receiver, we can compute the distance as:

$$d = (T_2 - T_1) \cdot c, \quad (1)$$

where c is the speed of light and $T_p \triangleq T_2 - T_1$ is the propagation time of the signal. To accurately estimate the distance, the devices need to be tightly clock synchronized, as a small mismatch of 1 ns can introduce a distance error of around 30 cm. Because synchronizing the sender and the receiver is usually unfeasible in practice, more messages are exchanged in order to reduce such errors, leading to the single- and the double-sided two-way ranging.

Single-Sided Two-Way Ranging (SS-TWR): The SS-TWR uses two messages per distance estimate, as shown in Fig. 1a. The propagation time is:

$$T_p = \frac{T_{round} - T_{proc}}{2}, \quad (2)$$

where T_{round} is the time spent in one message exchange and T_{proc} is the processing time on the responder side. It can be shown that the error in estimating T_p is [24]:

$$e_{T_p} = e_1 \cdot T_p + \frac{1}{2} T_{proc} (e_1 - e_2), \quad (3)$$

where e_1 and e_2 are the clock drift errors of the initiator and responder, respectively. The main source of errors in the SS-TWR are T_{proc} , which is in the range of hundreds of microseconds, and the clock drift, which can be up to ± 20 ppm in systems compliant with the IEEE 802.15.4 standard [6].

In the LRP PHY, a location-enhancing information postamble is introduced at the end of each message to estimate the clock drift error [6]. In addition, the processing time of LRP messages is shorter than the one of HRP. It is also more convenient to minimize the number of exchanged messages in the TWR since this reduces the time needed to obtain one




Fig. 2: The power spectral density of 3db and Decawave ICs.

distance measurement. Therefore, the SS-TWR is usually the method of choice for LRP devices.

Double-Sided Two-Way Ranging (DS-TWR): Because the HRP PHY does not include a postamble, it needs another method to minimize clock drift errors. This can be achieved by exchanging an additional message, as shown in Fig. 1b, leading to the DS-TWR. In this case, e_{T_p} is minimized if the processing times T_{proc1} and T_{proc2} are equal. However, this constraint is often hard to enforce in practice. An alternative DS-TWR has been proposed which minimizes clock drift errors even with asymmetric processing times [24]. This is the method currently used by Decawave [13].

III. EVALUATION SETUP

To evaluate how the system architecture influences the power consumption and distance measurements, we implement the ranging techniques on 3db and Decawave hardware. In the following, we describe the device setups.

3db Access: We integrate the 3db chip in an Arduino shield on top of an Arduino M0 board. The communication between the chip and the host MCU is performed via SPI. We use the lowest channel, centered at 6.52 GHz, and a peak data rate of 247 kb/s. The 10 dB bandwidth of a pulse is 380 MHz and, because pulse spectra partially overlap in BFSK modulation, the total system bandwidth is approximately 620 MHz. The packet duration is 400 μ s.

Decawave: We use the Decawave EVB1000 evaluation boards, which include a software kit for ranging applications. We configure the devices to communicate on a similar center frequency as 3db ones, of 6.49 GHz (Channel 5), using a 3 dB bandwidth of 499.2 MHz (equivalent to a 10 dB bandwidth of ≈ 662 MHz). We set a data rate of 110 kb/s, a PRF of 16 MHz, and a preamble length of 2048 symbols (Mode 3). The packet duration in this mode is 3487 μ s. The lower data rate allows longer range and increased link budget compared to higher-rate setups, while the PRF and preamble length were chosen to minimize NLOS effects.

The lower pulse bandwidth of 3db devices can, in theory, decrease the ranging precision, due to the lower time resolution. However, it is compensated by the frequency diversity added by the BFSK modulation.

Both ICs were configured to operate within UWB regulations, which specify a maximum transmit level of -41.3 dBm/MHz [18]. Fig. 2 shows their measured power spectral densities (PSDs). The maximum transmit levels of




Fig. 3: The power consumption profiles of 3db and Decawave initiators and responders. The average power consumption of 3db devices is about 10 times lower than the one of Decawave devices in all modes (idle, transmission, reception).

TABLE I: The average power consumption of 3db and Decawave devices in the transmission (TX), reception (RX), and idle modes.

Device	Average power consumption [mW]		
	TX	RX	Idle
3db Access	20.69	40.70	6.60
Decawave	194.54	492.45	68.78

the DW1000 and the 3db IC are -45.31 dBm/MHz and -43.86 dBm/MHz, respectively, but they do not include the antenna gain. The antenna gains of 3db and Decawave devices are 2 dBi and 3.3 dBi (at the center frequency), respectively, resulting overall in almost equal PSDs.

IV. EVALUATION

In this section, we measure the average power consumption, channel impulse response (CIR), range, and distance measurement accuracy and precision of 3db and Decawave devices.

A. Power Consumption

We measured the current consumption of 3db and Decawave devices when performing the SS-TWR and DS-TWR, respectively, with a Keysight DC Power Analyzer. The 3db IC is powered with 1.25 V and the supply voltage of Decawave is 3.3 V. We configured the 3db devices to perform one SS-TWR every approximately 14 ms. The Decawave profile was obtained with the default firmware of the EVK1000 kit and the DecaRanging application which uses the DS-TWR method. Decawave recommends the use of guard times on the order of hundreds of ms between each message [25]. In our setup, Decawave devices perform one DS-TWR every 500 ms.

Fig. 3 shows the power consumption of the 3db and Decawave initiators and responders (note the logarithmic y-axis). We isolate the receive (RX), transmit (TX), and idle modes and compute the average power consumption of each state, presented in Table I.

First, each device consumes more power in the RX mode than in the TX mode. Because of the large signal bandwidth, receivers need analog-to-digital converters (ADCs) operating at high sampling rates (on the order of Gsamples/s). This increases the processing load at the receiver, demanding more energy consumption. The RX-to-TX average power consumption ratio is approximately 2 in the case of 3db devices and 2.5 in the case of Decawave.

Overall, the average power consumption of the 3db IC is 10 times lower than the one of the Decawave IC. Note that the average power consumption of Decawave devices in the *idle* mode is about 70 % higher than the one of 3db devices in the *receive* mode, also the most power-hungry state.

Besides these states, both devices have sleep modes which consume $1 \mu\text{A}$ (sleep) or $50\text{--}100 \text{ nA}$ (deep sleep) in the case of Decawave [13] and 500 nA in the case of 3db. When used together with the regular operational modes, they can increase the battery life of the device.

The power consumption profile is a starting point for evaluating the energy consumption of an UWB-based localization system (LS). A localization system (LS) consists of multiple, fixed devices with known positions (called anchors) which localize mobile devices (called tags). A key challenge is minimizing the energy consumption of the *tag*, which is usually battery-powered. To avoid synchronizing the tag and the anchors, the tag can *initiate* the message exchange and stay in the idle or sleep mode between rangings. Using the SS-TWR implies, in this case, that the *tag* estimates the distance (or the location). Alternatively, if the tag initiates a DS-TWR, the anchors are the last entities in the message exchange, so the *anchor* estimates the distance (or the location).

To illustrate the energy efficiency of a tag in a LS, let us consider the most favorable scenario for each device, in which the tag is the initiator. We disregard the time spent in the idle mode³ and compute the energy consumption only when the device is in the TX or RX mode. Because the packet duration of the 3db chip is 10 times shorter than the one of Decawave, a 3db tag would consume 0.028 mJ per ranging, while a Decawave tag would need 3.55 mJ. Therefore, a 3db tag can consume 125 less energy than a Decawave tag.

Decawave devices can indeed be more energy-efficient in the high-rate mode (6.8 Mb/s). In the best case, this allows a 20x shorter packet duration but about 1.4x higher current consumption [26]. So even in this mode the 3db IC consumes at least 9x lower energy, without taking into account the idle time. The 6.8 Mb/s mode also reduces the range, as we will see in Section IV-C.

³The time spent in the idle mode is subject to the desired location update rate and the chosen guard times. Since they can be chosen freely to a certain extent, we neglect them in the energy computation.




Fig. 4: Averaged CIRs and their path variations obtained by placing pairs of 3db and Decawave devices in LOS and NLOS of each other in exactly the same positions.

B. Channel Impulse Response

Time-based ranging methods precisely estimate the moment at which the signal is received. This corresponds to the leading edge (LE) of the first peak in the CIR. In the 3db IC, the CIR is obtained by directly integrating the pulse envelopes. The DW1000 IC estimates the CIR by accumulating over time the correlation of the received signal and a known preamble sequence [13]. Both CIRs have a sampling period of 1 ns.

The measurement quality depends on how the signal propagates between the transmitter and the receiver. In the LOS scenario, there is a direct, unobstructed path between two devices, leading to minimal errors. In NLOS propagation, walls or large objects block the direct path and therefore attenuate or block the LOS signal. Copies of the signal reflected on surrounding objects might still reach the receiver, but the additional delay can introduce large ranging errors in NLOS scenarios. We should therefore distinguish between these two cases when analyzing data.

We conducted an experiment in which we placed pairs of 3db and Decawave devices in a LOS and a NLOS scenario, in exactly the same positions. In the NLOS case, a concrete pillar was blocking the direct path between the devices. We acquired approximately 150 CIR realizations in each recording. Fig. 4 shows the aligned and averaged CIRs (truncated to 110 samples) and their path variations. The 3db and Decawave CIRs contain in total 256 and 992 samples, respectively.

Peaks in the CIR correspond to replicas of the signal arriving through multiple paths. Because the pulse bandwidth of 3db devices is significantly lower than the one of Decawave devices (380 MHz vs. 662 MHz, respectively), pulses in the 3db CIR are wider. A larger bandwidth increases the time resolution of the device, which improves measurement precision.

The ratio between the amplitude of the first path compared to later replicas is influenced by the RF front-end linearity and therefore depends on the receiver implementation. One class of TOA estimators identifies the maximum-amplitude path and searches backward for the first sample that exceeds the noise floor and is smaller than the maximum amplitude [27]. In NLOS, delayed paths with high amplitudes can produce local

TABLE II: The useful range of Decawave and 3db devices.

Device	Data rate	Range [m]
3db Access	247 kb/s	116
Decawave	110 kb/s	105
Decawave	6.8 Mb/s	80

minima estimates. Therefore, TOA estimation should take into account NLOS scenarios for the best accuracy.

C. Distance Measurements

Range: We measured the range of Decawave and 3db devices with no obstruction between the transmitter and the receiver on a marked running track in increments of 5 m. At each point, we recorded the number of timeouts (messages without a response). We define the useful range as the distance at which measurements have less than 10% timeout probability. Both devices use the maximum transmission power within UWB regulations. The range can be extended by increasing the transmission power and tuning the channel and preamble length settings (albeit with higher energy consumption). We also measured the range of Decawave devices in the 6.8 Mb/s mode. TABLE II shows the measured ranges. The range of 3db devices and Decawave devices in the 110kb/s mode exceeds 100 m. The high-rate mode (6.8 Mb/s) of Decawave has a lower range, of 80 m. Therefore, although the high-rate mode of Decawave reduces the energy consumption, it also decreases the range.

Accuracy and Precision: To characterize distance measurements, we performed an extensive measurement campaign using 3db devices. Because Decawave devices have been included in numerous studies [10]–[12], [15], we rely on already-published results for the comparison. Except for [12], which does not specify the Decawave settings used, all the references use the 110 kb/s data rate.

Our database includes over 12,000 measurements acquired with the 3db IC in indoor spaces (e.g. large offices, small rooms, and hallways) labeled as either LOS or NLOS, covered in almost equal proportions. To reflect typical real-life situations, we acquired data when devices are stationary or moving at walking speed. In NLOS measurements, the obstruction was caused by the human body, walls, or pillars, out of which the latter two sharply attenuate the direct path. In all NLOS cases, the signal can still arrive at the receiver through reflections—in other words, we did not perform experiments in which the devices are in different rooms.

Since we compare 3db and Decawave devices using measurements obtained in different experiments, there is a chance that the results differ not because of the devices but due to the conditions under which they were acquired (i.e. type of NLOS or multipath environment). On the other hand, this weakness is also a strength, since by relying on Decawave results from more sources, it is more likely that they are more general. When acquiring 3db Access measurements, we strove to recreate all LOS and NLOS conditions considered in the papers which used Decawave devices [10]–[12], [15].




Fig. 5: The PDF and boxplots of LOS and NLOS distance measurements errors of the 3db IC. The median and IQR are, respectively, 9.04 cm and 14.9 cm in LOS and 47.74 cm and 56.3 cm in NLOS.

TABLE III: Mean and standard deviation of ranging errors in LOS and NLOS with 3db and Decawave devices. We split LOS measurements based on whether the antennas were facing each other or had other poses.

Device	LOS error [cm]		NLOS error [cm]
	Facing antennas	Different antenna poses	
3db Access	4.8 ± 8.4	11 ± 15	62.5 ± 104
Decawave	[10]	4.1 ± 3.2	$7.3 \pm 3.6^*$
	[11]	-	0 ± 15
	[15]	0.3 ± 5.5	-
			34 ± 35

* Average over errors with different antenna poses

Fig. 5 shows the probability distribution function (PDF) of LOS and NLOS errors and TABLE III summarizes their statistics. In Fig. 5, we restricted the domain to $[-0.5, 3]$ m for better visualization but the errors extend up to 19.95 m. Measurements outside the shown boundary occur only in NLOS and represent 0.2% of their total. These outliers are caused by a miscalculation of the TOA in NLOS scenarios in the firmware (the issue was later fixed).

In LOS, the irregular radiation pattern of UWB antennas can cause errors up to ± 0.4 m with certain antenna alignments [28]. Therefore, we distinguish measurements where the antennas mounted on the devices were facing each other, which yield the smallest errors, from measurements acquired with other antenna poses. The 3db Access measurements with facing antennas were acquired on a running track (the only outdoor measurements) where we varied the distance between the devices in steps of 8.5 m up to 51 m. In comparison with Decawave, the ranging accuracy and precision of 3db devices are at most 5.2 cm higher with the fixed antenna pose and at most 12 cm higher with different antennas poses.

NLOS scenarios cause the reported distance to be greater or equal than the true one, so the distribution of 3db measurement errors is heavy-tailed and no longer Gaussian-shaped, as shown in Fig. 5. In NLOS, the bias depends on the additional path traveled by the signal and can therefore vary between experiments with different room plans and furnishing. In this case, the error *spread* given by the standard deviation or the interquartile range (IQR) better describes the performance.

In TABLE III, we report the mean and standard deviation of 3db NLOS measurements in order to be consistent with cited Decawave results. However, note that these statistics are shifted upwards because of the infrequent but large outliers,

so the median and IQR characterize the distribution more robustly. The IQR of the NLOS error distribution is 56.3 cm. If we eliminate the outliers (which, in a real application, can be filtered out), the standard deviation is halved to 42.57 cm.

The reported NLOS standard deviation of Decawave devices is 35 cm in office environments [15], between 3.1–18.7 cm with human body shadowing [10], and around 20 cm when the path is blocked by panels made of different materials or concrete walls [12]. NLOS measurements with 3db devices were obtained with *all* these types of obstructions. Therefore, we should compare the 3db error spread with an average over the reported Decawave standard deviations, which is around 20 cm. The IQR of 3db NLOS errors is therefore 2–3 larger than the spread of Decawave errors.

V. DISCUSSION

The results in Section IV warrant a discussion about how the accuracy of distance measurements obtained with the LRP PHY can be improved while still benefiting from the same low energy consumption. We identify several improvements that can be done at the hardware, system, and application level in Sections V-A, V-B, and V-C, respectively.

A. Hardware Improvements

As noted in Section III, 3db devices have almost half the pulse bandwidth of Decawave devices. Since the time resolution of UWB devices is proportional to the pulse bandwidth, increasing the latter could improve the TOA estimation accuracy of 3db devices in both LOS and NLOS conditions.

At the moment, 3db devices can obtain the CIR either from the preamble *or* the postamble for a given measurement. If both CIRs were available at the same time, we could use a similarity metric to detect significant differences in their shape, which usually indicates a highly dynamic environment (i.e. NLOS). Measurements acquired in such conditions could then be discarded or further processed.

In NLOS situations, ranging errors are exacerbated by the clock drift estimation. 3db devices detect the TOA of a packet during the preamble and also during the postamble (at the end of a frame). The clock offset between two devices is computed based on the difference between the two TOAs. Since in NLOS the CIR is highly variable, the postamble TOA might include not only the clock drift error but also a delay caused by the excess path traveled by the signal. In this case, the clock offset estimation will incorrectly compensate for this excess path, increasing the ranging error even more. Discarding measurements in which the postamble and preamble CIRs are very different could therefore reduce the magnitude of NLOS errors, shortening the tail of their PDF.

B. System Improvements

System-level improvements comprise aspects that can be implemented (usually in the firmware) with current hardware capabilities. A firmware issue caused large outliers in 3db measurements acquired in NLOS, which significantly decreased the accuracy and precision of NLOS measurements.

The issue has since been solved, which should show an improvement in future measurements.

As already noted, the multipath components (MPCs) in the CIR significantly affect the TOA estimation. In NLOS, the first path can be very weak (close to the noise floor) and arrive tens of nanoseconds before the strongest MPC [29], so customizing the first path detection algorithm for this scenario could further improve NLOS results.

C. Application Improvements

Localization applications can improve their accuracy through NLOS detection and mitigation techniques using data readily available from the hardware, such as the CIR and its statistics [30]. In addition, filters (e.g. the Extended Kalman Filter) can be used to remove outliers in distance measurements frequently encountered in NLOS situations.

VI. CONCLUSION

We compared the power and energy consumption and distance measurement statistics of 3db Access and Decawave UWB devices, which implement the LRP and HRP PHYs, respectively. In LOS propagation, 3db devices have slightly higher bias and lower precision than Decawave devices. In NLOS scenarios, the error spread of 3db devices is 2–3 times larger than the one of Decawave. On the other hand, 3db devices have 10x lower average power consumption and 125x lower energy consumption per distance measurement compared to Decawave devices. Therefore, the LRP PHY is suitable for applications which can tolerate a loss in ranging accuracy and precision for higher energy efficiency.

In the future, we will optimize the leading-edge detection of 3db devices to reduce the NLOS error spread and we will integrate the devices in a localization system. The reduced energy consumption can provide location awareness to previously constrained devices (millirobots, body sensors) and generate novel localization applications (for instance, in swarm robotics). It would be interesting to study how the short air time and presumably high device density will impact the access to the medium and the localization update rate of future positioning systems.

REFERENCES

- [1] D. Lymberopoulos and J. Liu, "The Microsoft indoor localization competition: Experiences and lessons learned," *IEEE Signal Processing Magazine*, vol. 34, pp. 125–140, Sept. 2017. Conference Name: IEEE Signal Processing Magazine.
- [2] 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 *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pp. 61–72, IEEE, 2019.
- [3] "Apple U1 TMKA75 ultra wideband (UWB) chip analysis." TechInsights Inc. Report. Accessed on 15-04-2020.
- [4] "IEEE 802.15.4z task group." <http://www.ieee802.org/15/pub/TG4z.html>. Accessed on 15-04-2020.
- [5] 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)*, pp. 1–6, IEEE, 2019.
- [6] "IEEE 802.15.4 standard for low-rate wireless networks," 2015.
- [7] "Application note APR001. UWB regulations. A summary of worldwide telecommunications regulations governing the use of ultra-wideband radio." Decawave Ltd. Accessed: 15-04-2020.
- [8] "Application of IEEE standard 802.15.4," May 2014.
- [9] K. Witrisal, G. Leus, G. J. Janssen, M. Pausini, F. Troesch, T. Zasowski, and J. Romme, "Noncoherent ultra-wideband systems," *IEEE Signal Processing Magazine*, vol. 26, pp. 48–66, July 2009. Conference Name: IEEE Signal Processing Magazine.
- [10] Q. Tian, K. I.-K. Wang, and Z. Salcic, "Human body shadowing effect on UWB-based ranging system for pedestrian tracking," *IEEE Transactions on Instrumentation and Measurement*, vol. 68, pp. 4028–4037, Oct. 2019. Conference Name: IEEE Transactions on Instrumentation and Measurement.
- [11] B. J. Silva and G. P. Hancke, "Ranging error mitigation for through-the-wall non-line-of-sight conditions," *IEEE Transactions on Industrial Informatics*, pp. 1–1, 2020. Conference Name: IEEE Transactions on Industrial Informatics.
- [12] A. Schenck, E. Walsh, J. Reijniers, T. Ooijsvaar, R. Yudianto, E. Hostens, W. Daems, and J. Steckel, "Information theoretic framework for the optimization of UWB localization systems," in *2018 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, Sept. 2018. ISSN: 2471-917X.
- [13] "DW1000 user manual. Version 2.11." Decawave Ltd.
- [14] A. R. J. Ruiz and F. S. Granja, "Comparing Ubisense, Bespoon, and Decawave UWB location systems: Indoor performance analysis," *IEEE Transactions on Instrumentation and Measurement*, vol. 66, no. 8, pp. 2106–2117, 2017.
- [15] A. R. Jiménez and F. Seco, "Comparing Decawave and Bespoon UWB location systems: Indoor/outdoor performance analysis," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, IEEE, 2016.
- [16] J. Wang, A. K. Raja, and Z. Pang, "Prototyping and experimental comparison of IR-UWB based high precision localization technologies," in *UIC-ATC-ScalCom*, pp. 1187–1192, IEEE, 2015.
- [17] "3db Access." www.3db-access.com. Accessed on 15-04-2020.
- [18] F. C. Commission et al., "Revision of part 15 of the commission's rules regarding ultra-wideband transmission systems," *First Report and Order, FCC 02-48*, 2002.
- [19] M. Weisenhorn and W. Hirt, "Robust noncoherent receiver exploiting UWB channel properties," in *2004 International Workshop on Ultra Wideband Systems Joint with Conference on Ultra Wideband Systems and Technologies. Joint UWBST & IWUWBS 2004 (IEEE Cat. No. 04EX812)*, pp. 156–160, IEEE, 2004.
- [20] H. Arslan, Z. N. Chen, and M.-G. Di Benedetto, *Ultra wideband wireless communication*. John Wiley & Sons, 2006.
- [21] J. G. Proakis and M. Salehi, *Digital communications*, vol. 4. McGraw-Hill New York, 2001.
- [22] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu, "Localization via ultra-wideband radios: A look at positioning aspects for future sensor networks," *IEEE signal processing magazine*, vol. 22, no. 4, pp. 70–84, 2005.
- [23] I. Guvenc, C.-C. Chong, and F. Watanabe, "Analysis of a linear least-squares localization technique in los and nlos environments," in *2007 IEEE 65th Vehicular Technology Conference-VTC2007-Spring*, pp. 1886–1890, IEEE, 2007.
- [24] D. Neirynek, E. Luk, and M. McLaughlin, "An alternative double-sided two-way ranging method," in *2016 13th workshop on positioning, navigation and communications (WPNC)*, pp. 1–4, IEEE, 2016.
- [25] "Ranging demo user guide. Understanding and using the DecaRanging ranging demo (PC) application. Version 4.5." Decawave Ltd.
- [26] "DW1000 datasheet. Version 2.09." Decawave Ltd.
- [27] H. Soganci, S. Gezici, and H. V. Poor, "Accurate positioning in ultra-wideband systems," *IEEE Wireless Communications*, vol. 18, no. 2, pp. 19–27, 2011.
- [28] A. Ledergerber and R. D'andrea, "Calibrating away inaccuracies in ultra wideband range measurements: A maximum likelihood approach," *IEEE Access*, vol. 6, pp. 78719–78730, 2018. IEEE Access.
- [29] Z. Sahinoglu, S. Gezici, and I. Guvenc, "Ultra-wideband positioning systems," *Cambridge, New York*, 2008.
- [30] J. Khodjaev, Y. Park, and A. Saeed Malik, "Survey of NLOS identification and error mitigation problems in UWB-based positioning algorithms for dense environments," *Annals of telecommunications*, vol. 65, pp. 301–311, June 2010.

PUBLICATION 2

High-Accuracy Ranging and Localization With Ultrawideband Communications for Energy-Constrained Devices

L. Flueratoru, S. Wehrli, M. Magno, E. S. Lohan, and D. Niculescu

IEEE Internet of Things Journal 9.10 (2021), pp. 7463–7480

DOI: 10.1109/JIOT.2021.3125256

Publication reprinted with the permission of the copyright holders.

High-Accuracy Ranging and Localization with Ultra-Wideband Communications for Energy-Constrained Devices

Laura Flueratoru, *Graduate Student Member, IEEE*, Silvan Wehrli, *Member, IEEE*, Michele Magno, *Senior Member, IEEE*, Elena Simona Lohan, *Senior Member, IEEE*, and Dragoș Niculescu

Abstract—Ultra-wideband (UWB) communications have gained popularity in recent years for being able to provide distance measurements and localization with high accuracy, which can enhance the capabilities of devices in the Internet of Things (IoT). Since energy efficiency is of utmost concern in such applications, in this work we evaluate the power and energy consumption, distance measurements, and localization performance of two types of UWB physical interfaces (PHYs), which use either a low- or high-rate pulse repetition (LRP and HRP, respectively). The evaluation is done through measurements acquired in identical conditions, which is crucial in order to have a fair comparison between the devices. We performed measurements in typical line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios. Our results suggest that the LRP interface allows a lower power and energy consumption than the HRP one. Both types of devices achieved ranging and localization errors within the same order of magnitude and their performance depended on the type of NLOS obstruction. We propose theoretical models for the distance errors obtained with LRP devices in these situations, which can be used to simulate realistic building deployments and we illustrate such an example. This paper, therefore, provides a comprehensive overview of the energy demands, ranging characteristics, and localization performance of state-of-the-art UWB devices.

Index Terms—Ultra-Wideband (UWB), Distance Measurement, Ranging, Accuracy, Energy Efficiency.

I. INTRODUCTION

ULTRA-WIDEBAND (UWB) communications have become increasingly popular in recent years for their high-accuracy ranging and localization capabilities, which makes

This work was supported by funding from European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints, <http://www.a-wear.eu/>). This article is based on a previous conference paper [1] presented at the IEEE Global Communications Conference (GLOBECOM), Dec. 2020. (*Corresponding author: Laura Flueratoru.*)

Laura Flueratoru is with the Department of Computer Science, University Politehnica of Bucharest, Romania, and with the Department of Electrical Engineering, Tampere University, Finland (e-mail: laura.flueratoru@upb.ro).

Silvan Wehrli is with 3db Access AG (e-mail: silvan.wehrli@3db-technologies.com).

Michele Magno is with the Project-Based Learning Center, ETH Zürich, Switzerland (e-mail: michele.magno@pbl.ee.ethz.ch).

Elena Simona Lohan is with the Department of Electrical Engineering, Tampere University, Finland (e-mail: elena-simona.lohan@tuni.fi).

Dragoș Niculescu is with the Department of Computer Science, University Politehnica of Bucharest, Romania (e-mail: dragos.niculescu@upb.ro).

Copyright (c) 20xx IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

them promising candidates for providing location services to devices in the Internet of Things (IoT), industrial deployments, or wireless sensor networks in general. More recently, UWB chipsets have been included in smartphones and it is estimated that 50% of the smartphones on the market will incorporate UWB chipsets by 2027 [2]. Given the fast adoption of UWB technology and its integration and interaction with devices in the IoT, it is crucial to evaluate both its ranging and localization performance and its energy efficiency in order to determine its suitability for different types of applications.

UWB devices provide time-of-flight (ToF) measurements with sub-nanosecond accuracy which can be used to estimate the distance between two devices. Distance measurements (or ranges) are the basis of the true-range multilateration algorithm which is used in many localization applications [3]. Therefore, evaluating ranging errors is often the first step in analyzing the localization accuracy of UWB localization systems.

The energy consumption of UWB devices depends on their architecture. The IEEE Standard for Low-Rate Wireless Networks 802.15.4 [4] specifies two UWB physical interfaces (PHYs), that use high- and low-rate pulse repetition (HRP and LRP, respectively). Transmitting pulses at low rates enables a more energy-efficient implementation of LRP PHYs, using non-coherent receivers, than the ones based on coherent receivers, which are typically used in HRP PHYs. Coherent receivers use the phase of the signal in the detection process, while non-coherent receivers can estimate the channel coefficients with lower synchronization constraints based on the envelope of the signal. This makes LRP UWB devices suitable for energy-constrained devices. So far, it has not been clear whether this advantage comes with a cost in the ranging and localization performance.

Although coherent and non-coherent UWB receivers have been compared from a theoretical standpoint in literature, these studies have relied on simulations rather than measurements [5], [6]. Previous work that analyzed the ranging accuracy of UWB devices through measurements [7], [8], [9] has focused mostly on the Decawave DW1000 IC [10], which implements the HRP PHY. Few works have analyzed commercially available LRP UWB devices and they targeted mostly their ranging accuracy without a detailed analysis of their power and energy consumption [11], [12]. To the best of our knowledge, only one paper included a comparison of HRP and LRP devices (developed by Decawave and Ubisense, respectively) [11] but only on their ranging and localization

performance, without regards to the energy efficiency of the devices. Therefore, the current literature on comparisons of the two types of PHYs, on the one hand, and on commercially available UWB LRP devices, on the other hand, is very scarce. In particular, LRP devices deserve more attention since they can be implemented with energy-efficient receivers and can therefore potentially enable ranging and localization applications on ultra-low-power devices.

In our previous work [1], we compared for the first time the power and energy consumption and the ranging performance of LRP and HRP devices using two commercially-available UWB devices: the Decawave DW1000 IC (HRP)¹ and the 3db Access 3DB6830C IC [14] (LRP)². For the ranging performance, we used a database of distance measurements acquired with 3db devices and compared their statistics with results obtained with Decawave devices from the literature. This paper goes one step further and compares the ranging and localization performance of LRP and HRP devices using real measurements acquired in *identical* settings. This last detail is crucial for a fair comparison of the devices since different environments can have a different impact on distance measurement errors. In addition, the current work offers a more in-depth analysis of the typical ranging errors of LRP devices in several scenarios.

Since indoor localization is often subject to multipath and shadowing phenomena, we analyzed the statistics of ranging errors in line-of-sight (LOS) and three non-line-of-sight (NLOS) scenarios, where the obstruction between the transmitter and the receiver was caused by a person, a gypsum wall (also called drywall panel), or a concrete wall. We derived statistical models for the error distributions obtained from measurements, which can be used to simulate realistic ranging and localization scenarios that would otherwise take days or weeks to implement and evaluate. We argue, in particular, that there are still unsolved problems about deploying a UWB-based localization system inside a building. Finding adequate LOS and NLOS error models, such as the ones proposed in our work, and using them to simulate the expected localization errors can help in this regard. For instance, many existing works [11], [15], [16], [17], [18], [19] consider only setups where the anchors are placed inside the same room because they yield the highest localization accuracy. However, this constraint is often hard to enforce in real deployments. For one, in highly compartmentalized spaces (for instance, office buildings) this would lead to a high anchor density, which in turn increases the deployment costs, the complexity (in terms of synchronization constraints, multiple access, anchor selection and placement strategy, etc.), and the total energy consumption of the localization system. Second, we show that this constraint might not be even needed, for instance when rooms are divided by shallow walls which cause only small localization errors.

¹Decawave has recently been acquired by the semiconductor company Qorvo [13], hence is in the process of changing its name to Qorvo. Since this change is rather recent, we still refer to the company and devices as "Decawave," this being the name under which they are still widely known.

²We will refer to the 3db 3DB6830C (Release 2016) and the Decawave DW1000 (Release 2014) as the 3db and Decawave ICs, respectively.

This paper, therefore, provides a comprehensive outlook on the typical power consumption and ranging performance of two state-of-the-art UWB devices, as well as their expected localization accuracy based on both real measurements and simulations. Our proposed error models can be used in future works to simulate custom building deployments and our measurements are publicly available³ to facilitate future research.

To summarize, the main contributions of this paper are the following:

- We analyze the average power consumption of 3db Access (LRP) and Decawave (HRP) devices in the receive, transmit, and idle modes and compute their energy consumption per distance measurement.
- We evaluate and compare the accuracy and precision of distance measurements of 3db Access and Decawave devices based on measurements recorded in identical settings in LOS and NLOS scenarios caused by drywall, a concrete wall, and the human body.
- We analyze the ranging performance of 3db Access devices on different channels (at 6.5, 7, and 7.5 GHz) and propose channel diversity strategies that can improve the ranging accuracy.
- We implement localization systems based on the two types of devices and evaluate their performance experimentally in both LOS and NLOS settings.
- We provide statistical models for the LOS/NLOS ranging errors of 3db devices and evaluate their performance in a simulated building deployment when anchors are either in the same room or in adjacent rooms separated by a gypsum or a concrete wall.

The rest of the paper is organized as follows. In Section II, we analyze the theoretical differences between LRP and HRP PHYs and introduce the basics of UWB ranging and localization. We present the experimental setup in Section III and evaluate the power consumption, distance measurement errors, and localization performance of the ICs based on measurements in Section IV. In Section V, we model the ranging errors obtained experimentally and show how they can be used to simulate a localization application. In Section VI, we present the state-of-the-art in UWB localization and propose several directions for future work. Finally, we draw the conclusions in Section VII.

II. BACKGROUND

In Section II-A, we first introduce UWB devices and the main types of receiver architectures used in commercial devices. UWB devices can perform ToF measurements with sub-nanosecond accuracy and can therefore measure the distance between two devices with centimeter-level accuracy. Distances between two devices have value in themselves (e.g., to find lost objects) but also as a first step in multilateration algorithms for localization. In this paper, we compare both the ranging and localization performance of two types of UWB devices. We introduce ranging and localization concepts with UWB devices in Section II-B and II-C, respectively.

³<https://doi.org/10.5281/zenodo.4686379>

A. UWB Device Architectures

The IEEE 802.15.4 standard [4] defines two types of physical interfaces with low and high pulse repetition frequency: LRP and HRP, respectively. The Decawave DW1000 UWB chip is compliant with the HRP PHY defined in the IEEE 802.15.4 standard [4]. It is perhaps the most widely-used UWB device, so we chose it to represent the HRP PHY class. Decawave has recently released a new UWB chipset, the DW3000 [20]. However, the new-generation chipsets are currently available only as engineering samples, which is why we focused on the old release. The 3db IC is compliant with the LRP PHY specified in the IEEE 802.15.4z amendment [21]. The chip is already being used for secure keyless car access but it has not been evaluated in high-accuracy applications yet.

UWB transmissions have to satisfy two constraints imposed by international regulations [22]: a maximum average power spectral density (PSD) of -41.3 dBm/MHz (averaged over 1 ms) and a maximum peak power spectral density of 0 dBm/50MHz. UWB devices can, therefore, transmit over a fixed period either few pulses at high power levels or many pulses with lower transmit power. The first situation falls under the LRP specification and is employed by 3db devices, while the latter is known as HRP and is used by Decawave. If optimally employed, both of these technologies benefit from an *equal* average transmitted RF energy.

Since the HRP PHY transmits individual pulses with lower energy than the LRP, the received pulse energy is also lower for the same path loss (same distance). Therefore, the HRP PHY needs more sophisticated techniques to extract weaker pulses from the receiver noise, typically performed with correlations over many samples. For this reason, Decawave devices use coherent receivers. Because 3db devices implement the LRP PHY, they can use a *non-coherent* receiver based on energy detection (ED) for signals modulated with binary frequency-shift keying (BFSK).

Coherent receivers use phase information in the detection process. They typically have low sensitivity to inter-symbol and co-user interference and benefit from the multipath diversity of the UWB channel [23]. At the same time, the receiver architecture demands high computational resources and hardware complexity [24]. For optimal reception, the coherent receiver needs to estimate the multipath delays, their complex-valued channel coefficients, and the pulse shape distortion [24]. A precise estimation of the carrier phase is crucial for recovering the baseband pulse since inaccuracies will result in signal power loss and crosstalk interference in signals modulated using phase-shift keying (PSK) [25]. For a carrier frequency of 8 GHz, a time shift of half of the pulse period flips the phase of the signal, so coherent UWB systems generally tolerate rotations only within $\pi/4$ of the signal phase (around 30 ps). These requirements increase the power consumption of coherent demodulators [23].

Non-coherent receivers estimate channel coefficients based on the envelope rather than on the phase and amplitude of the received signal, so they have lower synchronization constraints. The timing requirements of a non-coherent receiver are dependent only on the pulse envelope, which is related




Fig. 1. Message exchange in the single-sided two-way ranging.

to the pulse bandwidth. For instance, if the pulse bandwidth is 500 MHz, the non-coherent receiver needs to operate with a timing resolution of 1 ns and it does not need high RF carrier synchronization. Therefore, non-coherent receivers can be more energy-efficient than coherent ones but have a higher bit error probability [25]. Another disadvantage of the non-coherent architecture is that it cannot be used for precise angle-of-arrival (AoA) measurements with closely-spaced antennas.

B. Ranging Methods

The distance between two devices can be estimated based on the time of flight (ToF) of the signal. Using the transmission time (T_1) of the signal measured by the sender and the arrival time (T_2) at the receiver, we can compute the distance as [3]:

$$d = (T_2 - T_1) \cdot c, \quad (1)$$

where c is the speed of light and $T_p \triangleq T_2 - T_1$ is the propagation time of the signal. To accurately estimate the distance, the devices need to be tightly clock synchronized, as a small mismatch of 1 ns can introduce a distance error of around 30 cm. Because synchronizing the sender and the receiver is usually unfeasible in practice, more messages are exchanged in order to reduce such errors, such as in the single- or the double-sided two-way ranging (SS-TWR and DS-TWR, respectively).

The SS-TWR uses two messages per distance estimate, as shown in Fig. II-B. The propagation time is:

$$T_p = \frac{T_{round} - T_{proc}}{2}, \quad (2)$$

where T_{round} is the time spent in one message exchange and T_{proc} is the processing time on the responder side. It can be shown that the error in estimating T_p is [26]:

$$e_{T_p} = e_1 \cdot T_p + \frac{1}{2} T_{proc} (e_1 - e_2), \quad (3)$$

where e_1 and e_2 are the clock drift errors of the initiator and responder, respectively. The main source of errors in the SS-TWR are T_{proc} , which is in the range of hundreds of microseconds, and the clock drift, which can be up to ± 20 ppm in systems compliant with the IEEE 802.15.4 standard [4].

In the LRP PHY, a location-enhancing information postamble is introduced at the end of each message to estimate the clock drift error [4]. Besides, the processing time of LRP messages is shorter than the one of HRP. It is also more convenient to minimize the number of exchanged messages in the TWR since this reduces the time needed to obtain one distance measurement. Therefore, the SS-TWR is usually the method of choice for LRP devices.

The DS-TWR uses an additional message exchange to minimize clock drift errors. Although this is the ranging method typically employed in Decawave devices [10], the Decawave MDEK1001 kit that we used throughout this paper applies the SS-TWR [27] to improve the energy efficiency and reduce the air time. Although the HRP PHY does not include a postamble nor does it have a fixed processing time, the tags can estimate their clock drift with respect to the anchors based on periodic beacons transmitted by the anchors [27].

C. True-Range Multilateration

The true-range multilateration algorithm estimates the location of a mobile device (also called a tag) using distance measurements between the tag and fixed devices with known locations (also called anchors). The special case for 2D localization using three anchors is known as trilateration.

Let d_i denote the distance between anchor A_i and the tag, which can be written as:

$$d_i = \|\mathbf{x}_{A_i} - \mathbf{x}\| + v_i, \quad i = 1, \dots, N \quad (4)$$

where \mathbf{x} is the location of the tag, \mathbf{x}_{A_i} is the location of anchor A_i , and v_i is the measurement noise. The noise terms of all anchors are assumed independent.

In vector form, the measurement equation becomes:

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{v}, \quad (5)$$

where \mathbf{y} is the measurement vector (containing all measurements d_j , $j = 1, \dots, N$), \mathbf{v} the error vector, and h the vector-valued measurement function. The equation can be solved by the least-squared solution \mathbf{x}^* which minimizes $\|\mathbf{y} - h(\mathbf{x})\|$ [3].

Multiple algorithms for solving the nonlinear system of equations were compared in [28]. The regularized Gauss-Newton multilateration algorithm is an iterative algorithm which has a similar accuracy to several algorithms with closed-form solutions and a low computational complexity suitable for real-time applications [28]. For this reason, we used it to implement the localization systems evaluated in Section IV-E.

The algorithm needs an initial starting position \mathbf{x}_0 , which should be chosen as close as possible to the real location for a quick convergence. For the first iteration, the starting position can be set to the solution of a closed-form multilateration algorithm or to the latest location of the tag (if available). At each iteration k , the algorithm computes the Jacobian matrix:

$$J_k(\mathbf{x}) = \left[\frac{\mathbf{x}_{A_1} - \mathbf{x}}{\|\mathbf{x}_{A_1} - \mathbf{x}\|}, \dots, \frac{\mathbf{x}_{A_N} - \mathbf{x}}{\|\mathbf{x}_{A_N} - \mathbf{x}\|} \right]^T. \quad (6)$$

The solution at step $k + 1$ is $\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta\mathbf{x}$, where $\Delta\mathbf{x}_k$ is the least-squares solution to

$$-(\Sigma^{-\frac{1}{2}} J_k + c\mathbf{I})\Delta\mathbf{x}_k = (\Sigma^{-\frac{1}{2}}(h(\mathbf{x}_k) - d) + c(\mathbf{x} - \mathbf{x}_r)), \quad (7)$$

where \mathbf{x}_r is a regularization point taken as the mean of the anchors' coordinates and c is a regularization coefficient equal to the inverse of the standard deviation of a distribution centered at \mathbf{x}_r . The algorithm stops if the location increment is below a tolerance δ or if the algorithm reaches the maximum number of iterations.

III. EVALUATION SETUP

In the following, we describe the device setups used in the power, ranging, and localization measurements.

3db Access: The 3db chip is integrated into an Arduino shield on top of an Arduino M0 board. The communication between the chip and the host MCU is performed via SPI. We use the channel centered at 6.52 GHz and the peak data rate of 247 kb/s. The 10 dB bandwidth of a pulse is 380 MHz and, because pulse spectra partially overlap in BFSK modulation, the total system bandwidth is approximately 620 MHz. The packet duration is 400 μ s. The IC was configured to transmit at the maximum level of -43.86 dBm/MHz, so within UWB regulations [22].

Decawave: We use the Decawave MDEK1001 kit which includes the DW1000 UWB chip integrated into the DWM1001 module. The DWM1001 module also contains a Nordic Semiconductor nRF52832 BLE microprocessor mostly used for network communication and an STM LIS2DH12TR 3-axis motion detector. The kit's default PANS software supports only the mode 14 which uses channel 5 (at 6.49 GHz), a data rate of 6.8 Mb/s, a PRF of 64 MHz, and a preamble length of 128 symbols, corresponding to a packet length of 287 μ s. The default configuration is suitable for short-range communication. The devices have a 3 dB bandwidth of 499.2 MHz (equivalent to a 10 dB bandwidth of ≈ 662 MHz). In all the ranging measurements, one of the devices is configured as a tag in the low-power mode, while the other device is an initiating anchor.

In our previous work [1], we used Decawave devices integrated in the EVK1000 evaluation kit which allowed more configurations and used the DS-TWR. There, we decided to use the long-range mode (Mode 3), in order to attain a similar range as with 3db devices. Here, we favored the MDEK kit because it implements the SS-TWR which requires less message exchanges and is more energy efficient. The SS-TWR is also implemented by 3db devices, making the operation of the two devices similar.

IV. MEASUREMENT-BASED EVALUATION

In this section, we compare how 3db and Decawave devices compare in terms of power and energy consumption, coverage, distance measurement accuracy and precision, and localization performance. Distance measurements are important in themselves, for instance in proximity detection applications, but also because they are at the basis of true-range multilateration.

Section IV-A compares the power consumption of the two devices, Section IV-B their maximum range, and Section IV-C the accuracy and precision of their distance measurements. Section IV-D analyzes measurements acquired on multiple channels. In Section IV-E, we integrate the devices in localization systems and evaluate their performance.

A. Power Consumption

In this section, we compare the power and energy consumption of the two types of devices. Unfortunately, the Decawave MDEK1001 board allows for measuring the current consumption only of the DWM1001 module, which contains, besides

the DW1000 UWB chip, a BLE microprocessor, and a motion detector. We present the power consumption measurements of the DWM1001 module for the sake of completeness, but for the comparison between the two *chipsets*, we rely on the current consumption of the Decawave DW1000 UWB chip from the device datasheet [29]. We used the current consumption reported for mode 14 which is referenced to 3.3 V [29].

We measured the current consumption of the 3db chip and the DWM1001 module with a Keysight DC Power Analyzer. We isolated the most important modes, namely the idle, transmit (TX), and receive (RX) and computed their average current consumption. The input voltages of the 3db chip and the DWM1001 module were 1.25 V and 3.3 V, respectively.

Table I presents the average power consumption in each mode of the 3db IC, the DW1000 IC, and the DWM1001 module. As mentioned, for the comparison between the UWB chipsets, we rely on the current consumption of the Decawave IC provided in the *datasheet* [29] and on the *measured* current consumption of the 3db IC. Overall, the average power consumption of the 3db IC in the TX, RX, and idle mode is at least 9 times lower than the one of the Decawave IC. Only in the deep sleep mode the 3db chipset has $1.9\times$ higher power consumption than the Decawave chipset. Note that the average power consumption of the DW1000 chip in the *idle* mode is about 1.45 times higher than the one of 3db devices in the *receive* mode, also the most power-hungry state. The results suggest that, indeed, the LRP interface can be more power-efficient than the HRP one.

The power consumption profile is a starting point for evaluating the energy consumption of an UWB-based localization system. A key challenge in a localization system is minimizing the energy consumption of the *tag*, which is usually battery-powered. To avoid synchronizing the tag and the anchors, the tag can *initiate* the message exchange and stay in the idle or sleep mode between rangings. Using the SS-TWR implies, in this case, that the *tag* estimates the distance (or the location).

To illustrate the energy efficiency of a tag in a localization system, let us consider the most favorable scenario in which the tag is the initiator. We disregard the time spent in the idle mode, since it is subject to the desired location update rate and guard times, which can be chosen freely to a certain extent. We therefore compute the energy consumption only when the device is in the TX or RX mode. The packet duration of the DW1000 chip in Mode 14 is $287\mu\text{s}$ and the one of the 3db chip is $400\mu\text{s}$. Therefore, a Decawave tag will consume $180\mu\text{J}$ per SS-TWR during transmission and reception, while a 3db tag will consume $28\mu\text{J}$ (including the transition times), so 6.4 times less energy. When placed in the long-range mode (for instance, mode 3), the packet duration of Decawave devices increases to $3487\mu\text{s}$ which is about $10\times$ larger than that of 3db devices, causing them to consume at least $100\times$ more energy [1]. The difference between these modes is the *maximum range* at which the devices can communicate, so in the next section we will compare the range of 3db and Decawave devices.

TABLE I
THE AVERAGE POWER CONSUMPTION OF 3DB AND DECAWAVE DEVICES.

	Average power consumption [mW]			
	TX	RX	Idle	Deep sleep
3db Access [†]	20.7	40.7	6.6	$6.25 * 10^{-4}$ *
DW1000 [‡]	237.6*	392.7*	59.4*	$3.3 * 10^{-4}$ *
DWM1001 [‡]	297.7	507.2	47.9	3.9

[†] Referenced to 1.25 V.

[‡] Referenced to 3.3 V.

* Based on the device datasheet.

B. Range

In this section, we want to find the ratio of successful distance measurements between a transmitter (TX) and a receiver (RX) placed at distances between 5–220 m. Remember that one distance measurement using the SS-TWR involves the successful transmission of two messages, a poll (from TX to RX) and a response (from RX to TX). The devices are said to have a (maximum) range of d meters when the ratio between the number of successful distance measurements and the total number of initiated measurements up to the distance d is higher than a chosen ratio $P = 0.9$. We performed measurements outdoors, on the pathwalk shown in Fig. 2, in order to minimize the multipath interference from surrounding objects which is usually higher indoors. At discrete steps, the TX was programmed to send 60 messages (polls) every 200 ms. If the response from the RX does not arrive at the TX either because the RX did not receive the probe or because the response was lost, a timeout occurs and the distance measurement is unsuccessful. We define the packet delivery ratio (PDR) as the number of responses received by the TX over the number of transmitted messages.

The Decawave PANS software reports only the (successful) responses and produces no output when transmitted packets are not answered. To compute the PDR of Decawave devices, we use the transmission period of 200 ms to compute how many messages should have been exchanged between the first and the last successful message at every test point. The PDR is then the number of received messages during that period divided by the number of expected messages. 3db devices report when packets are unanswered and we compute the PDR of 3db devices as described before.

Fig. 3 shows the PDR for 3db and Decawave devices. It is important to note that the PDR is highly dependent on the orientation of the devices, since at long distances the irregular radiation pattern of the antennas can cause high packet losses along certain directions. The PDR of Decawave devices dropped to 0 after 25 m, which is more than half of the expected range of 80 m of Mode 14 reported in the DW1000 Datasheet [29] (Section 6.3). However, this range was provided for channel 2 at 4 GHz, so the path loss is expected to be higher (and hence the range lower) at the center frequency of 6.5 GHz used in our experiment. 3db devices have a PDR higher than 0.9 at almost all distances up to 194 m, except for the higher losses between 65–90 m. At those distances, we found that the PDR was highly influenced by the relative pose



Fig. 2. Location at which the range of the devices was measured. The devices were placed at distances between 5–220 m along the pathway.




Fig. 3. The packet delivery ratio of 3db Access and Decawave devices.

between the devices, most likely due to destructive multipath interference and the antenna radiation pattern.

In our previous work [1], we found that 3db devices had a PDR of 0.9 up to 116 m, but we did not measure the PDR at longer distances because of the limited space. Similarly, in this experiment, we did not measure the PDR beyond 220 m. In [1], we compared the maximum range of 3db devices with the one of Decawave devices operating in the *long-range* mode (Mode 3) and we found that they had a similar PDR over the covered area. Decawave devices in the long-range mode have a packet duration about $10\times$ larger than in the short-range mode and therefore also a higher energy consumption. In practice, this means that a 3db tag could operate over a similar area as a Decawave tag in the long-range mode but with $125\times$ less energy. A Decawave tag in the short-range mode will be more energy-efficient than in the long-range mode but more anchors will be needed to provide coverage over the same area.

C. Distance Measurements

In this section, we compare the distance measurements of 3db and Decawave devices acquired in identical settings. We considered four indoor settings: LOS inside a large office and NLOS caused by either a gypsum wall (12.5 cm thickness), a concrete wall (29 cm thickness), or a human body. Fig. 4 shows the settings. Decawave and 3db Access devices were placed at exactly the same locations and acquired an equal number of measurements on the 6.5 GHz channel at the same rate (every 0.6 s). In all ranging experiments from this section, the devices were calibrated to account for errors caused by

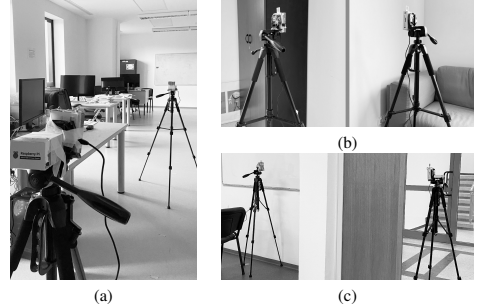


Fig. 4. Setup of ranging measurements in (a) LOS, (b) NLOS with a drywall, and (c) NLOS with a concrete wall. The UWB devices are placed on tripods. The NLOS with human body shadowing setup is identical with the LOS one, except that a person is standing right in front of the transmitter (the device further away).

TABLE II
SETUP OF RANGING EXPERIMENTS.

Scenario	Distances [m]	Sampling period [s]	Recording time per test point [min]
LOS	1, 2, ..., 8	0.6	2
NLOS with drywall	1, 2, ..., 6	0.6	4
NLOS with concrete wall	1, 2, 3	0.6	2
NLOS with human body	2, 5, 10	0.6	2

hardware, channel, or distance. The calibration method is described in Appendix A.

For the ranging datasets, at each test point, we recorded measurements for 2–4 min, which during the calibration phase was deemed enough to obtain a distribution with a mean error within ± 1 cm of the long-term one. The setup for each recording scenario is described in TABLE II. The 3db devices were configured to acquire measurements on all three channels at 6.5, 7, and 7.5 GHz, cycling through them every 0.2 s. Measurements on all channels will be later used in Section IV-D to investigate whether channel diversity improves the accuracy in certain situations. Because the MDEK1001 devices can use only the 6.5 GHz channel, we compare Decawave and 3db measurements acquired only on this channel.

TABLE III presents the mean, standard deviation, and inter-quartile range (IQR) of the distance errors of 3db and Decawave devices, computed as:

$$e_d = \hat{d} - d, \quad (8)$$

where \hat{d} is the measured and d is the true distance. Fig. 5a to 5d compare the PDF of the aggregated ranging errors at all distances for a particular LOS/NLOS scenario and Fig. 5e to 5h compare the error distributions of Decawave and 3db at each test point. The boxplots use Tukey's definition.

First, we notice that, at individual test points (Fig. 5e to 5h), the errors of Decawave devices have a smaller spread than those of 3db devices. After the calibration procedure detailed in Appendix A, 3db devices had distance errors of -0.05 ± 6.54 cm, while Decawave devices had errors of

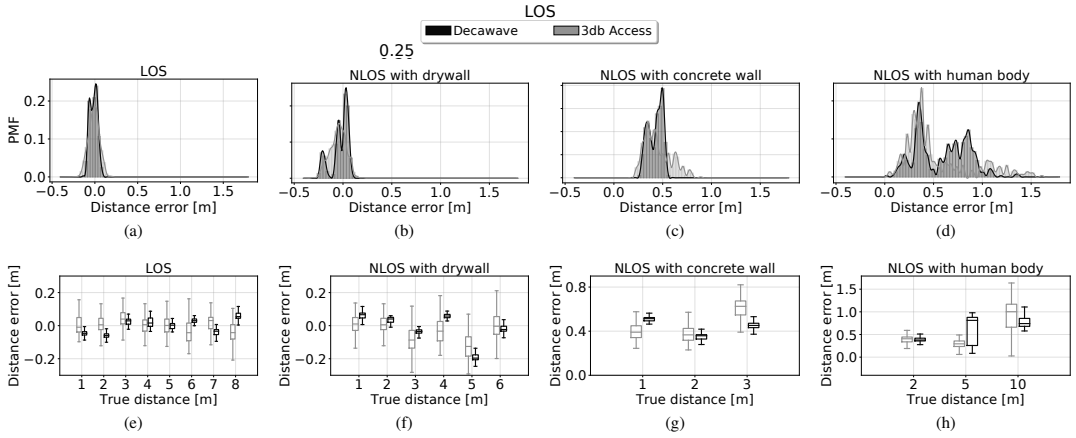


Fig. 5. Comparison of the distribution of aggregated ranging errors (Fig. 5a to 5d) and of the individual distributions at each distance (Fig. 5e to 5h) of 3db Access and Decawave devices in LOS, NLOS with drywall, NLOS with concrete wall, and NLOS with human body.

TABLE III
STATISTICS OF DISTANCE MEASUREMENT ERRORS.

Scenario	Device	Mean [m]	Standard deviation [m]	IQR [m]
LOS	3db Access	0.02	0.07	0.09
	Decawave	0.00	0.05	0.07
NLOS with drywall	3db Access	-0.04	0.08	0.12
	Decawave	-0.01	0.09	0.10
NLOS with concrete wall	3db Access	0.46	0.14	0.19
	Decawave	0.44	0.07	0.14
NLOS with human body	3db Access	0.55	0.32	0.29
	Decawave	0.60	0.26	0.46

0 ± 3.14 cm. Therefore, on the calibration data set, 3db devices had a bias 5 mm higher and a standard deviation about $2 \times$ larger than Decawave devices. In the LOS scenario, the location and test points were different from the ones in the calibration data set. Hence, we expect errors to be slightly higher than in the calibration setup. Over all test points, 3db Access devices had errors of 2 ± 7 cm and Decawave devices achieved errors of 0 ± 5 cm in the LOS scenario. The standard deviation of Decawave devices in the LOS scenario is higher than in the calibration data set because, at individual test points, the absolute average error is also higher.

Drywall is frequently used in modern buildings to delimit interior spaces. Surprisingly, it does not seem to cause a positive bias but a small negative one in both 3db and Decawave measurements, as can be seen in Fig. 5b. The errors caused by this type of obstruction are within several centimeters of LOS errors. This type of NLOS scenario is sometimes referred to in the literature as “soft” NLOS [30], since the LOS multipath component is still present in the CIR and the correct distance can be recovered. The fact that drywall does not introduce large errors is good news for proximity-detection and localization applications, because it means that ranging and localization errors will be small even if the devices

are in different rooms, if they are separated by drywall.

Thicker obstacles such as a wall or the human body can affect the signal in multiple ways. First, through this type of obstacles, the signals usually travel at a lower speed than through the air, which causes a delay in the round-trip time and hence an error in the distance measurement. Second, these obstacles can attenuate the direct path component or block it altogether, case in which copies of the signal reflected on surrounding objects can cause errors in the TOA estimation algorithm. These scenarios are also known as “hard” NLOS [30].

The aggregated distribution of ranging errors in hard NLOS scenarios (Figures 5c and 5d) is often heavy-tailed and no longer Gaussian-shaped. However, in most cases, the error distribution at each test point (i.e., at individual distances) is still approximately Gaussian, as shown in each boxplot from Fig. 5g and 5h. The biases depend on the particular environment and the multipath components that arrive at the receiver. Hence, at different distances, the bias can vary depending on how multipath components add up, which is why the aggregated NLOS distributions can be multi-modal.

With concrete wall and human body shadowing, the ranging errors are between 44–60 cm. In both hard NLOS scenarios, Decawave devices had a standard deviation 6–7 cm lower than 3db devices. Only with human body shadowing the IQR of Decawave errors is 17 cm higher than that of 3db because its error distribution, although shorter, has a fatter tail.

In conclusion, in all scenarios, both devices had mean errors within 2–5 cm of each other, with Decawave devices performing better in all scenarios except for the NLOS with human body shadowing one. In LOS and soft NLOS scenarios, the devices differed in the standard deviation and IQR by 1–2 cm, with Decawave devices obtaining a better performance in most cases. In the hard NLOS scenarios, Decawave devices had a lower spread than 3db devices by $1.23\text{--}2 \times$, except for the IQR in the case with NLOS with human body, which was $1.59 \times$ higher than the one of 3db devices. At individual test points, Decawave devices had $2 \times$ lower spread than 3db devices. Compared to our previous work [1], the performance




Fig. 6. Impact of channel diversity on distance accuracy when the signal passes through a concrete wall: at different distances, some channels perform better than the others.

of 3db devices has improved thanks to the refined calibration and to the correction of the firmware issues that previously caused large outliers in certain NLOS situations.

D. Channel Diversity

The analysis so far was based only on distance measurements acquired on the 6.5 GHz channel, since this was the only one available on the MDEK1001 devices. However, UWB devices can operate in more bands. The WiMedia Alliance defined 14 bands with 500 MHz bandwidth in the range of 3.1–10.6 GHz for UWB communications⁴. The use of the lower band between 3.5–4.5 GHz is often allowed only with interference mitigation techniques, while the 6–8.5 GHz band is less subject to regulations [31] and available in most countries. Since 3db Access devices can operate in the bands centered at 6.5, 7, and 7.5 GHz, it is useful to compare the performance on these channels and investigate whether distance measurements could benefit from channel diversity.

We programmed the 3db Access devices to acquire measurements on the 6.5, 7, and 7.5 GHz channels consecutively. The sampling period between measurements on successive channels is $T/3$, where T is the sampling period from TABLE II.

We noticed that, while LOS errors have the same characteristics irrespective of the channel (we calibrated the devices to operate this way), in hard NLOS situations some channels can experience better conditions at different locations. Fig. 6 presents such an example for NLOS with concrete wall shadowing: at 2 m distance, the 7 GHz channel has lower errors than the others, while at 3 m distance the 7 and 7.5 GHz channels had the highest accuracy. This can happen due to multipath interference, when copies of the signal traveling through multiple paths add up constructively or destructively at the receiver. The interference pattern depends on the frequency of the signal. Signals sent on different frequencies can have different propagation characteristics through obstacles. Since hard NLOS situations almost always cause positive biases, as we saw in the previous section, this prompts us to investigate whether taking the minimum or the mean of consecutive measurements (also called the min- and mean-select methods, respectively) acquired on different channels can improve the ranging accuracy.

⁴We will alternatively refer to the bands as (communication) channels.

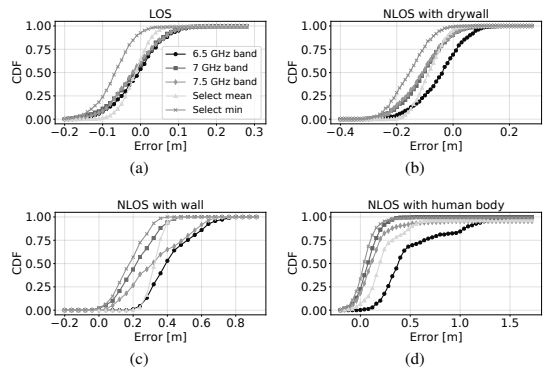


Fig. 7. The cumulative distribution function (CDF) of 3db Access ranging errors when using only measurements from the 6.5, 7, and 7.5 GHz band or the mean, median, or minimum of a set of consecutive measurements in all bands. The figure compares the CDF in (a) LOS, (b) NLOS with drywall, (c) NLOS with concrete wall, and (d) NLOS with human body. The legend in Fig. 7a is common to all subfigures.

Fig. 7 shows the cumulative distribution function (CDF) of measurement errors in the individual bands, as well as of errors when we select either the mean or minimum of three consecutive measurements acquired on all channels. In a regular LOS scenario, all channels perform similarly. The mean-select method leaves the bias still centered around 0 and decreases the standard deviation with approximately 2 cm. Instead, selecting the minimum measurement in LOS shifts the error distribution towards a negative mean, which decreases the accuracy. The same happens in the NLOS with drywall case.

In hard NLOS (Fig. 7c and 7d), however, the min-select method achieves a median error of approximately 18 cm with wall and 5 cm with human obstructions. Therefore, this simple channel diversity technique reduces the bias of hard NLOS measurements by more than 2× compared to using only the 6.5 GHz channel. The mean-select method also reduces the error compared to individual channels in “bad” conditions but to a lesser degree than the min-select.

It is not always desirable to use channel diversity. First, acquiring measurements on all channels increases the number of messages and thus the energy consumption. Second, the min-select method increased the accuracy by 2× in hard NLOS but decreased it in LOS. One method to take full advantage of channel diversity is to apply a NLOS detection technique [32] and acquire measurements on all channels only when the devices are in NLOS. In this way, additional measurements are triggered only when a higher ranging or localization accuracy is desired. Investigating the viability and efficiency of this method is left as future work.

E. Localization

Because UWB devices have become popular for indoor localization, in this section we compare the localization performance of the Decawave and 3db Access devices experimentally. We placed four anchors over an area of approximately 4.5×3.6 m shown in Fig. 8 and at heights between 1.2–1.8 m

TABLE IV
STATISTICS OF DISTANCE ERRORS IN LOCALIZATION RECORDINGS.

Case	Device	Mean distance error [cm]				Standard deviation of distance error [cm]			
		A1	A2	A3	A4	A1	A2	A3	A4
LOS	3db Access	7.4	1.1	12.3	8.6	6.8	6.7	11.3	6.9
	Decawave	6.3	1.5	14.6	7.8	9.1	6.7	13.2	14.0
NLOS at A4	3db Access	3.9	3.5	17.8	26.9	6.6	6.1	9.1	16.8
	Decawave	3.3	3.4	16.0	32.8	9.5	6.7	12.5	41.7



Fig. 8. Localization setup: the four anchors (A1 to A4) encompass an area of approximately 4.5×3.6 m and the tracking area is on the table. The ground truth was acquired using two HTC Vive base stations (BS1 and BS2) and a tracker that was colocated with the UWB tag, shown on the table.

(the anchors need to be at arm's length to acquire their ground truth location). We moved the tag by hand, at a height of approximately 40 cm above the table, along predefined points marked on the table. Since the tag was moved by hand, the paths in the two recordings were not identical, but very similar nevertheless, as can be seen from Fig. 9. This discrepancy should not significantly impact the comparison. During the first half of the trajectory, we oriented the tag towards anchors A1 and A4, while in the second half we oriented it towards anchors A2 and A3. We changed the orientation in order to vary the relative pose between the tag and the anchors, which can influence the ranging error [15].

The tags initiated the SS-TWR to each anchor. The 3db tag performed distance measurements to anchors A1 to A4 in their index order. The order in which the Decawave tag interrogated anchors changed throughout the recording according to the proprietary localization algorithm of Decawave.

Ground truth locations were acquired by an HTC Vive motion capture system using the setup described in [33], which has an average accuracy of at least 5 mm. The HTC Vive returns the location of a tracker which is colocated with the UWB tag. Then, a set of transformations is applied to recover the ground truth location of the tag. The anchor locations are also acquired using the HTC Vive system.

We recorded measurements in two scenarios: one in which all anchors were in LOS with the tag and one in which the direct path to one of the anchors (A4) was blocked by a person, so the tag was at all times in NLOS with one anchor. Given the results from Section IV-C, we expect a higher bias in the measurements coming from anchor A4 but not necessarily the same bias from TABLE III, since the bias depends also on the

particular room setup and environment. The NLOS distance error will introduce a *localization* error, which can be partially compensated by the correct distances received from the other anchors. Even in LOS, distance measurements can be affected by orientation errors caused by the irregular antenna radiation pattern [15].

During each recording, for each type of localization system (based on Decawave or 3db Access), we recorded the distances between each anchor and the tag, which were then given as input to a multilateration algorithm. As mentioned in Section II-C, for both localization systems, we used the Gauss-Newton multilateration algorithm strengthened with a regularization term. We initialized the algorithm with $\delta = 1$ mm, $k_{max} = 10$ iterations, $\mathbf{x}_r =$ the median of the anchors' locations, and $c = 10^{-1}$ (corresponding to a standard deviation of 10 m around \mathbf{x}_r , suitable for our setup). Although the Decawave MDEK1001 kit has its own localization algorithm, we did not use it for the comparison since the algorithm is closed-source and therefore we could not apply it on the anchor-tag distances given by 3db Access devices.

TABLE IV presents the mean and standard deviation of distance errors between the tag and each anchor A_i :

$$e_d = \hat{d}_{ij} - d_{ij}. \quad (9)$$

The true distance d_{ij} is computed as the Euclidean distance between the location of anchor A_i and each ground truth location of the tag, while \hat{d}_{ij} is the measured distance between each anchor and the tag.

The average difference between the mean distance error of the two devices is 1.15 cm and the average difference between the standard deviation of the distance error is 2.82 cm, with Decawave devices having smaller errors in about 50% of the cases. As previously mentioned, in the *localization* LOS scenario, the average bias is no longer null because of the changing orientation and the movement of the tag. All anchors have average errors under 10 cm when they are in LOS with the tag, with the exception of A3. Anchor A3 has higher distance errors than the others because of its close proximity to the concrete structure visible in Fig. 8. Although the structure does not completely obstruct the direct path between the tag and A3, it might cause the diffraction of the signal or other multipath effects, especially when the tag is close to the base station BS1 (the distance errors are higher in that area).

In NLOS, the effect of the human body shadowing is reflected in the distance error statistics of anchor A4. Similar to the ranging experiment, 3db Access devices have a mean ranging error lower by 5.9 cm than Decawave devices. However,

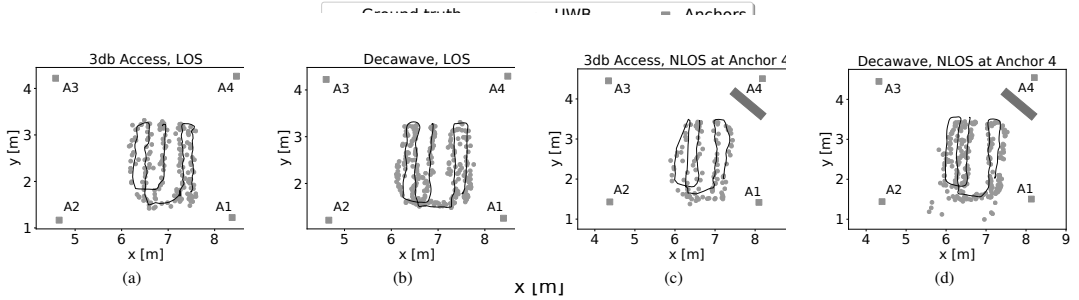


Fig. 9. Localization errors of 3db Access and Decawave devices in LOS and NLOS, where anchor 2 (A2) was obstructed by a human body.

TABLE V
STATISTICS OF MEASUREMENT-BASED LOCALIZATION ERRORS.

Case	Device	Localization error			
		Mean [cm]		Standard deviation [cm]	
		2D	3D	2D	3D
LOS	3db Access	9.7	36.8	5.4	14.8
	Decawave	12.4	42.7	7.8	19.1
NLOS	3db Access	18.9	72.2	10.5	34.1
	Decawave	22.9	89.1	23.5	43.7




Fig. 10. The CDF of 2D localization errors obtained from real measurements with 3db and Decawave devices.

unlike in the ranging experiment, here the standard deviation of the Decawave distance errors in NLOS (i.e., between the tag and anchor A4) was $2.5\times$ higher than the one of 3db Access devices. It is worth noting, though, that even in the ranging experiment in NLOS with human body, 3db Access devices had $1.59\times$ lower IQR than Decawave devices.

We compute the *localization* error as the Euclidean distance between the true location and the estimated one, either in 2D or 3D. For the 3D case, the error is:

$$e = \sqrt{(x - \hat{x})^2 + (y - \hat{y})^2 + (z - \hat{z})^2}, \quad (10)$$

where (x, y, z) are the Cartesian coordinates of the true location and $(\hat{x}, \hat{y}, \hat{z})$ are the Cartesian coordinates of the estimated location. For the 2D case, only the x and y coordinates corresponding to the plane parallel to the ground are taken into account.

The localization algorithm always computes the location in 3D but we analyze the 2D and 3D errors separately because the selected multilateration algorithm is prone to large errors

on the z axis (height). This happens especially when measurements are noisy since the geometric dilution of precision (GDOP) on the z axis is large. Therefore, it is important to distinguish errors in the 2D plan parallel to the ground, which usually need to be the smallest, from 3D errors.

TABLE V presents the mean and standard deviation of the localization error of both devices in 2D and 3D. Fig. 10 shows the CDF of the localization error of the devices only in 2D. In LOS, the two devices have the mean and standard deviation of localization errors within at most 5.9 cm of each other, leading to a similar performance. In the NLOS scenario, the mean and standard deviation of the localization errors of the Decawave-based localization system are higher by 4–16.9 cm than 3db's because of the higher distance errors between the tag and the NLOS anchor (A4). In 3D, the average localization errors are significantly higher than in 2D, with about 30 cm in LOS and 60 cm in NLOS, due to the measurement noise and high GDOP on the z axis. In the 2D case, 90% of the LOS errors are under 20 cm, while in NLOS 75% of the localization errors are under the same threshold.

V. ERROR MODELING AND SIMULATION

The localization results in Section IV-E are useful for comparing the two types of devices and for providing an estimate of the expected localization error in a small setup. We are now interested in evaluating the expected performance of a localization system that would be deployed on a larger scale (e.g., on the entire floor of an office building, in a shopping center, in a home). In such a setup, we expect a larger distance between anchors and possibly a lower anchor density than in our small-scale experiment. Moreover, while most of the existing literature assumes all anchors to be in the same room and preferably in LOS with the tag for the best accuracy, we argue that in real deployments such a requirement would be too costly in terms of price, setup effort, and maintenance. In such cases, administrators might prefer a localization system with lower accuracy but also lower setup costs.

In this section, we model the ranging errors obtained with 3db devices in the same LOS and NLOS scenarios as in Section IV-C, but with augmented data sets. We argue that the proposed statistical models can be used to simulate realistic localization scenarios that would otherwise take days or weeks to implement and evaluate. Section V-A describes the proposed

TABLE VI
DATABASE OF MEASUREMENTS WITH 3DB ACCESS DEVICES

Scenario	Distances [m]	Nr. locations	Nr. meas.
LOS	1, ..., 8	3	18 344
NLOS with drywall	1, ..., 6	2	6600
NLOS with concrete wall	1, ..., 10.5	5	14 868
NLOS with human body	1, ..., 10	3	8144

models and explains why a customized approach is needed to model errors obtained with different types of obstructions. In Section V-B, we evaluate the localization error of a simulated localization system based on 3db Access devices, where the anchors are in different rooms, which would be characteristic for a low-cost deployment. We also analyze the effect that different types of walls (made of concrete or gypsum) have on the localization error.

A. Error Modeling

In this section, we propose error models for distance measurements acquired in the same scenarios from Section IV-C: LOS, NLOS with drywall, NLOS with a concrete wall, and NLOS with human body shadowing. For more statistically significant results, we augmented the former data sets with more measurements in at least two locations, so that the results are not biased by the multipath profile of a single room. TABLE VI lists the number of measurements from each scenario, the range of distances covered, and at how many locations we acquired measurements. The NLOS with concrete wall dataset includes measurements performed through walls with thickness between 25.5–67.5 cm. For the LOS and NLOS with human body scenarios, we also included the measurements of anchors in LOS and NLOS, respectively, with the tag.

To model errors, we fitted some of the most well-known continuous distributions (the complete list is available at [34]) and computed the sum of squared errors (SSE) between the empirical PDF (\hat{y}_i) and each fitted PDF at discrete points i :

$$SSE = \sum_i (y_i - \hat{y}_i)^2. \quad (14)$$

The best parameters $\hat{\theta}$ for a specific probability distribution $p(x)$ are found by maximizing a likelihood function:

$$\hat{\theta} = \arg \max_{\theta} p(x|\theta), \quad (15)$$

over the entire parameter space. In other words, maximum likelihood estimation (MLE) selects the parameters under which the observed data is the most probable. The returned parameters are not guaranteed to be globally optimal. Where necessary, we provided good initial guesses for the optimization to improve the fit.

We present the distributions that minimized the SSE and their parameters obtained through MLE. If there were more distributions that achieved similar SSEs, we chose the most well-known and studied distributions. TABLE VII shows the distributions that best fit experimental data obtained in the four scenarios, their PDF, as well as the parameters of the best fit.

As illustrated in Fig. 11a, LOS errors can be modeled with a Gaussian distribution with a mean of 0.4 cm and a standard deviation of 7.1 cm, whose PDF is given in Eq. (11) from TABLE VII. The calibration process presented in Appendix A removed biases caused by different channels, hardware, and distances such that in regular LOS conditions distance measurement errors are approximately centered around 0 m.

Errors obtained in NLOS with a gypsum wall between the devices can also be modeled by a Gaussian distribution, as shown in Fig. 11b. Although errors obtained by 3db Access devices on the 6.5 GHz channel through this obstacle had a left-sided tail (see Fig. 5a), when aggregating data from all channels and from an additional experiment, the errors converge to a Gaussian with a bias of -4.3 cm and a slightly larger standard deviation than in LOS, of 9.2 cm. It is not clear why the gypsum wall causes negative biases. Its relative permittivity was found to be between 2.7–3.1 [35], higher than the relative permittivity of air, so the signal should travel at a lower speed through the obstacle, causing a delay. Since this delay is not reflected in the measurements, additional investigation is needed to determine whether other environmental factors are responsible for this bias. The main take-away is that gypsum walls introduce errors comparable to LOS propagation.

The scenarios where two devices are in NLOS with a concrete wall or with human body shadowing can be categorized as “hard” NLOS scenarios and introduce larger errors with heavier tails, as can be seen from Fig. 11c and 11d, respectively. We obtained the best fits for the Burr distributions type XII, also known as the Singh–Maddala distribution (Eq. (12)). The Burr type XII is part of the family of log-logistic distributions used to model data that increases in an initial phase and then decreases, such as wealth distribution, survival analysis, or mortality rate [36]. Its shape is similar to the more well-known log-normal distribution but can better handle heavier tails [36], as it is currently the case with our hard NLOS data. For the sake of completeness, we presented both the Burr and log-normal fits for the NLOS with concrete wall and human body shadowing scenarios, but also because the log-normal distribution has fewer parameters and is easier to interpret. The log-normal distribution has also been previously used in literature to model NLOS scenarios [37]. The PDF of the log-normal function is given in Eq. (13) with a parametrization in which s , μ , and σ are also known as the shape, location, and scale parameters, respectively. Given that Burr distributions are a better fit than the log-normal one for our experimental data, this suggests that such NLOS obstructions might introduce heavier tails than previously thought, especially in the case of human body shadowing where the Burr type III is a noticeably better fit than the log-normal (Fig. 11d).

The chosen distributions can be used to simulate different localization scenarios. Our analysis shows that different types of obstructions can introduce very different errors and that a one-size-fit-all error model for NLOS propagation would likely lead to unrealistic results. Therefore, when evaluating the expected localization accuracy of a particular setup, different error models should be taken into account depending, for instance, on the crowdedness of the room or its wall structure.




Fig. 11. The histograms of distance errors based on data from the ranging experiment in LOS, NLOS with drywall, NLOS with concrete wall, and NLOS with human body, and the distributions that best model them.

TABLE VII
PARAMETERS OF THE DISTRIBUTIONS USED TO MODEL DISTANCE ERRORS.

Scenario	Distribution	PDF	Shape	μ	σ	
LOS	Gaussian	$f(x \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$	(11)	–	0.004	0.071
NLOS drywall	Gaussian	See Eq. (11)	–	–0.043	0.092	
NLOS concrete wall	Burr XII	$f(x c, d, \mu, \sigma) = cd \left(\frac{x-\mu}{\sigma}\right)^{c-1} / \left(1 + \left(\frac{x-\mu}{\sigma}\right)^c\right)^{d+1}, x \geq 0, c, d > 0$	(12)	$c = 9.64$ $d = 0.98$	–0.46	0.72
	Log-normal	$f(x s, \mu, \sigma) = \frac{1}{s(x-\mu)\sqrt{2\pi}} \exp\left(-\frac{\ln^2\left(\frac{x-\mu}{\sigma}\right)}{2s^2}\right), x > 0, s > 0$	(13)	$s = 0.17$	–0.53	0.81
NLOS human	Burr XII	See Eq. (12)	$c = 32.84$ $d = 0.24$	–1.63	1.66	
	Log-normal	See Eq. (13)	$s = 0.44$	–0.30	0.50	




Fig. 12. Simulation setup with four anchors. We consider a LOS case, in which there is no wall between the anchors, and two NLOS cases, in which the anchors are separated by a gypsum or concrete wall shown in the figure. We consider locations of the tag spread uniformly within the tracking area, in steps of 0.25 m (the figure shows steps of 0.5 m for better visibility).

B. Building Deployment

In this section, we illustrate an example in which the proposed statistical models can be used to evaluate the expected localization accuracy. We consider the setup from Fig. 12, with four anchors placed in the corners of a space of 9×20 m, which is the area of our office space together with a meeting room. We consider a LOS scenario, in which there are no separating walls in the tracking area such that the anchors and the tag are at all times in LOS, and two NLOS scenarios when there




Fig. 13. The CDF of simulated localization errors with four anchors when the anchors are at all times in LOS with the tag or when the anchors are in different rooms delimited by a drywall or a concrete wall.

is either a concrete or a drywall separating the area into two rooms of 9×13 m and 9×7 m (corresponding to the office and meeting room, respectively). The anchors are placed close to the ceiling: two (opposing) anchors at a height of 3 m and the other two at 2.7 m. We consider all possible locations of the tag over the tracking area in steps of 0.25 m, similar to the points shown in Fig. 12. At each test point, we simulate the measured distance between the tag and each anchor by adding an error term to the true distance. The error term is sampled from the distributions from TABLE VII, based on the condition: LOS, NLOS with drywall, or NLOS with concrete wall. We run the simulation for each scenario five times and compute the error

statistics over the errors obtained in all runs.

The CDFs of the 2D localization errors obtained in the three scenarios are presented in Fig. 13. As expected based on the ranging experiments and error modeling, drywall does not introduce significant errors. The localization error when the space is separated by drywall is almost the same as if the wall were not there. The concrete wall, on the other hand, introduces a median 2D localization error of about 25 cm. In some cases, this might be an acceptable localization error, given that half as many anchors are needed for the entire space than if four anchors (the minimum necessary for 3D localization) were deployed in each room.

In a building with many rooms or delimitations (for instance, an office building), placing sets of four anchors in each room can be detrimental even if it improves the localization accuracy. First, devices in adjacent rooms (or even further away from each other) can be in communication range of each other so the anchors and users need to use different channels for communication and/or synchronize their transmissions in order to not interfere with each other. These issues are similar to challenges in the placement of base stations in cellular networks [38]. Second, the tag would also need to select a subset of surrounding anchors with which to range based on the link quality between them (for instance, using the techniques described in [39]) and based on the geometry formed by the anchors, since the localization accuracy is the highest within the convex hull determined by the anchors [40]. In addition, a higher number of deployed anchors increases the deployment costs. If the walls in a building introduce only small errors, these issues can be largely avoided by allowing walls within the tracking area encompassed by the minimum number of anchors. To establish whether this is the right approach for a given space, more work is needed, such as evaluating how two or more walls of different types influence the ranging error and developing a flexible, realistic simulator that outputs the best anchor configuration for a building plan while taking into account the cost–accuracy trade-off.

VI. DISCUSSION

A. Related work

1) *The Accuracy and Precision of UWB Ranging and Localization*: The reported NLOS bias and standard deviation of Decawave devices was 34 ± 35 cm in office environments with different types of obstacles [41], 15.6 ± 7.4 cm with human body shadowing [7], around 56 ± 20 cm when the path is blocked by concrete walls [9], and about 5 ± 15 cm with shallow obstructions (such as plasterboard, wood, or steel) [9]. The results for walls and shallow obstructions agree with our own observations but we obtained much larger errors with human body shadowing with both Decawave and 3db devices. In [16], an UWB localization system using Decawave devices achieved a localization accuracy of 0.21 m in 2D and 0.24 m in 3D in an industrial environment.

In [11], the authors compared the ranging and localization performance of three UWB devices which implement either the HRP PHY (Decawave DW1000), the LRP PHY (Ubisense), or a proprietary PHY (BeSpoon). Decawave

achieved the best average localization error of 0.5 m, followed by BeSpoon with 0.71 m, and Ubisense with 1.93 m. The article did not investigate the power and energy consumption of the devices. Ubisense devices have also been used in [12], where they achieved sub-15 cm average error in LOS and about 50 cm through a metallic enclosure. Another brand of LRP devices is PulsON (models P220 and the P400 series), but they implement a coherent interface. They have been used in [42] for detecting a human target through the wall, and integrated in localization systems in [17], [43] that yielded localization errors under 20 cm. PulsON and another brand of LRP devices from Multispectral Solutions Inc. (MSSI, now known as Zebra) were compared in [44]. PulsON and MSSI devices had average biases of 10 cm and 50 cm, respectively. The authors also noted the linear dependence on the errors with the distance and proposed a first-order linear model to correct them.

Custom non-coherent UWB transceivers were proposed in [18] in an FPGA implementation and in [45], [46] as integrated solutions, out of which only the former achieves errors lower than 20 cm. To the best of our knowledge, the solutions are not compliant with the IEEE 802.15.4 standard nor are they available commercially.

A method for using channel diversity to improve the ranging accuracy (but only in LOS) was proposed in [19]. The authors leverage the constructive interference phenomenon to reduce the number of measurements needed for a single distance measurement, which can also decrease the energy consumption when applying diversity methods.

In our work, we focused on a basic localization algorithm, namely the regularized Gauss-Newton trilateration, since we focused on comparing the localization performance of two different devices. More complex algorithms which leverage constraints common in localization problems can increase the accuracy even further. For instance, in [47], the authors proposed an optimization-based localization algorithm suitable for IoT devices and implemented it using Decawave UWB chipsets. The proposed method achieves 2–3× higher localization accuracy than standard trilateration methods.

2) *The Energy Efficiency of UWB devices*: To the best of our knowledge, we are the first to compare the power and energy consumption of devices implementing the LRP and HRP PHYs as defined in the IEEE 802.15.4 standard [4]. Energy efficiency is starting to be a concern in UWB devices and, recently, concurrent ranging has been proposed as a solution to minimize the energy consumption in localization tasks at the application level [48]. Using this method, a tag can compute its location based on the time difference of arrival between multiple anchors' responses which are processed in a *single* message (rather than $2 \times N$ messages, where N is the number of anchors, as in SS-TWR-based trilateration). In [49], the authors analyzed the energy efficiency of UWB localization systems depending on the association and synchronization demands. In the new DW3000 UWB chipset, Decawave claims to have improved the power consumption, but it is not yet clear by how much [20]. There is also a newly-released, improved version of the 3db 3DB6830C chip, namely the ATA8352 chip [50].

One option to decrease the energy consumption when the tag does not initiate the ranging is to use an ultra-low-power wake-up receiver to keep the tag active only before a message exchange and in the idle or deep sleep mode otherwise [51].

3) *NLOS Detection and Mitigation Techniques*: There is a large body of literature dedicated to NLOS detection and mitigation strategies [32]. Most works detect the NLOS condition using the statistical properties of the channel impulse response (CIR) of signals acquired in NLOS. These methods require knowledge of the CIR, which by default is not dumped by the device, and introduce additional processing times. Our open-source measurement database also includes the dumped CIRs of 3db Access devices, which could be used to evaluate the performance of these methods. It is noteworthy, however, that in our previous work [1], the multipath components in the CIRs of 3db Access devices were found to be wider than those of Decawave because of their lower pulse bandwidth. Therefore, algorithms based on CIR characteristics might not generalize well to all types of devices. Other works proposed NLOS detection and mitigation methods that do not require prior knowledge on LOS/NLOS statistics, for instance based on sparse pseudo-input Gaussian processes (SPGP) [52] or on fuzzy theory [53], which show promising results. Our measurement database contains also the CIRs of 3db devices, which can be used to test CIR-based NLOS detection and mitigation algorithms.

In [8], the authors proposed a method to detect and correct NLOS biases when the signal propagates through a wall based on the wall's relative permittivity and thickness, which reduces the NLOS error by 53%. In the future, we will investigate whether this model fits the data from our experiments with both concrete and gypsum walls.

4) *Comparison with other Localization Technologies*: The power consumption of GNSS modules ranges from 12–72 mW for super-low-power modules [54] to 160 mW in a typical smartphone [55]. The accuracy of GPS receivers is about 2.5 m for high-end receivers [56], 2–5 m in smartphones [57], and at least centimeter-level for real-time kinematic (RTK) GPS [58]. Therefore, UWB devices can provide significantly higher localization accuracy indoors than GNSS receivers without enhancements; in terms of power consumption, LRP devices are comparable to super-low-power GNSS receivers, while the average power consumption of HRP devices exceeds the one of GNSS receivers in smartphones. However, given that the reception times might vary between the two technologies, their energy consumption per location might also be different and needs to be assessed in the future.

Bluetooth Low-Energy (BLE) has been also used in localization applications in recent years. Its power and energy consumption is similar to the one of 3db devices [59] but its accuracy is at best 0.7–1 m with sensor fusion [60]. Wi-Fi also achieves decimeter-level accuracy [61] but has higher energy consumption [59] than both BLE and UWB.

B. Future Work

As phone manufacturers started including UWB chips in smartphones, one vision would be the instrumentation of

entire buildings with anchors to provide seamless positioning indoors. As briefly mentioned throughout the paper, there are still unsolved challenges to reach this goal. First, if anchors are placed in the same room in order to avoid NLOS measurements, buildings with many rooms or cubicles will need many localization cells (formed by the minimum number of anchors for 2D or 3D localization). If the cells are in range of each other, transmissions within multiple cells need to be synchronized or allocated to different bands, which increases the administration efforts. Moreover, since localization accuracy degrades at the edge of a cell, a hand-over protocol needs to be implemented at the tag to decide which anchors to select for localization at a given time. Allowing walls within one cell does not necessarily decrease the localization accuracy if the walls are shallow (such as gypsum walls) and can reduce the deployment effort and costs. However, more work is needed to model the ranging errors and maximum range through an arbitrary number of walls. Towards this end, a simulator for building deployments of UWB localization systems that recommends the optimal number of anchors and their placement for the desired accuracy–cost trade-off would be needed.

In this work, we considered only range-based localization, where each distance measurement is obtained through at least two message exchanges between the tag and each anchor. This scheme does not scale well with many users and anchors because of the large number of messages involved and the need to schedule uplink transmissions. Moreover, this method is privacy-sensitive since anchors have access to at least part of the ranging information, which can be used to track the user with a certain precision. Instead, a GPS-like localization system based on the time-difference of arrival (TDOA) of signals, in which anchors act as satellites and passive tags use their broadcasts to locate themselves can, in theory, scale to an unlimited number of users and is more privacy-friendly. The drawback is that, in this case, anchors need to synchronize their transmissions. Wired synchronization introduces the lowest errors but is unlikely to be adopted because of the high deployment costs. Wireless synchronization, on the other hand, leads to localization errors in the range of decimeters even when anchors are in LOS with the tag [62] and requires anchors to be in LOS with each other. More work is needed, for instance, to determine whether a calibration protocol can allow wireless-synchronized TDOA systems where anchors are in different rooms to obtain a similar accuracy with the case in which anchors are in LOS with each other.

VII. CONCLUSION

This paper provided an outlook on the power and energy consumption, distance measurement statistics, and localization performance of 3db Access and Decawave devices, representative of the two types of UWB physical interfaces, LRP and HRP, respectively. Both devices have ranging and localization errors on the same order of magnitude. Decawave devices generally show better performance in LOS and through-the-wall NLOS conditions, while 3db devices have slightly better performance in NLOS with human body shadowing. For a

similar maximum range, Decawave devices have $125\times$ higher energy consumption than 3db Access devices, while in the short-range mode (which decreases the range by at least $8\times$) they have $6.4\times$ higher energy consumption than 3db Access. Therefore, devices implementing the LRP PHY might be more suitable for ultra-low power applications, while the HRP PHY might be a better choice for the highest ranging accuracy.

We evaluated the performance of 3db Access and Decawave devices in multiple LOS and NLOS (caused by a person, dry-wall, or concrete wall) scenarios and provided models for the error distribution of 3db Access distance measurements. These models can be further used to simulate realistic deployments of localization systems which would otherwise take days or even weeks to evaluate. We illustrated their applicability by simulating a localization scenario in which the anchors are placed in different rooms separated by drywall or a concrete wall. Results suggest that drywall causes negligible errors and anchors do not need to be in the same room to obtain high localization accuracy.

More research is needed to evaluate the impact of multiple walls on the ranging accuracy and to create suitable models of the maximum range and distance measurement errors with an arbitrary number of walls. Using such models, a simulator of UWB localization systems could then recommend the minimum number of anchors for the desired localization accuracy and their placement. More work is also needed to select the best anchors at a particular moment based on their link to the tag and to synchronize UWB transmissions between clusters of anchors within range of each other.

APPENDIX DISTANCE CALIBRATION

Before comparing the ranging accuracy and precision of Decawave and 3db devices, a calibration step is also necessary since, as we will show, the raw distance error can depend on the hardware, on the channel used, and even on the distance.

We collected distance measurements between different device pairs placed 2, 5, and 10 m apart, for at least 30 min, which was deemed a long enough time period to capture the long-term distribution of distance measurements. In a different experiment, we acquired measurements during 24 h to obtain the long-term error distribution. The mean error obtained in windows longer than 2 min was within ± 1 cm of the long-term mean error. We chose a longer window, of 30 min, to obtain more stable distributions.

We recorded measurements on all the available channels (6.5 GHz, 7 GHz, and 7.5 GHz) using 3db devices and only on the 6.5 GHz channel with Decawave, since it is the only available channel when using the default Decawave software for the MDEK1001 kit. Although the software compensates the antenna delay, in most cases distance measurements with Decawave devices still had a non-negative bias that was eliminated in the calibration process.

Our findings suggest that the measurement bias of UWB devices varies with the distance, the channel, and the pair of devices used. The dependency on the distance for the same channel can be seen in Fig. 14 for the 6.5 GHz channel.

The TOA estimation error increases with lower SNR [63], so without proper calibration distance measurement errors increase with the measured distance. Fig. 14 also shows how the ranging error varies between channels because they have different amplitude saturation points. Fortunately, the channel is known at the time of the measurement and these errors can be compensated.

The mean error also depends on the pair of devices used. This can occur due to the different clock offset or antenna delay of the hardware devices. Although these parameters can be estimated, the clock offset can still deviate in time or the antenna delay calibration might be imperfect. Fig. 15 compares the error distributions of distance measurements acquired by different 3db and Decawave device pairs on the 6.5 GHz channel. All pairs have a common transmitter (the tag used for localization in Section IV-E). For distances up to 10 m, the mean error can vary with 0.6 m for 3db devices and 0.2 m for Decawave devices, but note that this range includes also the linear dependence of the distance error on the SNR.

An arbitrary non-negative measurement bias is not desirable in ranging or localization. Therefore, such errors are often eliminated during a calibration step. The desired measurement model in LOS is a zero-mean Gaussian, whose standard deviation is at least partly determined by the hardware capabilities. It is worth noting that while the mean distance error varies with up to 0.6 m when we change the hardware, channel, or distance, the standard deviation of each error distribution is almost constant for each device (see Table VIII). Since the raw error distribution is already Gaussian, we need to correct only its bias.

We consider each set of measurements s_{ij} for a given pair of devices $p_i, i = 1, \dots, 4$ and channel c_j , where $j = \{0, 1, 2\}$ corresponds, respectively, to channels at 6.5, 7, and 7.5 GHz. Each set s_{ij} contains an equal number of approximately 3000 measurements (equivalent to a recording time of 30 min) taken at each distance $d_k \in \{2, 5, 10\}$ m. We assume a simple linear dependence of the measured distance on the true distance, which is found by minimizing the squared error:

$$E = \sum_{n=0}^N |x_n \cdot p_0 + p_1 - y_n|^2, \quad (16)$$

where N is the total number of measurements, y_n is the measured distance at the true distance x_n for all $n = 1, \dots, N$ and p_0 and p_1 are the polynomial coefficients.

Once the polynomial coefficients for a set s_{ij} are computed, the measurements of that set can be corrected as follows:

$$x_c = \frac{x_m - p_1}{p_0}, \quad (17)$$

where x_c and x_m are, respectively, the calibrated and the raw measurements.

There are several caveats to this approach. First, we do not derive error coefficients for each factor that introduces errors (channel, pair, and true distance). Rather, the accumulated error is corrected for a particular set.

Second, while it is feasible to obtain the calibration coefficients for a particular channel by performing measurements at different distances, the device-dependent calibration is harder




Fig. 14. Comparison of the distribution of distance measurement errors of the same pair of 3db Access devices, on the 6.5, 7, and 7.5 GHz channels, at distances of 2, 5, and 10 m between the two devices. On the 6.5 GHz channel, the mean error increases approximately linearly with the distance, while on the other two channels it is constant with the distance.




Fig. 15. Comparison of the distribution of distance measurement errors of four pairs of (a) 3db Access and (b) Decawave devices on the 6.5 GHz channel. The same transmitter (tag) was used in all pairs and only the receiver (anchor) was changed. The error varies with both the distance and the hardware.

to perform in a real application. Localization systems deployed inside buildings should be able to offer accurate locations even to unknown users, so one could not calibrate each anchor-tag pair. Even if all the users were known, for instance in a privately-deployed localization system, the tag population can easily reach hundreds or thousands devices, which again makes pair-wise device calibration impractical. Hardware-dependent distance errors are an interesting research topic but outside the scope of this paper, which is why we calibrated each pair of Decawave and 3db Access devices used (including the tag and each anchor in the localization experiment from Section IV-E).

Third, it is worth noting that the distance errors of Decawave devices do not linearly increase with the measured distance, as in the case of 3db Access. Instead, the dependency seems non-linear. Due to the lack of a more appropriate straight-forward model, we still apply a linear fitting with the mention that there might be better models for this distribution. Even

TABLE VIII
DISTANCE ERROR STATISTICS BEFORE AND AFTER CALIBRATION.

Device	Case	Mean [cm]	Standard deviation [cm]
3db Access	Before calibration	-0.79	16.09
	After calibration	-0.05	6.54
Decawave	Before calibration	-7.56	4.79
	After calibration	0.00	3.14

so, the calibration reduces the measurement bias with 8 cm, as shown in TABLE VIII, which presents the mean and the standard deviation of all errors before and after calibration.

The calibration does not significantly reduce the *overall* average error of 3db Access devices because the sets with a positive bias balance out those with a negative bias (predominantly on channels 1 and 2 which, for brevity, are not shown), but there is a benefit in the bias reduction of the individual sets. Moreover, the calibration reduces the standard deviation of measurement errors by more than half.

ACKNOWLEDGMENT

The authors gratefully acknowledge funding from European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints, <http://www.a-wear.eu/>). This work does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of its content.

REFERENCES

- [1] L. Flueraoru, S. Wehrl, M. Magno, and D. Niculescu, "On the energy consumption and ranging accuracy of ultra-wideband physical interfaces," in *GLOBECOM 2020-2020 IEEE Global Communications Conference*, pp. 1-7, IEEE, 2020.
- [2] "UWB market outlook, eet asia." <https://www.eetasia.com/uwb-market-outlook/>. Accessed on 15-08-2020.
- [3] X. Shen, M. Guizani, R. C. Qiu, and T. Le-Ngoc, *Ultra-wideband wireless communications and networks*. John Wiley & Sons, 2007.
- [4] "IEEE 802.15.4 standard for low-rate wireless networks," 2015.
- [5] R. Hazra and A. Tyagi, "A survey on various coherent and non-coherent IR-UWB receivers," *Wireless Personal Communications*, vol. 79, no. 3, pp. 2339-2369, 2014.
- [6] G. Durisi and S. Benedetto, "Performance of coherent and noncoherent receivers for UWB communications," in *2004 IEEE International Conference on Communications (IEEE Cat. No. 04CH37577)*, vol. 6, pp. 3429-3433, IEEE, 2004.
- [7] Q. Tian, K. I.-K. Wang, and Z. Salcic, "Human body shadowing effect on UWB-based ranging system for pedestrian tracking," *IEEE Transactions on Instrumentation and Measurement*, vol. 68, pp. 4028-4037, Oct. 2019. Conference Name: IEEE Transactions on Instrumentation and Measurement.
- [8] B. J. Silva and G. P. Hancke, "Ranging error mitigation for through-the-wall non-line-of-sight conditions," *IEEE Transactions on Industrial Informatics*, pp. 1-1, 2020. Conference Name: IEEE Transactions on Industrial Informatics.
- [9] A. Schenck, E. Walsh, J. Reijniers, T. Ooijevaar, R. Yudanto, E. Hostens, W. Daems, and J. Steckel, "Information theoretic framework for the optimization of UWB localization systems," in *2018 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1-8, Sept. 2018. ISSN: 2471-917X.
- [10] "DW1000 user manual. Version 2.11." Decawave Ltd.

- [11] A. R. J. Ruiz and F. S. Granja, "Comparing Ubisense, Bespoon, and Decawave UWB location systems: Indoor performance analysis," *IEEE Transactions on Instrumentation and Measurement*, vol. 66, no. 8, pp. 2106–2117, 2017.
- [12] A. Shahi, A. Aryan, J. S. West, C. T. Haas, and R. C. Haas, "Deterioration of UWB positioning during construction," *Automation in Construction*, vol. 24, pp. 72–80, 2012.
- [13] "Qorvo completes acquisition of Decawave." <https://www.decawave.com/qorvo-completes-acquisition-of-decawave/>.
- [14] "3db Access." www.3db-access.com. Accessed on 15-08-2020.
- [15] A. Ledgergerber and R. D'Andrea, "Calibrating away inaccuracies in ultra wideband range measurements: A maximum likelihood approach," *IEEE Access*, vol. 6, pp. 78719–78730, 2018. IEEE Access.
- [16] M. Delamare, R. Boutteau, X. Savatier, and N. Iriart, "Static and dynamic evaluation of an UWB localization system for industrial applications," *Sci*, vol. 1, no. 3, p. 62, 2019.
- [17] S. Monica and G. Ferrari, "An experimental model for UWB distance measurements and its application to localization problems," in *2014 IEEE International Conference on Ultra-WideBand (ICUWB)*, pp. 297–302, IEEE, 2014.
- [18] M. Segura, H. Hashemi, C. Sisterna, and V. Mut, "Experimental demonstration of self-localized ultra wideband indoor mobile robot navigation system," in *2010 International Conference on Indoor Positioning and Indoor Navigation*, pp. 1–9, IEEE, 2010.
- [19] 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 ACM Conference on Embedded Network Sensor Systems CD-ROM, SenSys '16*, (Stanford, CA, USA), pp. 137–149, Association for Computing Machinery, Nov. 2016.
- [20] "Decawave DW3000 uwb chipset release." <https://www.decawave.com/product/dwm3000-module/>.
- [21] "IEEE 802.15.4z task group." <http://www.ieee802.org/15/pub/TG4z.html>. Accessed on 15-08-2020.
- [22] F. C. Commission *et al.*, "Revision of part 15 of the commission's rules regarding ultra-wideband transmission systems," *First Report and Order, FCC 02-48*, 2002.
- [23] M. Weisenhorn and W. Hirt, "Robust noncoherent receiver exploiting UWB channel properties," in *2004 International Workshop on Ultra Wideband Systems Joint with Conference on Ultra Wideband Systems and Technologies. Joint UWBS & IWUWBS 2004 (IEEE Cat. No. 04EX812)*, pp. 156–160, IEEE, 2004.
- [24] H. Arslan, Z. N. Chen, and M.-G. Di Benedetto, *Ultra wideband wireless communication*. John Wiley & Sons, 2006.
- [25] J. G. Proakis and M. Salehi, *Digital communications*, vol. 4. McGraw-Hill New York, 2001.
- [26] D. Neirynek, E. Luk, and M. McLaughlin, "An alternative double-sided two-way ranging method," in *2016 13th workshop on positioning, navigation and communications (WPNC)*, pp. 1–4, IEEE, 2016.
- [27] "DWM1001 system overview. Version 2.0." Decawave Ltd.
- [28] N. Sirola, "Closed-form algorithms in mobile positioning: Myths and misconceptions," in *2010 7th Workshop on Positioning, Navigation and Communication*, pp. 38–44, IEEE, 2010.
- [29] "DW1000 datasheet. Version 2.09." Decawave Ltd.
- [30] S. Venkatesh and R. Buehrer, "Non-line-of-sight identification in ultra-wideband systems based on received signal statistics," *IET Microwaves, Antennas & Propagation*, vol. 1, no. 6, pp. 1120–1130, 2007.
- [31] "Application note APR001. UWB regulations. A summary of worldwide telecommunications regulations governing the use of ultra-wideband radio." Decawave Ltd. Accessed: 15-08-2020.
- [32] J. Khodjaev, Y. Park, and A. Saeed Malik, "Survey of NLOS identification and error mitigation problems in UWB-based positioning algorithms for dense environments," *Annals of telecommunications*, vol. 65, pp. 301–311, June 2010.
- [33] L. Flueraoru, E. S. Lohan, J. Nurmi, and D. Niculescu, "HTC vive as a ground-truth system for anchor-based indoor localization," in *2020 12th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, pp. 214–221, IEEE, 2020.
- [34] "Continuous distributions available in the SciPy module." <https://docs.scipy.org/doc/scipy/reference/stats.html#continuous-distributions>.
- [35] S. S. Zhekov, O. Franek, and G. F. Pedersen, "Dielectric properties of common building materials for ultrawideband propagation studies [measurements corner]," *IEEE Antennas and Propagation Magazine*, vol. 62, no. 1, pp. 72–81, 2020.
- [36] C. Kleiber, *A Guide to the Dagum Distributions*, pp. 97–117. New York, NY: Springer New York, 2008.
- [37] A. F. Molisch, J. R. Foerster, and M. Pendergrass, "Channel models for ultrawideband personal area networks," *IEEE wireless communications*, vol. 10, no. 6, pp. 14–21, 2003.
- [38] J. G. Andrews, F. Baccelli, and R. K. Ganti, "A tractable approach to coverage and rate in cellular networks," *IEEE Transactions on communications*, vol. 59, no. 11, pp. 3122–3134, 2011.
- [39] Y. Chen, J. Yang, W. Trappe, and R. P. Martin, "Impact of anchor placement and anchor selection on localization accuracy," *Handbook of Position Location: Theory, Practice, and Advances*, pp. 425–455, 2011.
- [40] N. Salman, H. K. Maheshwari, A. H. Kemp, and M. Ghogho, "Effects of anchor placement on mean-CRB for localization," in *2011 The 10th IFIP Annual Mediterranean Ad Hoc Networking Workshop*, pp. 115–118, IEEE, 2011.
- [41] A. R. Jiménez and F. Seco, "Comparing Decawave and Bespoon UWB location systems: Indoor/outdoor performance analysis," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, IEEE, 2016.
- [42] A. Kumar, Z. Li, Q. Liang, B. Zhang, and X. Wu, "Experimental study of through-wall human detection using ultra wideband radar sensors," *Measurement*, vol. 47, pp. 869–879, 2014.
- [43] W. Yan, Z. Jing, and Z. Nailong, "The designing of indoor localization system based on self-organized WSN using PulsON UWB sensors," in *2015 2nd International Conference on Information Science and Control Engineering*, pp. 965–969, IEEE, 2015.
- [44] G. MacGougan, K. O'Keefe, and R. Klukas, "UWB ranging and ranging measurement accuracy," *Measurement Science and Technology*, vol. 20, no. 9, pp. 0957–0233, 2009.
- [45] L. Stoica, A. Rabbachin, and I. Oppermann, "A low-complexity noncoherent IR-UWB transceiver architecture with TOA estimation," *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 4, pp. 1637–1646, 2006.
- [46] G. Fischer, O. Klymenko, and D. Martynenko, "Time-of-arrival measurement extension to a non-coherent impulse radio UWB transceiver," in *2008 5th Workshop on Positioning, Navigation and Communication*, pp. 265–270, IEEE, 2008.
- [47] P. N. Beuchat, H. Hesse, A. Domahidi, and J. Lygeros, "Enabling optimization-based localization for iot devices," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5639–5650, 2019.
- [48] P. Corbalán and G. P. Picco, "Ultra-wideband concurrent ranging," *ACM Transactions on Sensor Networks (TOSN)*, vol. 16, no. 4, pp. 1–41, 2020.
- [49] J. Tiemann and C. Wietfeld, "Scalability, real-time capabilities, and energy efficiency in ultra-wideband localization," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 12, pp. 6313–6321, 2019.
- [50] "ATA8352 ultra-wideband (UWB) transceiver." <https://www.microchip.com/wwwproducts/en/ATA8352>. Accessed: 15-08-2020.
- [51] V. Niculescu, M. Magno, D. Palossi, and L. Benini, "An energy-efficient localization system for imprecisely positioned sensor nodes with flying UAVs," in *18th IEEE International Conference on Industrial Informatics Online Event (INDIN 2020)*, 2020.
- [52] X. Yang, "NLOS mitigation for UWB localization based on sparse pseudo-input gaussian process," *IEEE Sensors Journal*, vol. 18, no. 10, pp. 4311–4316, 2018.
- [53] K. Wen, K. Yu, and Y. Li, "NLOS identification and compensation for UWB ranging based on obstruction classification," in *2017 25th European Signal Processing Conference (EUSIPCO)*, pp. 2704–2708, IEEE, 2017.
- [54] "Ultra small super-low-power U-blox M8 GNSS system-in-package (SIP)." <https://www.u-blox.com/en/product/zoe-m8b-module>. Accessed on 15-08-2020.
- [55] A. Carroll, G. Heiser, *et al.*, "An analysis of power consumption in a smartphone," in *USENIX annual technical conference*, vol. 14, pp. 21–21, Boston, MA, 2010.
- [56] M. Rychlicki, Z. Kasprzyk, and A. Rosiński, "Analysis of accuracy and reliability of different types of GPS receivers," *Sensors*, vol. 20, no. 22, p. 6498, 2020.
- [57] C. Specht, P. Dabrowski, J. Pawelski, M. Specht, and T. Szot, "Comparative analysis of positioning accuracy of GNSS receivers of Samsung Galaxy smartphones in marine dynamic measurements," *Advances in Space Research*, vol. 63, no. 9, pp. 3018–3028, 2019.
- [58] Y. Feng, J. Wang, *et al.*, "GPS RTK performance characteristics and analysis," *Positioning*, vol. 1, no. 13, 2008.
- [59] M. Siekkinen, M. Hienkari, J. K. Nurminen, and J. Nieminen, "How low energy is Bluetooth Low Energy? Comparative measurements with ZigBee/802.15.4," in *2012 IEEE wireless communications and networking conference workshops (WCNCW)*, pp. 232–237, IEEE, 2012.

- [60] J. Röbesaat, P. Zhang, M. Abdelaal, and O. Theel, "An improved BLE indoor localization with Kalman-based fusion: An experimental study," *Sensors*, vol. 17, no. 5, p. 951, 2017.
- [61] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti, "Spotfi: Decimeter level localization using WiFi," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pp. 269–282, 2015.
- [62] 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 *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pp. 61–72, IEEE, 2019.
- [63] I. Guvenc and Z. Sahinoglu, "Threshold selection for UWB TOA estimation based on kurtosis analysis," *IEEE Communications Letters*, vol. 9, no. 12, pp. 1025–1027, 2005.



Laura Flueraștoru obtained her Master's degree in Electrical Engineering from ETH Zürich, Switzerland, in 2019, and her Bachelor's degree in Electronics and Telecommunications from University Politehnica of Bucharest, Romania, in 2017. She is currently pursuing a double Ph.D. degree at University Politehnica of Bucharest, Romania, and Tampere University, Finland, as a Marie Skłodowska-Curie Fellow in the European project A-WEAR. During her studies, she gained experience in both industry and research from internships at Freescale

Semiconductor, École Polytechnique Fédérale de Lausanne (EPFL), Schindler Group. Her research interests include indoor localization, ultra-wideband communications, wireless and mobile communications, embedded systems, signal processing, and machine learning.



Silvan Wehrli received the M.Sc. degree in electrical Engineering and Information Technology from ETH Zürich, Switzerland, in 2005, and the Ph.D. degree from the Electronics Laboratory, ETH Zürich, Switzerland for his thesis "Integrated Active Pulsed Reflector for an Indoor Positioning System." Currently, he is Vice President (VP) of Product Development at 3db Access. His expertise covers analog design, high-speed electronics and localization. Prior to 3db, Silvan used to work for Gigoptix/IDT in the area of high-speed optical communications. Silvan

has 12+ years academic and industry experience.



Michele Magno (Senior Member, IEEE) received the Master's and Ph.D. degrees in electronic engineering from the University of Bologna, Italy, in 2004 and 2010, respectively. Currently, he is a Senior Researcher at ETH Zürich, Switzerland, where he is the Head of the Project-Based Learning Center. He has collaborated with several universities and research centers, such as Mid University Sweden, where he is a Guest Full Professor. He has published more than 150 articles in international journals and conferences, in which he got multiple best paper

and best poster awards. The key topics of his research are wireless sensor networks, wearable devices, machine learning at the edge, energy harvesting, power management techniques, and extended lifetime of battery-operated devices.



Elena Simona Lohan received an M.Sc. degree in Electrical Engineering from University Politehnica of Bucharest (1997), a D.E.A. degree in Econometrics at Ecole Polytechnique, Paris (1998), and a Ph.D. degree in Telecommunications from Tampere University of Technology (2003). She is now a full Professor at Electrical Engineering unit at Tampere University (TAU) and a Visiting Professor at Universitat Autònoma de Barcelona (UAB), Spain. She is leading a research group on Signal processing for wireless positioning. She is a co-editor of the first

book on Galileo satellite system (Springer "Galileo Positioning technology"), co-editor of a Springer book on "Multi-technology positioning", and author or co-author in more than 220 international peer-reviewed publications, 6 patents and inventions. She is also an associate Editor for RIN Journal of Navigation and for IET Journal on Radar, Sonar, and Navigation. She is currently coordinating the A-WEAR MSCA European Joint Doctorate network in the field of wearable computing. Her expertise includes signal processing for wireless positioning and navigation, multipath and interference mitigation, and RF fingerprinting.



Dragoș Niculescu obtained a Ph.D. in Computer Science from Rutgers University (New Jersey) in 2004, with a thesis on sensor networks routing and positioning. He spent five years as a researcher at NEC Laboratories America in Princeton, NJ, working on simulation and implementation of mesh networks, VoIP, and WiFi-related protocols. At University Politehnica of Bucharest he is currently teaching courses in Mobile Computing and Services for Mobile Networking; also researching mobile protocols, UWB, and 802.11 networking.

PUBLICATION 3

Challenges in Platform-Independent UWB Ranging and Localization Systems

L. Flueratoru, E. S. Lohan, and D. Niculescu

In: *The 16th ACM Workshop on Wireless Network Testbeds, Experimental evaluation and Characterization (WiNTECH)*. ACM, 2022, pp. 1–8

doi: 10.1109/3556564.3558238

Publication reprinted with the permission of the copyright holders.

Challenges in Platform-Independent UWB Ranging and Localization Systems

Laura Flueratoru
laura.flueratoru@tuni.fi
University Politehnica of Bucharest,
Romania
Tampere University, Finland

Elena Simona Lohan
elena-simona.lohan@tuni.fi
Tampere University
Tampere, Finland

Dragoș Niculescu
dragos.niculescu@upb.ro
University Politehnica of Bucharest
Bucharest, Romania

ABSTRACT

The Ultra-Wideband (UWB) technology has grown in popularity to the point in which there are numerous UWB transceivers on the market that use different center frequencies, bandwidths, or hardware architectures. At the same time, efforts are made to reduce the ranging and localization errors of UWB systems. Until now, not much attention has been dedicated to the cross-platform compatibility of these methods. In this paper, we discuss for the first time the challenges in obtaining platform-independent UWB ranging and localization systems. We derive our observations from a measurement campaign conducted with UWB devices from three different developers. We evaluate the differences in the ranging errors and channel impulse responses of the devices and show how they can affect ranging mitigation methods customized for one device only. Finally, we discuss possible solutions towards platform-independent UWB localization systems.

CCS CONCEPTS

• **Hardware** → **Wireless integrated network sensors**; • **Networks** → **Location based services**.

KEYWORDS

ultra-wideband, ranging, indoor localization

ACM Reference Format:

Laura Flueratoru, Elena Simona Lohan, and Dragoș Niculescu. 2022. Challenges in Platform-Independent UWB Ranging and Localization Systems. In *Proceedings of (WiNTECH '22)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/3556564.3558238>

1 INTRODUCTION

Ultra-wideband (UWB) devices are increasing in popularity for their centimeter-level ranging and localization capabilities. At this date, there are at least four major smartphone brands equipped with UWB chipsets and at least six developers of UWB platforms. Therefore, localization systems must ensure interoperability among different device brands [22]. For this purpose, the FiRa™ Consortium [4] provides specifications and certifications that ensure interoperability between different UWB solutions.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WiNTECH '22, 17 October, 2022, Sydney, Australia

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9527-4/22/10.

<https://doi.org/3556564.3558238>

UWB systems can provide centimeter-level ranging and localization accuracy in line-of-sight (LOS) conditions but their performance is affected by obstacles between the transmitter (TX) and the receiver (RX), which is known as non-line-of-sight (NLOS) propagation. Multiple techniques have been proposed for dealing with errors caused by NLOS propagation [8, 21, 24]. Most of them rely on the channel impulse response (CIR) of the signal, since it offers rich information about the propagation path of the signal. However, so far, little to no attention has been dedicated to ensuring the proposed methods perform well on different UWB platforms.

To the best of our knowledge, we are the first to look at the challenges faced by ranging and localization systems when dealing with devices from different vendors. We derive our observations from a measurement campaign using UWB devices developed by three companies: Qorvo, TDSR, and 3db Access. We acquired measurements with pairs of devices from each brand at the same locations under different LOS and NLOS conditions. The measurements are made open-source [6] to facilitate future research in platform-independent UWB localization systems. The goal of the paper is to evaluate the differences in terms of distance errors and CIRs between different device models under the same propagation conditions (i.e., the same location, furniture arrangement, crowdedness, and, if applicable, obstacles).

First, we illustrate the danger in ignoring the cross-platform compatibility of a centralized localization system. We evaluate the performance of a neural network (NN) trained to achieve a good distance error prediction for one device when it is tested on measurements from different devices (acquired at the same locations), simulating the scenario in which the system disregards the users' device models. Our results show that an error mitigation technique that disregards the device model of incoming measurements may actually *degrade* the final performance of the system.

Second, we look at the root causes of cross-platform compatibility issues in UWB systems and we identify four main challenges:

- (1) The same environmental conditions lead to different distance errors for different devices, especially in NLOS conditions. Although great efforts are made to reduce ranging and localization errors through NLOS detection and mitigation, not all types of obstacles introduce significant ranging errors and, perhaps more interestingly, they do not introduce the same error in all devices even under the same conditions.
- (2) The same (NLOS) environmental conditions can lead to a different time of arrival (TOA) estimation in consecutive measurements even for the same device. Therefore, NLOS errors depend not only on the obstacle and environment, but also on the hardware and its TOA estimation algorithm.

- (3) CIRs acquired by different platforms at the same locations have different lengths, shapes, and statistics. This can be due to different center frequencies, pulse shapes and bandwidths, nonlinearities imprinted by the different front-end architectures on the CIR, antenna propagation characteristics, etc.
- (4) Different vendors provide different types of diagnostics (e.g., signal power, noise level), sometimes using incompatible or undisclosed units of measurement, which can hinder the cross-platform compatibility of localization methods based on such diagnostics [21].

Finally, we use these observations to make recommendations towards platform-independent UWB localization systems.

2 RELATED WORK

Measurements with different UWB devices at the same locations have been previously performed in [7, 11, 16, 19]. However, these studies were focused on comparing the performance of the devices in terms of ranging and/or localization accuracy or energy consumption. In this paper, we look not only at the *average* ranging error of each device, but also at the error distribution at selected locations, which gives insight into their different operating modes. To the best of our knowledge, we are the first to also compare the CIRs of different devices and to highlight how these differences can prevent the cross-platform compatibility of ranging and localization systems. We also compare other platforms than in the previous works and provide one of the few open-source datasets of this kind.

The localization accuracy of UWB-based systems can be improved by filtering or correcting individual distance measurements used in multilateration algorithms and/or by filtering or estimating locations directly. Distance errors can be reduced through NLOS detection and mitigation [24], data-driven TOA estimation algorithms [5], or models trained for distance error prediction [8, 21]. Other approaches estimate directly the location based on CIR features [12]. In recent years, ML-driven approaches for ranging and localization have taken precedence over statistics-based techniques [17]. However, all the cited solutions have been trained on data from only one device model. Previous work has also revealed that models trained in one environment have issues in adapting to different environments [3]. Given the diversity of UWB devices on the market, we deem necessary to evaluate the differences between various UWB devices that might affect the performance of error mitigation methods when applied on unknown devices.

3 DATA COLLECTION

In this section, we briefly present the experimental setup used to collect the data. The dataset is open-source and a detailed description of all the locations in which we acquired measurements can be found in the repository [6]. The goal of the measurement campaign is to verify the consistency of CIRs and ranging errors from different devices at the same locations.

We performed ranging measurements between pairs of devices belonging to each brand in several locations at Tampere University. We targeted LOS and NLOS scenarios, for the latter using as obstructions walls, pillars, furniture, a TV screen, and room divider panels.




Fig. 1: Examples of environments from the measurement campaign.

Table 1: Device settings used in the experiments: the center frequency (f_c), the pulse bandwidth (B_p), and the CIR sampling period (T_s^{CIR}).

Device	f_c [GHz]	B_p [MHz]	T_s^{CIR} [ns]
DW3000	6.5	600	1
TDSR	4.3	620	0.061
3db	6.5	380	1

We performed measurements in nine different spaces which included meeting rooms, offices, corridors, or cafeterias. Fig. 1 shows some of the locations in which we acquired measurements.

We used three models of UWB devices from different developers: the Qorvo (formerly Decawave) DW3000, the 3db Access 3DB6830C, and the TDSR P452A. For brevity, we will further refer to the devices, respectively, as DW3000 (or DW), 3db, and TDSR. Table 1 summarizes the characteristics and settings of the devices. DW3000 and 3db devices have the same center frequency of 6.5 GHz but different pulse bandwidths of 600 MHz and 380 MHz, respectively. TDSR devices have a different center frequency of 4.3 GHz but a bandwidth close to the DW3000 (620 MHz). DW3000 and 3db devices are compliant with the IEEE 802.15.4z standard [1], while TDSR devices implement a proprietary physical interface.

We took several measures to ensure that the devices from different manufacturers are placed at the same locations for each measurement. First, we mounted the devices on tripods such that the center of their antennas are aligned with respect to the tripod. Second, we marked the locations of the tripod legs on the floor for each test point. We estimate that, during the measurement process, there can be a maximum error of 1 cm between the antenna centers of two different devices.

At each test point, we kept a minimum recording time of 30 s in order to capture the CIR variations over multiple measurements. Because of some software limitations, in a part of the measurements the ranging update rate of DW3000 was lower than the one of 3db and TDSR devices (2 s vs. approx. 300 ms, respectively). In total,

our database contains approx. 30,000 measurements with 3db and TDSR devices each and 2,800 measurements with Decawave. The proportion of LOS and NLOS measurements is approx. 50/50.

We need to preprocess the CIRs in order to compare them. TDSR and 3db devices provide the real part of the CIR, while the DW3000 stores the complex CIR. To unify the three representations, we use the absolute value of the CIR. The TDSR CIRs have a sampling period of 61 ps so we downsampled them to 1 ns using linear interpolation. We also rescale the CIRs of each device to have amplitudes between $[0, 1]$, taking into account the global minimum and maximum of the entire dataset for each device. The rescaling enables us to compare the CIR statistics from different devices and does not change the statistics of CIRs from the same dataset (i.e., from one manufacturer).

4 CROSS-PLATFORM PERFORMANCE OF AN ERROR PREDICTION MODEL

In this section, we want to highlight the danger in disregarding the cross-platform compatibility of a centralized localization system (LS). In many LSs, the anchors collect ranging measurements from a user, then transmit them to a central server, which further processes them and estimates the user's location [8]. Let us consider that the processing step consists of a NN which predicts the distance error¹ based on the CIR and uses the prediction to refine the estimated distance. Since the data collection process is difficult for one device, let alone for multiple brands of devices, we can assume that the NN was trained on a dataset acquired with only one type of device, as it is frequently the case in the literature [5, 21, 24]. However, in a public space, the system might have to deal with measurements from users with different UWB devices. So we evaluate the possible outcome if the NN naively tried to predict errors based on measurements from devices that it was not trained on. We will compare the performance of the model when applied on measurements from the device used for training vs. on measurements from different devices, both *acquired at the same locations*.

Let $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be our training set, where \mathbf{x}_i is the feature vector (a portion of the CIR) from the feature space \mathcal{X} and y_i is an instance of the target variable Y (the distance error) defined on the domain \mathcal{Y} . Our goal is to learn the function $f: \mathcal{X} \rightarrow \mathcal{Y}$, that maps a CIR to its distance error.

We use a fully connected network with three dense layers, each using a rectified linear unit (ReLU) activation function plus an additional fully connected layer. Each layer uses 256 filters. We used the Focal-R loss function proposed in [23] for imbalanced regression to deal with the long-tailed distribution of the target variable. We used an input size of 40 CIR samples aligned to the TOA, keeping 30 samples before and 10 samples after the TOA. Although using more samples after the TOA might show benefits, we preferred to use the minimum common CIR length among all devices in order to avoid zero-padding shorter CIRs.

We remind that our dataset includes multiple measurements performed with the three brands of devices from the same physical location, from multiple locations inside the same room, and from

¹Among the possible error mitigation strategies, we chose predicting the distance *error* to better control the target variable distribution compared to directly predicting the (correct) distance. NLOS detection is another popular method but, as per Section 5.1, it often discards many useful measurements.




Fig. 2: CDF of the original distance errors ($|e_i|$) vs. the corrected errors ($|\hat{e}_i - e_i|$) using models trained on the same or different device(s).

different rooms. To avoid overfitting, we include measurements from different locations in the training, test, and validation sets. We select at random one location for the test set, one for the validation set, and assign the rest to the training set. We repeat this procedure 10 times, generating 10 different splits each time. The same split is used for all devices, meaning that split number $k \in \{1, \dots, 10\}$ from each device contains exactly the same combinations of locations and rooms inside the train, validation, and test sets. Therefore, if a model M trained on the set $D_k^{\text{train}}(d_1)$ from device d_1 achieves a good performance on the test set $D_k^{\text{test}}(d_1)$ but a bad performance on $D_k^{\text{test}}(d_2)$ from device d_2 , the difference in the performance will be due to the hardware since $D_k^{\text{train}}(d_1)$ and $D_k^{\text{test}}(d_2)$ contain measurements acquired at exactly the same locations.

We evaluate the initial distance errors and the errors after correction using the model's predictions. Let $e_i = \hat{d}_i - d_i$ be the distance error of measurement $i \in \{1, \dots, N\}$, i.e., the difference between the measured distance \hat{d}_i and the true distance d_i . The average initial (absolute) error is $\mu_{\text{init}} = \sum_{i=1}^N |e_i|$, where we aggregate the errors over all test set splits. The average corrected (absolute) error is $\mu_{\text{corr}} = \sum_{i=1}^N |\hat{e}_i - e_i|$, where \hat{e}_i is the error predicted by the model for measurement i . If the prediction perfectly matches the true error, then the distance error is completely mitigated.

Fig. 2 shows the CDF of the original distance errors of all devices vs. the corrected errors using the trained NNs. First, when using the same device for training and testing, the error after correcting the distance using the model's prediction is 22–27% smaller in the mean and 31–42% smaller in the standard deviation compared to the initial error. This shows that the model is able to generalize to unknown locations from similar environments.

When applying the error prediction model on an unknown device, i.e., that the model was not trained on, the average *corrected* error is always higher than the initial error, in most cases 3–8× larger. This shows that models not trained on a particular device might *degrade* the performance of a ranging system if we disregard information about the platforms used for training and testing. Therefore, the cross-platform compatibility of error mitigation techniques in LSs should be taken into account from the design stage.




Fig. 3: (a) CDF of errors in LOS and NLOS for all devices. (b) Distribution of errors from selected locations with different obstacles.

5 CHALLENGES

In this section, we analyze the collected data and highlight the challenges in obtaining cross-platform compatible UWB-based ranging and localization systems.

5.1 Same Conditions, Different Distance Errors for Different Devices

NLOS propagation is defined by the absence of a *visual* direct propagation path between the TX and the RX. Obstructions can introduce distance errors in UWB measurements in two ways. First, objects can attenuate the signal traveling through the direct path, making it indistinguishable from noise. When this happens, the device can incorrectly estimate the TOA as corresponding to a later reflection which has a higher amplitude than the direct path, causing a time (and distance) error. Second, some obstacles can decrease the propagation speed of signals traveling through them. In this case, the direct signal will arrive later than it would have arrived through air, introducing a distance error.

Fig. 3a shows the cumulative distribution function (CDF) of distance errors obtained with the three devices in LOS and NLOS scenarios. Indeed, the errors obtained in NLOS scenarios are always larger than those obtained in LOS. However, it is worth noting that at least 50% of NLOS measurements are under 15 cm, which is usually deemed an acceptable ranging error [19]. Even if we developed a perfect NLOS detection algorithm, discarding or assigning lower weights to measurements with small errors might decrease the accuracy of the localization system, especially in cases with few anchors available. Therefore, a binary LOS/NLOS classification might not always be the best strategy for distance error mitigation.

Table 2 shows the median error of each device for all types of obstacles considered here. While it is true that certain types of obstacles (e.g., concrete walls, TV screen) generally introduce higher errors than others (e.g., door) [11, 20], we noted that not all devices yielded the same ranging errors even when placed under exactly the same conditions. Fig. 3b shows the distribution of distance errors (illustrated as box plots using Tukey’s definition) obtained by the three devices at one location for several obstacles. There is a high variability in the distance errors obtained by different devices at the same location and, sometimes, a high spread in the errors obtained by a *single* device at one location.

Table 2: Median distance errors with various obstacles.

Device	Median distance error [m]					
	LOS	TV screen	Fridge	Door	Pillar	Half wall
DW	0.06	0.45	0.31	0.10	0.57	0.08
TDSR	0.00	0.22	0.33	0.08	0.28	0.01
3db	-0.02	1.89	0.00	-0.23	0.42	0.05

5.2 Same Conditions, Different TOAs

It is perhaps useful to explain the root cause of the differences in the measured distance of various devices. Since the distance between two devices is computed based on the (round-trip) time of flight between the TX and the RX, any error in the TOA estimation will also yield a distance error.

Although the TOA estimation algorithms used by the devices are closed-source, the “Applications of the IEEE 802.15.4 standard” document [9] describes some approaches that are likely followed by the manufacturers. In LOS, the maximum peak of the CIR usually corresponds to the direct path, so the (correct) TOA is straight-forward to obtain. In NLOS, however, the direct path can be attenuated by an obstruction and later paths, which correspond to longer travel times, often have higher amplitudes. Therefore, the receiver needs to also consider other peaks in the vicinity of the strongest one as possible candidates for the first path. Popular approaches to implement the back-search include a sequential linear cancellation scheme [18] or threshold-based techniques [10]. Both techniques are model-based and require assumptions about the propagation conditions, such as the power ratio between two consecutive paths, the number of MPCs to be subtracted, the maximum peak to earlier peak ratio, or the peak to average power ratio.

We found that the most common reason for differences in ranging accuracy between the devices is the different TOA estimation, especially in the case of direct paths that have low amplitudes (close to the noise threshold). In our measurements, the “global” CIR shape in NLOS propagation is relatively stable for a single device and it is only the estimated TOA that fluctuates in some NLOS scenarios.

To illustrate this, let us consider a series of CIRs acquired with one pair of devices from the same brand in a NLOS scenario with a TV screen between the transmitter and the receiver. The devices are left unmoved during the experiment. We use the DW3000 for




Fig. 4: The plots show stacked CIRs (with their amplitude encoded through the color depth) acquired during a series of measurements in which there was a TV screen between two DW3000 devices. The circle markers represent the TOA estimated by the hardware and its color encodes the distance error of that measurement. Fig. 4a shows a portion of the unaligned CIRs (i.e., as dumped in the buffer by the device), Fig. 4b shows the CIRs aligned to the TOA, and Fig. 4c shows the CIRs aligned to the maximum value of the cross-correlation function computed between the first CIR and all of the following.

this example, but all devices present similar patterns. Fig. 4 shows series of 30 CIRs stacked vertically with different alignments. The MPCs from each CIR have color-encoded amplitudes and the TOA is denoted by a circle marker. The color of the circle marker encodes the distance error of the measurement corresponding to that CIR.

Fig. 4a shows the “raw” CIR alignment. In theory, the TOA can occur at any sample in the internal buffer which stores the CIR. In practice, we found that all DW3000 TOAs in our dataset occur between samples 710 and 750 in a quasi-random manner. In Fig. 4a, batches of consecutive CIRs appear to be shifted versions of each other, but the pattern changes every 5–10 measurements. In this representation, we cannot distinguish a clear pattern between CIRs that lead to large errors or between the MPCs of consecutive CIRs.

In most works, the CIRs are cropped around the TOA identified by the device, as in Fig. 4b. This provides a convenient representation for LOS/NLOS classification or for distance error prediction, since CIRs are aligned to a common denominator and deviations from it (e.g., the maximum peak occurring earlier or later in the buffer, shorter or longer tails) can indicate different error magnitudes. Using this representation, we could believe that different CIR shapes are correlated with different magnitudes of ranging errors.

Instead, if we align the CIRs according to the index which yields the maximum cross-correlation value, i.e., according to their similarity, we obtain Fig. 4c. In this figure, it is evident that, in fact, consecutive CIRs acquired in the same scenario are almost identical. What differs in this representation is the estimated TOA. The measurements with the lowest distance errors have the earliest TOA. However, because these early paths have amplitudes close to the noise threshold, they are ignored in some measurements and the TOA is identified as corresponding to a later reflection. Note that, since the devices perform two-way ranging, they must identify the correct TOA during both the poll *and* the response messages. In the plot, we see only the CIRs of the initiator, but large errors can also be caused by an incorrect TOA estimation at the responder. Therefore, the devices can experience different errors even at the same location because of the different SNR of the first path (which is influenced by the hardware architecture and signal processing steps) or because of the internal TOA estimation algorithm.

5.3 Same Conditions, Different CIRs for Different Devices

At first sight, CIRs seem convenient for platform-independent learning problems, since they provide a representation of the paths through which the signal travels from the TX to the RX. Therefore, we would perhaps expect CIRs acquired by different devices at the same location to be similar. However, we will see why this is not the case. We focus on two main differences: the effective duration and the shape of CIRs.

5.3.1 Effective duration. One crucial difference between CIRs acquired by different devices is their length. Let us call the *effective duration* of the CIR the portion starting from the estimated TOA until the end of buffer in which CIRs are stored, since this is usually the portion of most interest. The buffers have a length of 1016 samples for DW3000, 1632 samples for TDSR, and 256 samples for 3db. However, because we downsampled the TDSR from 61 ps to 1 ns, they will have a shorter (effective) duration than the other devices. In practice, we found that TDSR devices capture, on average, the paths which arrive within 17 ns of the TOA. In comparison, DW3000 and 3db devices have an average effective length of 278 ns and 125 ns, respectively. This difference is evident in Fig. 5, which shows an example of CIRs acquired with the three devices at the same location in one LOS scenario.

Different effective CIR lengths can cause issues in NNs which receive as input the raw CIR from different platforms. If the networks are designed to work on CIRs with a pre-defined length, shorter CIRs must be padded until the desired length, which can change the input sample distribution and the model’s performance on those inputs. Alternatively, recursive NNs can be used on inputs of different lengths.

5.3.2 CIR shape. The received signal can be modeled as [13]:

$$r(t) = s(t) * h(t) + n(t), \quad (1)$$

where $s(t)$ is the transmitted pulse which is convolved with the CIR $h(t)$. The signal $n(t)$ represents sensor or environmental noise (so not related to the propagation path), usually modeled as zero-mean white Gaussian noise.

The CIR can be decomposed into K multipath components (MPCs) with delays τ_k , $k = 1, \dots, K$ and amplitudes α_k and a stochastic process $v(t)$ which results in diffuse multipath caused by scattering

and diffraction:

$$h(t) = \sum_{k=1}^K \alpha_k \delta(t - \tau_k) + v(t). \quad (2)$$

The estimated CIR, denoted by $\hat{h}(t)$, is obtained by decorrelating the received signal $r(t)$ with the known template pulse $s(t)$.

Because we acquired measurements with the three platforms at the same location, we would perhaps expect to see MPCs with the same delays τ_k in their CIRs. However, because the devices have different center frequencies and/or pulse bandwidths, there will be different constructive or destructive interference patterns reflected in $v(t)$ or even in the observed delays of the main MPCs τ_k . In addition, the amplitudes α_k are influenced by the different front-end circuits. The various signal processing components (e.g., low-noise amplifiers, mixers, automatic gain controls, analog filters, analog-to-digital converters) and digital processing can introduce different linear and/or non-linear distortions depending on the architecture. Therefore, even if the channel conditions and environment are the same, we will, in fact, see different patterns in the CIRs acquired by different platforms.

To characterize CIRs acquired at the same location with different devices, we look at the average number of main peaks (or MPCs) and the average delay between the first and the last significant peaks, which indicate whether we can identify the same main MPCs in all CIRs and how long it takes until their amplitudes decay to an insignificant level. We also compute the energy, mean excess delay (MED), and root-mean square (RMS) delay spread of the CIRs, which have been previously used to characterize CIRs in LOS/NLOS detection problems [17]. The MED and RMS delay spread are, respectively, the first and second moments of the power delay profile of the signal and characterize its delay statistics. The RMS delay spread captures the temporal dispersion of the signal's energy.

We search for the number of significant MPCs in each CIR, which we define as the peaks that exceed 25% of the maximum amplitude of the CIR with a minimum time separation between peaks of 2 ns (to avoid detecting peaks belonging to the same path). We compute the average number of significant peaks (N_p) and the average time delay between the first and the last significant peak (μ_δ). The peak search is performed on the raw CIR, while the energy, MED, and RMS delay spread are computed on the CIR starting from the TOA until the end of the buffer, to mitigate the influence of the TOA index in the raw CIR buffer on the CIR statistics.

Fig. 5 shows the most significant peaks in a triplet of CIRs acquired at the same location with different devices and Table 3 shows the average CIR statistics over all locations. 3db CIRs have more significant MPCs than the other two devices and their amplitudes take longer to decay, reflected in a higher μ_δ and energy. TDSR CIRs have the shortest duration and therefore capture the fewest significant peaks, so they have the lowest energy, MED, and RMS delay spread.

In the literature, a low signal energy and high delay statistics have been associated with signal attenuation occurring in NLOS propagation [17]. However, if CIRs acquired in the *same* conditions with different brands of devices have, on average, different statistics,




Fig. 5: CIRs and their main peaks acquired with each device at the same location in a LOS scenario.

Table 3: Statistics of the CIRs: number of significant paths (N_p), delay between the first and the last significant path (μ_δ), energy (E , quantized), mean excess delay (τ_{MED}), and RMS delay spread (τ_{RMS}).

Device	N_p	μ_δ [ns]	E [-]	τ_{MED} [ns]	τ_{RMS} [ns]
DW3000	3	11.2	2.4	13.5	21.7
TDSR	2	5.3	1.6	8.2	3.0
3db	5	24.4	8.5	13.3	13.8

these can introduce issues in error mitigation methods customized for one device but applied on different brands of devices.

5.4 Different Devices, Different Diagnostics

Since the CIRs are long (256–1632 samples for the devices we used), for some applications it can be too time consuming and computationally expensive to process them. Therefore, some works instead propose using other diagnostics provided by the devices (for instance, power or noise figures) for ranging/localization error mitigation [21]. However, this can pose issues for platform-independent localization systems because not all devices provide the same metrics and often not in the same units of measurement.

For instance, the DW3000 provides the power, maximum amplitude, and phase of arrival (POA) computed on the preamble for regular ranging and, additionally, the power and POA computed on the scrambled timestamp sequence for the secure ranging mode. The manufacturer provides formulas for converting the power to dBm [15]. TDSR devices provide the maximum value in the leading edge (LE) window of the received CIR and the noise level. However, the unit of measurement of these parameters is not specified in the documentation [14], so it is not clear if these parameters could be compared, for instance, with the ones of the DW3000 chipset. TDSR devices also provide the noise amplitude and a coarse and a precise estimate of the range, of the distance error, and of the tag's

velocity. The 3db chip provides the peak and the LE amplitudes (but not expressed in dBm) and other additional diagnostics [2].

Given the different diagnostics provided by the devices and the lack of a common measurement unit for the power and noise figures, it would be difficult to use this kind of information in a platform-independent UWB localization system. Therefore, the manufacturers should invest more effort to provide unified metrics for additional diagnostics.

6 DISCUSSION

As we showed in Section 4, error mitigation techniques developed for only one UWB hardware should implement a fallback plan to avoid degrading the accuracy of measurements acquired by other devices. An alternative that requires further research is to develop error mitigation solutions that can work across multiple platforms. Although different platforms have different CIRs and ranging errors under the same conditions, there is a common reason for all errors: the incorrect TOA estimation. Therefore, there might exist a set of device-agnostic CIR features related to the estimated vs. the correct TOA that could yield the correct distance (or TOA) errors across multiple devices.

An alternative to centralized solutions is to implement error correction only at the edge and use models customized for each target device. However, in this case, more work is needed to create environment-independent models [3], since it is unreasonable to collect data from every new location with every device on the market. Data-driven TOA estimation shows promising results [5] but it must be lightweight enough to be implemented on-chip and fast enough to minimize clock drift errors.

7 CONCLUSIONS

In this paper, we provided an overview of the challenges faced in obtaining platform-independent UWB ranging or localization systems. We derived our observations from a database of measurements acquired with three brands of UWB devices at exactly the same locations. We evaluated the differences in their ranging errors, CIRs, and diagnostics. We show that applying error mitigation models on devices not included in the training phase might result in a severe degradation of the ranging accuracy.

The cross-platform compatibility of error mitigation methods should not be taken for granted; instead, ranging and localization systems deployed in environments where users can have different UWB platforms must take into account cross-platform compatibility from the design phase. Future research should focus more on learning methods that can perform equally well on a wide range of devices or that degrade gracefully in the presence of an unknown device. Alternatively, error mitigation techniques could be moved to the edge and customized on the target device with the caveat that they must be environment-independent.

ACKNOWLEDGMENTS

The authors gratefully acknowledge funding from European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints, <http://www.a-wear.eu/>). This

work does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of its content.

REFERENCES

- [1] 2020. IEEE Standard for Low-Rate Wireless Networks. *Amendment 1: Enhanced Ultra Wideband (UWB) Physical Layers (PHYs) and Associated Ranging Techniques (IEEE Std 802.15.4z) (2020)*.
- [2] 3db Access AG. 2018. 3DB6830D User Guide. Version 4.0.
- [3] Valentin Barral, Carlos J Escudero, José A García-Naya, and Pedro Suárez-Casal. 2019. Environmental cross-validation of NLOS machine learning classification/mitigation with low-cost UWB positioning systems. *Sensors* 19, 24 (2019), 5438.
- [4] FiRa Consortium. 2022. Retrieved 19 August 2022 from firaconsortium.org
- [5] Tobias Feigl, Ernst Eberlein, Sebastian Kram, and Christopher Mutschler. 2020. Robust ToA-Estimation using Convolutional Neural Networks on Randomized Channel Models. In *IPIN*. IEEE, 1–8.
- [6] Laura Fluoratoru. 2022. Dataset: Ultra-wideband ranging measurements acquired with three different platforms (Qorvo, TDSR, 3db Access). <https://doi.org/10.5281/zenodo.6984698>
- [7] Laura Fluoratoru, Silvan Wehrli, Michele Magno, Elena Simona Lohan, and Drago Niculescu. 2021. High-accuracy ranging and localization with ultra-wideband communications for energy-constrained devices. *IEEE IoT-J* (2021).
- [8] Jaron Fontaine, Matteo Ridolfi, Ben Van Herbruggen, Adnan Shahid, and Eli De Poorter. 2020. Edge inference for UWB ranging error correction using autoencoders. *IEEE Access* 8 (2020), 139143–139155.
- [9] IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs). 2014. Application of IEEE Std 802.15.4.
- [10] Ismail Guvenc, Sinan Gezici, and Zafer Sahinoglu. 2008. Ultra-wideband range estimation: Theoretical limits and practical algorithms. In *IEEE Int'l Conf. on UWB*, Vol. 3. IEEE, 93–96.
- [11] Antonio Ramón Jiménez and Fernando Seco. 2016. Comparing Decawave and Bespoon UWB location systems: Indoor/outdoor performance analysis. In *IPIN*. IEEE, 1–8.
- [12] Sebastian Kram, Maximilian Stahlke, Tobias Feigl, Jochen Seitz, and Jörn Thielecke. 2019. UWB channel impulse responses for positioning in complex environments: A detailed feature analysis. *Sensors* 19, 24 (2019), 5547.
- [13] Erik Leitinger, Paul Meissner, Christoph Rüdiger, Gregor Dumphart, and Klaus Witrals. 2015. Evaluation of position-related information in multipath components for indoor positioning. *IEEE JSAC* 33, 11 (2015), 2313–2328.
- [14] TDSR LLC. 2020. TDSR RangeNet Application Programming Interface (API) Specification. Version 320-0313L.
- [15] Decawave Ltd. 2021. DW3000 User Manual. Version 1.1.
- [16] G MacGougan, K O'Keefe, and R Klukas. 2009. UWB ranging and ranging measurement accuracy. *Measurement Science and Technology* 20, 9 (2009), 0957–0233.
- [17] Stefano Marano, Wesley M Gifford, Henk Wymeersch, and Moe Z Win. 2010. NLOS identification and mitigation for localization based on UWB experimental data. *IEEE JSAC* 28, 7 (2010), 1026–1035.
- [18] Yihong Qi, Hisashi Kobayashi, and Hirohito Suda. 2006. On time-of-arrival positioning in a multipath environment. *IEEE Trans. on Vehicular Technology* 55, 5 (2006), 1516–1526.
- [19] Antonio Ramón Jiménez Ruiz and Fernando Seco Granja. 2017. Comparing Ubisense, Bespoon, and Decawave UWB location systems: Indoor performance analysis. *IEEE TIM* 66, 8 (2017), 2106–2117.
- [20] Anthony Schenck, Edwin Walsh, Jonas Rejniers, Ted Ooijselaar, Risang Yudanto, Erik Hostens, Walter Daems, and Jan Steckel. 2018. Information Theoretic Framework for the Optimization of UWB Localization Systems. In *IPIN*. 1–8.
- [21] Lorenz Schmid, David Salido-Monzú, and Andreas Wieser. 2019. Accuracy assessment and learned error mitigation of UWB ToF ranging. In *IPIN*. IEEE, 1–8.
- [22] Maximilian Schuh, Hannah Brunner, Michael Stocker, Markus Schuß, Carlo Alberto Boano, and Kay Römer. 2022. First Steps in Benchmarking the Performance of Heterogeneous Ultra-Wideband Platforms. In *CPS-IoTBench*. IEEE, 34–39.
- [23] Yuzhe Yang, Kaiwen Zha, Yingcong Chen, Hao Wang, and Dina Katabi. 2021. Delving into deep imbalanced regression. In *Int'l Conf. on Machine Learning*. 11842–11851.
- [24] Kegen Yu, Kai Wen, Yingbing Li, Shuai Zhang, and Kefei Zhang. 2018. A novel NLOS mitigation algorithm for UWB localization in harsh indoor environments. *IEEE Trans. on Vehicular Technology* 68, 1 (2018), 686–699.

PUBLICATION 4

Self-Learning Detection and Mitigation of Non-Line-of-Sight Measurements in Ultra-Wideband Localization

L. Flueratoru, E. S. Lohan, and D. Niculescu

In: *2021 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*.

IEEE, 2021, pp. 1–8

DOI: 10.1109/IPIN51156.2021.9662532

Publication reprinted with the permission of the copyright holders.

Self-Learning Detection and Mitigation of Non-Line-of-Sight Measurements in Ultra-Wideband Localization

Laura Flueraoru^{*†}, Elena Simona Lohan[†] and Dragoş Niculescu^{*}

^{*} Computer Science Department, University Politehnica of Bucharest, Romania

[†] Electrical Engineering Unit, Tampere University, Finland

Abstract—Non-line-of-sight (NLOS) propagation is one of the main error sources in indoor localization, so a large body of work has been dedicated to identifying and mitigating NLOS errors. The most accurate NLOS detection methods often rely on large training data sets that are time-consuming to obtain and depend on the environment and hardware. We propose a method for detecting NLOS distance measurements without manually collected training data and knowledge of channel statistics. Instead, the algorithm generates LOS/NLOS labels for sets of distance measurements between fixed sensors and the mobile target based on distance residuals. The residual-based detection has 70–80% accuracy but has high complexity and cannot be used with high confidence on all measurements. Therefore, we use the predicted labels and the channel impulse responses of the measurements to train a classifier that achieves over 90% accuracy and can be used on all measurements, with low complexity. After we train the classifier during an initial phase that captures specifics of the devices and of the environment, we can skip the residual-based detection and use only the trained model to classify all measurements. We also propose an NLOS mitigation method that reduces, on average, the mean and standard deviation of the localization error by 2.2 and 5.8 times, respectively.

Index Terms—Non-line of sight (NLOS), localization, positioning, ultra-wideband (UWB), machine learning

I. INTRODUCTION

Indoor localization has garnered attention in recent years for its useful applications such as navigating in public spaces, offering customized location-based services and interactions with the environment, or controlling and monitoring industrial robots and indoor drones. One popular localization method uses several sensors (called anchors) with fixed and known positions that communicate with a mobile target (called tag). The method estimates the location of the tag based on time or distance measurements between each anchor and the tag. Non-line-of-sight (NLOS) propagation, in which an object or a person obstructs the direct path between two devices, affects most localization methods. In this case, the observed time of flight (TOF) or distance between two devices is larger than without the obstacle, which causes a localization error.

This work was supported by funding from European Union’s Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints, <http://www.a-wear.eu/>). The work was also partly supported by a grant from the Romanian National Authority for Scientific Research and Innovation, UEFISCDI project PN-III-P3-3.6-H2020-2020-0124.

Detecting and correcting NLOS distance errors has been widely studied in the literature [1], [2], [3], [4], [5], [6]. The methods with the highest detection accuracy rely on the statistics of the channel impulse response (CIR) of the signal [2], [3], [4]. However, they require extensive measurement campaigns to learn the statistics of LOS and NLOS measurements. Such measurements are time-consuming, require some expertise, depend on the environment and on the used hardware, and need to be repeated frequently in order to capture the environment dynamics. Therefore, collecting training data before every deployment and maintaining databases up-to-date are demanding tasks, which are usually infeasible in practice.

In this work, we propose an NLOS detection and mitigation method that does *not* require manually-acquired training data nor channel statistics. Fig. 1 shows the main steps of our approach. When a tag is first deployed in an area, it starts the initial phase of the algorithm, in which it collects distance measurements and CIRs from all the anchors in the area. In this first step, the measurements are labeled as LOS, NLOS, or ambiguous using *residual analysis*. If there are more anchors than the minimum necessary for 2D or 3D localization, we can compute locations using each subset of anchors. Since NLOS anchors¹ introduce higher location errors, they also have higher distance residuals, defined as the difference between the measured distance and the Euclidean distance between the anchor and the estimated position of the tag. The method labels anchors as LOS/NLOS if their average residuals can be grouped in two one-dimensional clusters (intervals). The labeling step has an acceptable accuracy (70–80%) but it cannot be used on all sets of measurements, since for small NLOS errors the separation between LOS and NLOS residuals is not clear. Moreover, the residual-based labeling step requires computing the location using all anchor combinations, which scales with $\mathcal{O}(2^N)$ for N anchors.

Therefore, we introduce the second step, in which we train a Random Forest (RF) classifier using the labels predicted in the first step and the CIR features of the measurements. The model can recover the correct class boundary even with noisy labels, reaching a higher classification accuracy (>90%) than the residual-based labeling. After the model is trained, we can directly classify all distance measurements as LOS/NLOS

¹“NLOS anchor” is shorthand for “anchor in NLOS with the tag.”




Fig. 1. Flowchart of the proposed method for the detection and mitigation of NLOS distance measurements.

using the RF and skip the residual-based labeling. Classifying samples using RF has a constant complexity and can be used in all localization instances.

We further propose a location-correction method based on identified LOS/NLOS measurements which does not discard NLOS measurements. We evaluate the accuracy and localization error of the proposed method on a database of measurements acquired with UWB devices, which therefore resembles a real localization setup.

II. RELATED WORK

NLOS identification methods based on channel statistics have been well studied in literature [2], [3], [4] and they achieve very high classification accuracy (>90%). However, they need extensive measurement campaigns to collect training data, which are rarely feasible in practice.

NLOS errors can be directly mitigated using semi-definite programming [5], [7] without assuming any measurement statistics. However, these methods are usually more computationally expensive than plain localization algorithms [5]. Other methods that do not require error statistics use additional hardware, such as inertial measurement units [8]. In [9], the authors proposed a NLOS mitigation technique for dense NLOS environments that does not need training but it assumes the measurement variance to be known. The error was corrected using two extended Kalman filters (EKFs) alternatively depending on the LOS/NLOS condition.

In [10], the authors propose an unsupervised NLOS identification method. The biggest difference from our approach is that the method in [10] can classify data only in bulk (so not online) because it needs a collection of data points to obtain the distribution of features with Gaussian mixture models.

In [11] and [12], the authors proposed NLOS identification methods based on pre-trained models. In [11], a convolutional neural network (CNN) is trained in one environment and updated with data from a new environment. The method is validated in two similar office environments, so it is not clear how well the model can be transferred between two very different environments, e.g., a mall and an office. This is also an open research question for our method. However, we do not rely on a pre-trained model but can train it online. In [12], a pre-trained model is improved by retraining using unlabeled samples. This approach can be used to improve the accuracy of our model (discussed in Section VI).

The closest work to ours is [6], where the authors used anchor residuals instead of CIR features to train a classifier. The simulated error for NLOS measurements was sampled

from a uniform distribution between 0.75–3.5 m, so NLOS measurements were easily distinguishable from LOS ones, which had zero-centered normally-distributed errors. However, in our measurement campaigns (described in Section IV), we found that typical NLOS errors with UWB devices are spread between a smaller range of 0.25–0.8 m, so it was harder to accurately identify NLOS measurements using only anchor residuals. Therefore, we propose a classification in two steps: first using anchor residuals, then using CIR features.

Another work that applied residual analysis to identify NLOS errors is [13]. The authors used residual analysis to identify NLOS errors, a voting algorithm to correct these errors, and a fuzzy C-means algorithm to classify NLOS errors into “hard” and “soft” NLOS. A Kalman filter (KF) and an unscented Kalman filter (UKF) filtered the two types of NLOS errors and corrected the location estimates. The authors in [13] focused more on NLOS error mitigation than on NLOS identification and they did not mention the accuracy of the classification method alone. Compared to [13], we also provide a NLOS detection method which can be useful in detecting obstacles, creating building maps, or estimating crowd densities. We show that we can also reduce localization errors with our NLOS detection method.

III. NON-LINE-OF-SIGHT DETECTION

In this section, we present the basics of anchor-based localization (Section III-A) and the main steps of the unsupervised NLOS detection method, also highlighted in Fig. 1. First, we use residual analysis to obtain LOS/NLOS labels using only distance measurements (Section III-B). By repeating this step for multiple locations, we create a database of CIR features and their predicted labels. These are given as training data to an RF classifier (Section III-C). Once trained, the classifier can directly classify subsequent measurements.

A. Localization

Range-based localization estimates the coordinates \hat{x} of a target (also called a tag) using the distance measurements between the tag and N anchors with known locations \mathbf{x}_{A_i} , where A_i is the i^{th} anchor, for $i = 1, \dots, N$. When the direct path between two devices is unobstructed, also known as line-of-sight (LOS) propagation, the distance between two devices can be recovered as $d = c \cdot \text{TOF}$, where c is the speed of light. If an object or person blocks the direct path, the signal usually travels at a lower speed than through the air. This causes a larger TOF than without the obstacle and the estimated distance is larger than in reality. Obstacles can also

completely block the direct signal; in such cases, reflections which travel longer paths than the direct signal can arrive at the receiver and also cause time delays. The last two scenarios are known as non-line of sight (NLOS) and cause distance and localization errors in UWB localization.

The anchor–tag distances can be written as:

$$d_i = \|\mathbf{x}_{A_i} - \mathbf{x}\| + v_i, \quad i = 1, \dots, N \quad (1)$$

where \mathbf{x} is the true location of the tag, $\|\cdot\|$ is the Euclidean norm, and v_i is a noise term. In vector form, this becomes:

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{v}, \quad (2)$$

where \mathbf{y} is a vector which contains all distance measurements d_i , $i = 1, \dots, N$, \mathbf{v} is the error vector, and h is a vector-valued measurement function.

The location can be found using the least squares method:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{y} - h(\mathbf{x})\|. \quad (3)$$

We chose the regularized Gauss-Newton multilateration algorithm from [14], because it has low computational complexity and localization errors comparable to closed-form solutions. The algorithm needs an initial location, which can be obtained with a closed-form solution. For each iteration k , the algorithm computes the Jacobian matrix:

$$J_k(\mathbf{x}) = \begin{bmatrix} \frac{\mathbf{p}_1 - \mathbf{x}}{\|\mathbf{p}_1 - \mathbf{x}\|}, \dots, \frac{\mathbf{p}_N - \mathbf{x}}{\|\mathbf{p}_N - \mathbf{x}\|} \end{bmatrix}^T. \quad (4)$$

The solution at each iteration is $\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta\mathbf{x}$ and $\Delta\mathbf{x}_k$ is the least-squares solution to:

$$-(\Sigma^{-\frac{1}{2}} J_k + c \mathbf{I}_N) \Delta\mathbf{x}_k = (\Sigma^{-\frac{1}{2}} (h(\mathbf{x}_k) - d) + c(\mathbf{x} - \mathbf{x}_r)), \quad (5)$$

where \mathbf{I}_N is the unitary matrix of size $N \times N$, \mathbf{x}_r is a regularization location and c is a regularization coefficient equal to the inverse of the standard deviation of a distribution centered at \mathbf{x}_r . The algorithm stops when the norm of the incremental location is smaller than the tolerance δ or when it reaches the maximum number of iterations k_{max} .

B. Unsupervised Labeling with Anchor Residuals

If the distance measurements d_i are noiseless, i.e. $v_i = 0$ in Eq. (1), then in the 2D case the tag is found at the intersection between circles centered at the anchors, with a radius equal to d_i , for $i = 1, \dots, N$ (see Fig. 2a). If the distance measurements are noisy, the circles do not intersect in a single point anymore and the tag's location is (ideally) found inside the intersection area of the circles, as shown in Fig. 2b. The *residual* of anchor A_i is defined as [1]:

$$r_i = d_i - d'_i, \quad (6)$$

where d_i is the measured distance between anchor A_i and the tag and d'_i is the distance between the anchor and the estimated location $\hat{\mathbf{x}}$ of the tag. Intuitively, an anchor's residual is likely to be higher when the anchor is in NLOS with the tag, since this causes a higher distance measurement error [1].

For M -dimensional localization, at least $M+1$ anchors are needed to solve the system of equations from Eq. (2). If N




Fig. 2. In the ideal (2D) case, the tag is found at the intersection point between the circles centered at the anchors' locations with the radius equal to each anchor–tag distance. When the measured distances are noisy, the circles do not intersect in a single point anymore. An anchor's residual r_i is the difference between the measured distance d_i and the distance between the estimated location and the anchor's location d'_i .




Fig. 3. Average anchor residuals, computed over all subsets in which an anchor is used. The residuals of LOS and NLOS anchors are linearly separable in many cases.

anchors with $N > M + 1$ are available, we can form subsets of K anchors, where $M + 1 \leq K \leq N$, and compute the residual of a subset S as the mean of the squared residuals of all anchors:

$$R_S = \frac{1}{|S|} \sum_{i=1}^{|S|} r_{i,S}^2, \quad (7)$$

where $|\cdot|$ denotes the cardinality of a set and $r_{i,S}$ is the residual of the i^{th} anchor in this subset. In [1], the final location estimate is obtained as the weighted average of the intermediate locations obtained in all subsets, where the weight is the inverse of a subset's residual.

Starting from the insight that anchor residuals are higher when anchors are in NLOS with the tag, we propose using the average residual of an anchor, computed over all subsets in which it is used, to identify NLOS anchors:

$$R_i = \frac{1}{|\{S|A_i \in S\}|} \sum_{\{S|A_i \in S\}} r_{i,S}. \quad (8)$$

Because NLOS anchors have higher distance measurement errors, their average residuals are also usually higher than those of LOS anchors. In many cases, residuals coming from LOS and NLOS anchors form two 1D clusters (or intervals) which can be separated. Fig. 3 shows such an example, for a simulation of eight anchors when the tag and some anchors are separated by a wall. The wall introduces a lognormally-distributed error with a median of 24 cm and a standard




Fig. 4. Examples of kernel density estimation (KDE) applied on anchor residuals for the shaping parameter $h = 0.08$.

deviation of 1.8m (the parameters were obtained from a measurement campaign [15]). In many cases, the Residual values of LOS and NLOS anchors can be clearly delimited.

To find the threshold that separates LOS from NLOS residuals, we use kernel density estimation (KDE) to get the distribution of the average anchor residuals (for a given location). The kernel density estimator of a series of independent and identically distributed samples $\{R_1, \dots, R_N\}$ is:

$$\hat{f}_h(R) = \frac{1}{Nh} \sum_{i=1}^N K\left(\frac{R - R_i}{h}\right), \quad (9)$$

where K is a non-negative function called the kernel (we used a Gaussian kernel) and $h > 0$ is a smoothing parameter. Fig. 4 shows three examples of KDE applied on anchor residuals, for $h = 0.08$. When LOS and NLOS anchor residuals form two different clusters, the distribution has two maxima and a single minimum, like in Fig. 4a. In this case, we label anchors whose residual is higher than the minimum as NLOS and the rest as LOS. There can also be ambiguous cases. In Fig. 4b, the residuals are uniformly spread over the interval and the distribution has only one maximum. In Fig. 4c, on the other hand, there are more than one local minima found.

The parameter h determines the smoothness of the fitted distribution. If h is too small, the estimated distribution will contain spurious data artifacts, similar to the distribution in Fig. 4c. If h is too high, the estimator cannot capture the underlying data structure, leading to an estimate similar to Fig. 4b. Therefore, too large or too small of a smoothing parameter leads to ambiguous cases in which the data cannot be labeled. Because we aim for a high LOS/NLOS detection accuracy, we label only unambiguous cases in which the distribution has exactly one minimum. In the current implementation, we use the ambiguous cases only for validation but in the future we could use them to retrain the RF model and increase its accuracy (see the discussion in Section VI). In Section III-B, we will discuss the choice of h which maximizes the labeling accuracy and the percentage of instances classified.

Because not all anchor residuals can be unambiguously split into two intervals, we cannot apply the residual-based labeling on all measurements. Also, we need to compute locations using all anchor subsets, which scales with 2^N . Therefore, we use the labels provided by the residual method to train an ML model that can classify *all* distance measurements.

C. Model Training

The identification and mitigation of the LOS/NLOS condition using features based on the CIR of the signal has been studied in [2], [3], [4]. Supervised classification methods have very high accuracy ($> 90\%$) but need training data, i.e., a database of distance measurements labeled as LOS/NLOS and their CIRs. We replace the manual labeling with the unsupervised labeling based on anchor residuals, as described in Section III-B. Because our training set has label noise (LOS measurements labeled as NLOS and vice-versa), we want to train a machine learning model robust to label noise. We chose a Random Forest (RF) classifier, since it is an ensemble machine learning (ML) algorithm that performs well with noisy labels [16]. The RF is a collection of decision trees which outputs the class predicted by most of the individual decision trees through bootstrap aggregating.

The purpose of this paper is not to find the best ML algorithm for the task, but to demonstrate the general idea, that we can detect NLOS measurements without training data using residual-based labeling. We leave as future work an exhaustive search through more ML models suitable for data sets with noisy labels that further increase the NLOS detection accuracy.

We train the model using CIR features known to characterize well LOS/NLOS conditions [2], [3]: the energy of the received signal, the maximum amplitude of the signal, the mean excess delay, the RMS delay spread, the kurtosis, and the difference between the TOA and the time at which the signal has the maximum amplitude ($\Delta T(\text{TOA}, \text{Max})$).

IV. EVALUATION SETUP

We simulate a localization scenario based on a database of real UWB measurements to evaluate the feasibility and performance of the proposed method for NLOS error detection and correction. Fig. 5 describes the simulation flow. We start from a setup with an area of 9×20 m and $N \in \{5, 6, 7, 8, 9\}$ anchors distributed approximately uniformly on the perimeter of the area. When N is odd, one anchor is in the center of the area, while the others are on the perimeter. We consider a grid of ≈ 1700 true locations of the tag spaced 20 cm apart within the area encompassed by the anchors. In a real deployment with a reasonable location update period of 100 ms, 1700 locations could be obtained in under 3 min. For each true location of the tag, we choose at random $\lfloor N * Q \rfloor$ anchors to be in NLOS with the tag, where $Q \in \{0, 0.3, 0.5\}$.

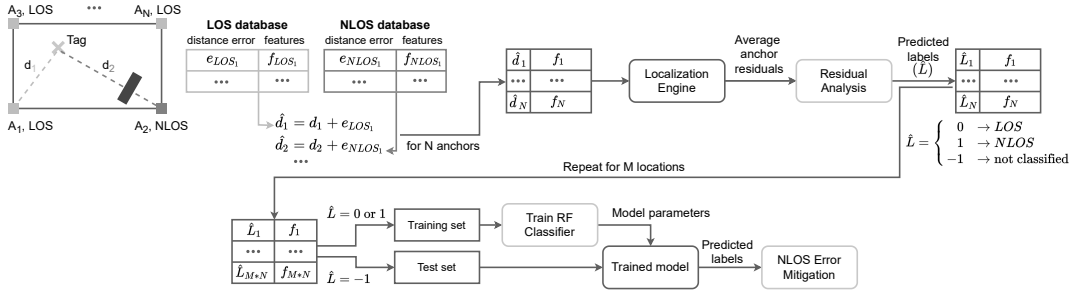


Fig. 5. Flowchart of the simulation framework for NLOS detection and mitigation.

We simulate the distance measurements between each anchor and the tag by adding to the true distance a distance error selected from a measurement database, based on whether the anchor is in LOS or NLOS with the tag. We also store the CIR corresponding to the selected measurement. Although we could simulate the distance errors for each obstacle based on proposed models [15], it is harder to simulate the CIR for a particular type of obstruction. Therefore, for a realistic setup, we preferred selecting the distance error and CIR of a real measurement from a database. We use a database of distance measurements and CIRs acquired with UWB devices developed by 3db Access, which was partly described in [15]. The measurements were acquired in LOS, in NLOS with human body shadowing, or in NLOS with concrete wall shadowing at various indoor locations. The NLOS database aggregates the measurements with both types of obstructions. TABLE I shows the number of measurements acquired in each scenario and the range of covered distances.

We feed the N distances to a localization engine, which computes the 2D location and average anchor residuals over all anchor subsets. For localization, we used the Gauss-Newton multilateration algorithm strengthened with a regularization term. We initialized the algorithm with $\delta = 1$ mm, $k_{max} = 10$ iterations, \mathbf{x}_r = the median of the anchors' locations, and $c = 10^{-1}$ (corresponding to a standard deviation of 10 m around \mathbf{x}_r , suitable for our setup).

The residual analysis block receives as input the average anchor residuals and predicts the labels of each anchor-tag measurement. The label can be either LOS, NLOS, or ambiguous (in case the density of the anchor residuals does not have exactly one minimum). We repeat the procedure for M locations and build a database of $M * N$ predicted labels and the corresponding CIR features.

We split the database into a training set, which contains the measurements predicted as LOS/NLOS, and a test set, which contains the ambiguous measurements. We use the training set to train a Random Forest classifier, which learns the LOS/NLOS CIR features based on the labels predicted by the residual analysis. Once the model is trained, we can use it to directly classify all measurements, without going through the residual analysis procedure. Finally, we mitigate NLOS

 TABLE I
 DESCRIPTION OF THE DATABASE WITH DISTANCE MEASUREMENTS.

Scenario	Distances	Number of measurements
LOS	1–8 m	16 814
NLOS with concrete wall	1–10.5 m	13 770
NLOS with human body	1–10 m	8462

measurements and reduce the localization error.

V. EVALUATION

We now evaluate the performance of the residual-based labeling (Section V-A), of the trained RF classifier (Section V-B), and of the NLOS mitigation method (Section V-C).

A. Residual-based LOS/NLOS Labeling

The first step is the unsupervised labeling using anchor residuals, described in Section III-B. We can alternatively formulate the labeling as a detection problem, where the detected event is a NLOS measurement. To evaluate the performance of the labeling method, we use the balanced accuracy, which is the arithmetic mean of the true positive and true negative rates (TPR and TNR, respectively):

$$\text{Balanced accuracy} = (\text{TPR} + \text{TNR}) / 2, \quad \text{where} \quad (10)$$

$$\text{TPR} = \text{TP} / (\text{TP} + \text{FN}), \quad \text{and} \quad (11)$$

$$\text{TNR} = \text{TN} / (\text{TN} + \text{FP}). \quad (12)$$

TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives.

Because we label only cases where the density estimate of the anchor residuals has exactly one minimum, we are also interested in the percentage of instances labeled (denoted by P_C), defined as the number of locations for which the algorithm provides a label for all anchors. This ratio and the accuracy depend on the KDE shaping parameter (h). If h is too small or too large, the estimation is oversmoothed or undersmoothed, respectively, resulting in few classified instances (because the distribution has either too many or no




Fig. 6. The balanced accuracy and percentage of classified instances of the residual-based labeling as a function of the KDE shaping parameter for 30 % and 50 % NLOS anchors out of $N \in \{5, 6, 7, 8, 9\}$ anchors.

minima at all). Therefore, we want to find the KDE shaping parameter which maximizes the accuracy and P_C .

Fig. 6 shows the balanced accuracy and percentage of classified instances against h for 30 % and 50 % NLOS anchors. We omitted the case when all anchors are in LOS with the tag because less than 1 % instances are classified in this case. This is desirable for the RF classifier because it keeps the number of LOS and NLOS measurements balanced.

Fig. 6a and 6b show that h changes the balanced accuracy with at most 15 % for $Q = 30\%$ NLOS anchors and at most 10 % for $Q = 50\%$ NLOS anchors. On the other hand, Fig. 6c and 6d show that h has a marked impact on the percentage of classified instances. P_C is the highest for $h = [0.03, 0.06]$ (depending on the number of anchors) and decreases for values outside the interval. Because, for different values of h , the change in accuracy is small but the change in P_C is large, we want a performance score based on these two metrics which increases the impact of P_C . Therefore, we chose as the aggregated performance score the harmonic mean between the balanced accuracy and P_C .

In practice, we usually have mixed NLOS conditions which can change over time. Therefore, we extend the performance score to be the harmonic mean of the balanced accuracy and percentage of classified instances for both $Q = 30\%$ and 50% and choose the KDE shaping parameter which maximizes this score. The optimum shaping parameter is $h = 0.04$ for $N = 5$ to 8 anchors and $h = 0.03$ for $N = 9$ anchors. TABLE II shows the LOS, NLOS, and balanced accuracy and the percentage of classified instances for the optimum h .

The classification accuracy is higher for 30 % NLOS anchors than for 50 %. When more anchors are in NLOS with the tag, the location estimate is more skewed and the residuals of all anchors (not only of NLOS anchors) are larger. In this case, LOS and NLOS anchor residuals are harder to distinguish.

We note that the accuracy slightly decreases for more anchors. This is because, with more anchors (out of which only a few are in NLOS), it is harder to find a value of h low enough to delimit the few NLOS anchors from the LOS ones but high enough to avoid an oversmoothed distribution that leads to more than two intervals.

B. Supervised Classification

We now evaluate the accuracy of the RF classification applied on measurements labeled with residual analysis. The training set consists of labeled measurements and their features. We aggregated the measurements for $Q = 30\%$ and

TABLE II
PERFORMANCE OF UNSUPERVISED RESIDUAL-BASED LABELING FOR THE OPTIMUM h .

N	Q [%]	Accuracy [%]			Classified instances [%]
		LOS	NLOS	Balanced	
5	30	68.2	98.7	83.4	37.6
	50	66.8	82.1	74.4	40.2
6	30	68.3	98.9	83.6	29.3
	50	72.0	76.5	74.2	39.6
7	30	64.7	84.5	74.6	34.6
	50	62.4	79.5	70.9	37.6
8	30	74.4	81.1	77.7	32.8
	50	69.9	73.3	71.6	38.1
9	30	70.2	79.2	74.7	36.8
	50	66.0	72.5	69.2	35.2

50 % NLOS anchors, since in practice we can have mixed NLOS conditions. We train a classifier for each number of anchors. The test set contains all measurements which were not labeled in the previous step, i.e., where the density of anchor residuals did not have exactly one minimum. The training set has approximately 6,500–10,000 samples and an almost equal number of LOS and NLOS samples. We use stratified K-fold cross-validation with $K = 4$ folds to identify the best model’s parameters from a specified subset.

Fig. 7 shows the LOS, NLOS, and balanced accuracy of the model. The NLOS detection accuracy slightly exceeds 90 % in all cases, while the LOS accuracy exceeds 95 %. Compared with only the residual-based labeling, we gain 10–20 % accuracy. It is perhaps surprising that the RF classification accuracy exceeds 90 % even when the labeling accuracy can be as low as 63 %. This is because noisy labels resemble outliers or anomalies and ML models can usually recover to a certain extent the correct class boundaries [16].

C. NLOS Mitigation

We now devise a strategy for handling NLOS measurements in order to improve the localization accuracy. We present the NLOS mitigation procedure as a pseudocode in Algorithm 1 and describe it in the following.

If there are enough LOS anchors to compute one location (at least $D + 1$ anchors for D dimensions) and the set of anchors is not degenerate (i.e., the anchors are not collinear), we can use only the LOS anchors to compute the location.




Fig. 7. The accuracy of the trained random forest (RF) classifier using data labeled with the residual method.

Algorithm 1 NLOS Mitigation

```

1:  $\mathbf{x}_{A_i}$  = Location of anchor  $A_i$ ,  $i = 1, \dots, N$ 
2:  $d_i$  = Distance from anchor  $A_i$  to the tag
3:  $S_{LOS} = \{A_i | A_i \text{ is in LOS with the tag}\}$ 
4:  $D$  = Number of dimensions
5: if ( $|S_{LOS}| \geq D + 1$ ) and ( $S_{LOS}$  is not degenerate) then
6:    $\hat{\mathbf{x}}$  = compute_location( $\{d_i | A_i \in S_{LOS}\}$ )
7:   for  $i = 1, \dots, N$  do ▷ Distance correction
8:     if  $A_i$  in NLOS then
9:        $r_i = \|\mathbf{x}_{A_i} - \hat{\mathbf{x}}\|$  ▷ Anchor residual
10:       $\tilde{d}_i \leftarrow d_i - r_i$  ▷ Corrected distance
11:    return  $\hat{\mathbf{x}} \leftarrow$  compute_location( $\{d_i, i = 1, \dots, N\}$ )
12: else
13:    $S$  = The set of all anchor subsets using all LOS
14:   anchors and all combinations of NLOS anchors
15:   for  $S_k$  in  $S$  do
16:      $\hat{\mathbf{x}}_k =$  compute_location( $\{d_i | A_i \in S_k\}$ )
17:      $R_k = \frac{1}{|S_k|} \sum_{\{i | A_i \in S_k\}} r_i^2$  ▷ Subset residual
18:   return  $\hat{\mathbf{x}} = \left( \sum_{k=1}^{|S|} \hat{\mathbf{x}}_k R_k^{-1} \right) / \left( \sum_{k=1}^{|S|} R_k^{-1} \right)$ 

```

However, if there are few LOS anchors and their placement is not ideal (for instance, the tag falls outside the convex hull of the anchors), the location estimated using only the LOS anchors can sometimes have large errors. Therefore, we noticed that we obtain better location estimates if we correct NLOS measurements and use them for localization. For correction, we first estimate the intermediate location using only the LOS anchors. We compute the residuals of the NLOS anchors based on the intermediate location. Then, we subtract the residuals from the measured distances of NLOS anchors. We estimate the final location using the distance measurements of LOS anchors *and* the corrected distances of NLOS anchors.

If the set of LOS anchors is degenerate or there are not enough LOS anchors to compute the location, we must use some NLOS anchors to compute the tag’s location. Because we cannot correct the NLOS measurements as in the previous case, we generate all subsets S_k containing all LOS anchors and all combinations of NLOS anchors such that $|S_k| \geq D + 1$. For each subset, we compute the intermediate location and the subset’s residual using Eq. 7. The final location is the weighted linear combination of all intermediate locations, similar to the method proposed in [1], except that we do not use all possible

TABLE III
LOCALIZATION ERROR WITH NLOS MITIGATION

N	Method	Q = 30% NLOS		Q = 50% NLOS	
		Mean [m]	Std. dev. [m]	Mean [m]	Std. dev. [m]
5	No mitigation	0.14	0.27	0.20	0.39
	Proposed	0.07	0.06	0.11	0.14
6	No mitigation	0.11	0.08	0.19	0.31
	Proposed	0.06	0.04	0.10	0.17
7	No mitigation	0.14	0.30	0.21	0.64
	Proposed	0.07	0.05	0.08	0.07
8	No mitigation	0.14	0.36	0.22	0.67
	Proposed	0.06	0.05	0.08	0.06
9	No mitigation	0.13	0.35	0.20	0.64
	Proposed	0.05	0.04	0.07	0.14

subsets, thus reducing the complexity of the algorithm.

We compute the localization error as the Euclidean distance between the estimated location $\hat{\mathbf{x}}$ and the true location \mathbf{x} :

$$e = \|\hat{\mathbf{x}} - \mathbf{x}\|. \quad (13)$$

We evaluate the localization error on the test data set from Section III-C. For each set of anchor–tag measurements, we predict the LOS/NLOS condition using the trained model and then apply Algorithm 1 to mitigate the NLOS errors. TABLE III compares the localization errors obtained with the proposed mitigation algorithm (denoted by “proposed”) with those obtained when using all anchor–tag distances, without mitigation. The localization errors of the proposed method have 1.8–2.8× smaller mean and 1.8–11.6× smaller standard deviation after NLOS mitigation. On average, the algorithm reduces the mean and standard deviation by 2.2 and 5.8 times, respectively. Therefore, the proposed method can successfully mitigate localization errors caused by NLOS propagation.

VI. DISCUSSION AND FUTURE WORK

One remaining question is whether it is worth deploying more anchors than the minimum necessary (e.g., three anchors for 2D localization). In buildings with rooms separated by thick walls, if we want to provide high-accuracy location services everywhere, we have to deploy at least three anchors in each room. However, the tag can still be in the range of anchors in other rooms, so it has to decide which anchors are in LOS. Even within the same room, some anchors might be shadowed by surrounding objects, so it is worth having extra anchors. In TABLE III, we see that when we correct NLOS errors, the localization error decreases for more anchors. Therefore, more anchors than the minimum are often needed in practical deployments.

So far, we have not discussed which entity should train the classifier: the localization engine (LE) or the tag. There are arguments for both sides. On the one hand, NLOS errors depend on the environment in which the LE operates. For instance, in an industrial setting with metallic objects, NLOS errors might be larger than in an office. Therefore,

the LE could collect anchor–tag distances from tags operating in an area and train a model which can then classify all measurements at the LE or the tag, depending on which entity computes the location. On the other hand, the learned CIR features can also depend on the hardware of the end device. For instance, the CIR can have different shapes depending on the type of UWB device [15]. Therefore, models trained on features from one hardware model might not generalize well to others. One future research direction is to evaluate how well a model trained on one type of hardware or at one location generalizes to other device models and environments. If CIR features are indeed model- or environment-specific, one option is to train an initial model and periodically update it with new data from different devices and locations. This is also beneficial if the training is performed by the tag, since it requires less data storage and the model can be updated online. Since the residual labeling step outputs labels only for a part of the input measurements, the accuracy of the model could be improved using semi-supervised learning methods. For instance, the model can be retrained using its most confident predictions [17], which can also speed up the training process.

The localization error can also be reduced by applying Chen’s residual weighting algorithm [1] or variations of it, without going through the labeling and training process. However, computing the location using all anchor combinations has a complexity of $\mathcal{O}(2^N)$ for N anchors. In our case, the complexity is high only during the short training phase. After this, samples can be classified with the RF with a constant complexity of $\mathcal{O}(kp)$, where k is the number of decision trees and p is the maximum depth of one tree. In practice, the execution time of the proposed NLOS mitigation method (excluding the training phase) was faster than the residual weighting algorithm for $N \geq 8$ anchors.

NLOS identification has interesting applications beyond reducing localization errors, especially when it can be done without supervision, as in our proposal. Our method can potentially be used to build maps of a building by aggregating the locations at which the tag is consistently in NLOS with certain anchors. Crowd density estimation is another interesting possible application. Since human body shadowing introduces large distance errors, any increase in the number of detected NLOS measurements could suggest that a room gets more populated. Note that this method does not require all users to be connected to the localization network. Finally, the trained model can be applied on individual distance measurements, so it can be useful in peer-to-peer proximity applications (e.g. contact tracing, object finding).

VII. CONCLUSIONS

We proposed a method for detecting NLOS measurements in localization systems without manually-acquired training data or knowledge of channel statistics. The method predicts the LOS/NLOS labels using the measured distances between each anchor and the tag. We use the predicted labels and the CIR features of the measurements to train a classifier, which has over 90 % classification accuracy. We also proposed a NLOS

mitigation technique which reduces, on average, the mean and spread of the localization error by at least $2\times$.

REFERENCES

- [1] P.-C. Chen, “A non-line-of-sight error mitigation algorithm in location estimation,” in *WCNC. 1999 IEEE Wireless Communications and Networking Conference*, vol. 1, pp. 316–320, IEEE, 1999.
- [2] S. Marano, 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*, vol. 28, no. 7, pp. 1026–1035, 2010.
- [3] 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*, vol. 60, no. 6, pp. 1719–1728, 2012.
- [4] J. Schroeder, S. Galler, K. Kyamakya, and K. Jobmann, “NLOS detection algorithms for ultra-wideband localization,” in *2007 4th Workshop on Positioning, Navigation and Communication*, pp. 159–166, IEEE, 2007.
- [5] R. M. Vaghefi, J. Schloemann, and R. M. Buehrer, “NLOS mitigation in TOA-based localization using semidefinite programming,” in *2013 10th Workshop on Positioning, Navigation and Communication (WPNC)*, pp. 1–6, IEEE, 2013.
- [6] B. J. Silva and G. P. Hancke, “Non-line-of-sight identification without channel statistics,” in *IECON 2020 The 46th Annual Conference of the IEEE Industrial Electronics Society*, pp. 4489–4493, IEEE, 2020.
- [7] G. Wang, W. Zhu, and N. Ansari, “Robust TDOA-based localization for IoT via joint source position and NLOS error estimation,” *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8529–8541, 2019.
- [8] S. Li, M. Hedley, I. B. Collings, and D. Humphrey, “Joint trajectory and ranging offset estimation for accurate tracking in NLOS environments,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 1, pp. 3–14, 2019.
- [9] U. Hammes and A. M. Zoubir, “Robust MT tracking based on M-estimation and interacting multiple model algorithm,” *IEEE Transactions on Signal Processing*, vol. 59, no. 7, pp. 3398–3409, 2011.
- [10] J. Fan and A. S. Awan, “Non-line-of-sight identification based on unsupervised machine learning in ultra wideband systems,” *IEEE Access*, vol. 7, pp. 32464–32471, 2019.
- [11] J. Park, S. Nam, H. Choi, Y. Ko, and Y.-B. Ko, “Improving deep learning-based uwb los/nlos identification with transfer learning: An empirical approach,” *Electronics*, vol. 9, no. 10, p. 1714, 2020.
- [12] T. Wang, K. Hu, Z. Li, K. Lin, J. Wang, and Y. Shen, “A semi-supervised learning approach for uwb ranging error mitigation,” *IEEE Wireless Communications Letters*, 2020.
- [13] L. Cheng, Y. Li, Y. Wang, Y. Bi, L. Feng, and M. Xue, “A triple-filter NLOS localization algorithm based on fuzzy c-means for wireless sensor networks,” *Sensors*, vol. 19, no. 5, p. 1215, 2019.
- [14] N. Sirola, “Closed-form algorithms in mobile positioning: Myths and misconceptions,” in *2010 7th Workshop on Positioning, Navigation and Communication*, pp. 38–44, IEEE, 2010.
- [15] L. Flueratoru, S. Wehrli, M. Magno, E. S. Lohan, and D. Niculescu, “High-accuracy ranging and localization with ultra-wideband communications for energy-constrained devices,” *IEEE Internet of Things Journal*, pp. 1–1, 2021.
- [16] B. Frénay and M. Verleysen, “Classification in the presence of label noise: a survey,” *IEEE transactions on neural networks and learning systems*, vol. 25, no. 5, pp. 845–869, 2013.
- [17] I. Triguero, S. García, and F. Herrera, “Self-labeled techniques for semi-supervised learning: taxonomy, software and empirical study,” *Knowledge and Information systems*, vol. 42, no. 2, pp. 245–284, 2015.

PUBLICATION 5

On the High Fluctuations of Received Signal Strength Measurements with BLE Signals for Contact Tracing and Proximity Detection

L. Flueratoru, V. Shubina, E. S. Lohan, and D. Niculescu

IEEE Sensors Journal 22.6 (2021), pp. 5086–5100

DOI: 10.1109/JSEN.2021.3095710

Publication reprinted with the permission of the copyright holders.

On the High Fluctuations of Received Signal Strength Measurements with BLE Signals for Contact Tracing and Proximity Detection

Laura Fluera^{1,2,*}, Viktoriia Shubina^{2,1,*}, Dragoş Niculescu¹, Elena Simona Lohan²
¹Computer Science Department, Politehnica University of Bucharest, Bucharest, Romania
²Electrical Engineering Unit, Tampere University, Tampere, Finland
 Emails: ¹name.surname@tuni.fi and ²name.surname@upb.ro

Abstract—This paper presents a measurement-based analysis of the Received Signal Strength (RSS) of Bluetooth Low Energy (BLE) signals, under Line-of-Sight (LOS) and Non-Line-of-Sight (NLOS) scenarios, performed in tandem at two universities in Tampere, Finland, and Bucharest, Romania. We adopted the same hardware and methodology for measurements in both places, and paid particular attention to the impact of RSS on various environmental factors, such as LOS and NLOS scenarios and interference in 2.4 GHz band. In addition, we considered the receiver orientation and the different frequencies of BLE advertising channels. We show that snapshot RSS measurements typically have high variability, not easily explainable by classical path-loss models. A snapshot recording is defined here as one continuous recording at fixed device locations in a static setup. Our observations also show that aggregated RSS data (i.e., considering several snapshot measurements together) is more informative from a statistical point of view and more in agreement with current theoretical path-loss models than snapshot measurements. However, in BLE applications such as contact tracing and proximity detection, the receivers typically have access only to snapshot measurements (e.g., taken over a short duration of 10–20 minutes or less), so the accuracy of contact-tracing and proximity detection can be highly affected by RSS instabilities. In addition to presenting the measurement-based BLE RSS analysis in a comprehensive and well-documented format, our paper also emphasizes open challenges when BLE RSS is used for contact tracing, ranging, and positioning applications.

I. INTRODUCTION AND MOTIVATION

Proximity-based applications have become increasingly popular in recent years. Estimating the distance between two devices can be used to find lost objects, to share files between nearby devices, to enable smart homes to react to owners' location, or to fight against a pandemic. In the past year, digital contact-tracing applications (shortly called apps) have received increasing attention to prevent the spread of COVID-19 and many countries have developed such digital apps. Detailed overviews of existing contact-tracing apps can be found in our previous works [1], [2] and in other recent works [3]–[7].

The most popular technologies that enable proximity-based applications are Wi-Fi, Bluetooth Low-Energy (BLE), Ultra-Wideband, and Global Navigation Satellite Systems (GNSS). BLE is the most promising candidate of them since it

offers the lowest power consumption and is supported by most mobile devices and operating systems [8]. While BLE-based positioning can reach meter-level accuracy when both angle-of-arrival (AOA) and received signal strength (RSS) information are combined [9], most consumer devices such as mobile phones and wearables are not equipped with direction-finding capabilities and they rely only on RSS measurements for proximity detection. RSS measurements from any wireless signal (BLE, Wi-Fi, cellular, etc.) are known to fluctuate due to the presence and movement of people in the signal's path [10], the presence of multipath [11], the switches between carrier frequencies of sub-channels used in the transmitted signal [11], the antenna polarization [12], the orientation of the transmitter (TX) and receiver (RX) [13], and the chipset model [14].

While there are currently many studies about the RSS variability in Wi-Fi signals, e.g., [15]–[19], most such studies focus only on one source of fluctuations or investigate the aggregated effect of multiple error sources. In contrast, this paper documents the (in)stability of the BLE RSS over time, over space (with different multipath characteristics), with different hardware, on different advertising channels, at different distances, with different device orientations, and with different type of obstructions between device pairs. We isolated these factors and evaluated their impact individually. In addition, we documented a new error source, namely the influence on Wi-Fi–BLE combo chipsets on the RSS. Based on an extensive measurement campaign, we provide recommendations that can partly mitigate BLE fluctuations caused by these factors.

We also provide open-access data that accompanies this study in order to aid future research. During the COVID-19 pandemic, open-access BLE RSS data sets have proven essential for the research community. However, most such data sets, e.g., [20]–[22], have limited documentation or do not analyze the behavior of the BLE RSS with all the aforementioned instability sources. Therefore, a more thorough investigation on BLE RSS instabilities documented by open-source data is still needed.

This paper offers a comprehensive analysis of BLE RSS instabilities, fluctuations, and challenges in BLE-based proximity detection and contact tracing. We based the analysis on two extensive measurement campaigns performed

* The first two authors had equal contribution to this paper.

in parallel at Tampere University (TAU) in Tampere, Finland and at University Politehnica of Bucharest (UPB) in Bucharest, Romania between January–March 2021. The tandem measurements were conducted with exactly the same type of devices to eliminate the possible fluctuations coming from different hardware models as well as possible calibration issues. Our measurements will be available, upon the paper publication, in open-access at the A-WEAR research community on Zenodo¹.

The main contributions and findings of this paper are:

- Offering an extensive measurement-based analysis of BLE RSS fluctuations and showing that current single-slope path-loss models from the literature do not capture these effects.
- Comparing snapshot (or single recording) measurements with aggregated recordings and showing that, when enough RSS data is aggregated, the statistics converge to stable models;
- Analyzing the effect of BLE advertising channels on RSS fluctuations and showing that the aggregated RSS from all BLE advertising channels has significantly higher fluctuations than on individual BLE channels. This is an important challenge in current BLE-based proximity apps, where channel information is usually not available;
- Analyzing the effect of non-line-of-sight (NLOS) propagation on the BLE RSS;
- Analyzing the effect of relative orientations between the transmitter and the receiver on the RSS;
- Analyzing the same-chip Wi-Fi interference with BLE.

Based on our state-of-the-art review (Section II), we believe that these high BLE RSS fluctuations have not yet been reported and documented to their full extent in the current literature and that there are still several challenges to be overcome when dealing with snapshot BLE RSS measurements, as those used in contact-tracing and proximity detection applications. Therefore, this paper documents BLE RSS fluctuations and raises several research questions about the applicability of classical path-loss models in the line-of-sight (LOS) and NLOS propagation of BLE signals.

II. STATE-OF-THE-ART OVERVIEW

In Section II-A we provide an overview of the state-of-the-art in digital contact tracing and proximity detection apps based on BLE signals, which are increasingly relevant in our times. In Section II-B we discuss the main factors that cause BLE RSS variability and the most important studies that have investigated them. In Section II-C we summarize the findings and state the key points that differentiate our work from past research.

A. BLE-based Contact Tracing and Proximity Detection Principles

Digital contact tracing is a particular case of proximity detection, used as an identification and follow-up solution

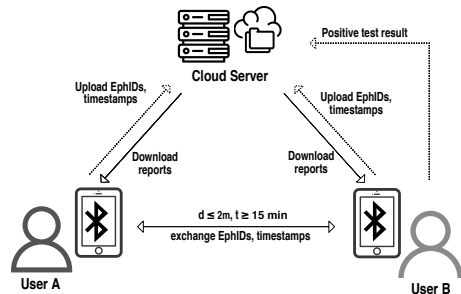


Fig. 1. An illustration of the contact-tracing chain with users *A* and *B* exchanging BLE signals at a distance of at most 2 m and interacting with the cloud server to receive the anonymized reports for crossing paths with infected users.

aiming to break the transmission chains of airborne infections within communities.

In a digital contact-tracing chain such as in Fig. 1, smartphones and wearables are commonly assigned with permanent and temporary identifiers generated by each device for privacy-preserving purposes. The server owns the complete list of the users reporting their confirmed cases of infection, which includes both permanent and ephemeral IDs. Periodically, the user devices receive anonymized data with user reports of confirmed test results from the server, such as the case of the user *A* in Fig. 1, and then locally estimate the risk of having been exposed to the infection.

A device equipped with a BLE chipset starts to log the ephemeral IDs and timestamps of other users when these are nearby (within a distance d) for a certain time window (e.g., typical thresholds used in many apps nowadays are 15 min time window and $d = 2 \text{ m}$ distance, which is currently deemed a safe distance). The infection risk is computed based on the time spent in proximity with a confirmed case.

By nature, BLE signals are susceptible to the environment and therefore require calibration and averaging. When the range is estimated with a certain error, there is a higher risk of generating false positives, when a user appears to be closer than in reality, or false negatives, when the actual distance is less than the estimated one. These errors could also appear if a wall or a door blocks the space between the devices, leading to NLOS propagation, when in fact the infection risk is low. Therefore, it is crucial to accurately estimate the range between two users. When the estimation fully relies on BLE RSS measurements, it is therefore important to understand the various causes of BLE RSS fluctuations.

B. Related Studies on BLE RSS Variability

BLE was primarily designed for communication purposes and its use as a ranging technology has appeared only recently. As any wireless signal, BLE signals are susceptible to environment dynamics such as multipath, signal scattering, shadowing, refraction, or attenuation. In addition, the difficulty of evaluating the exact distance between two persons might

¹<https://doi.org/10.5281/zenodo.4643668>




Fig. 2. An illustration of the 2.4GHz ISM band channels. Advertising channels 37, 38, and 39 are scattered deliberately to avoid interference with Wi-Fi.

be exacerbated by noisy measurements, faulty BLE chipsets, low transmit power, low received signal strength, or infrequent scanning intervals [2].

One of the key challenges of digital contact tracing, which is the scope of on-going research, is the high false positive rate. This occurred, for example, when experts from The Alan Turing Institute used the GAEN system to build the National Health Service (NHS) COVID-19 app [23]. The authors reported a problem of high false positive rates in detecting distances between users staying apart from 2 to 4 m; in other words, 2 m distance proved to be a reliable threshold both for epidemiological safety measures and for BLE performance specifications. Another critical goal for contact tracing is accurate LOS and NLOS detection, yet many factors are still unknown regarding BLE signal propagation. In the following, we outline some of the most important challenges in proximity detection based on BLE RSS and the state-of-the-art concerning them.

1) *Advertising on different carrier frequencies:* BLE uses 40 radio frequency (RF) channels, 2 MHz wide each and assigned with a unique index illustrated in Fig. 2. BLE channels are divided into two groups: advertisement channels (indexed 37, 38, and 39) and data channels (indexed from 0 to 36). In BLE, the three advertisement channels indexed 37, 38, and 39 with center frequencies at 2.402, 2.426, and 2.48 GHz, respectively, are scattered over the 2.4 GHz band to avoid interference with other devices operating in the Industrial, Scientific, and Medical (ISM) band. Based on the analysis and modeling of these advertising channels in [24], channel 39 was deemed the most reliable, since it is further away from the center frequency of a main Wi-Fi channel, whereas channels 37 and 38 overlap with one, respectively two Wi-Fi channels.

The impact of advertising channels on the RSS is twofold. First, according to path-loss models, the RSS is inversely proportional to the squared carrier frequency. Second, embedded antennas usually do not have a flat response over the entire bandwidth, resulting in different gains depending on the frequency [25]. The difference between RSS values acquired at the same location on different channels was found to be as high as 15 dB in [11] or almost 6 dB in [25], therefore decreasing the RSS-based ranging accuracy.

Knowing the channel on which a beacon was transmitted

can improve distance estimates [11], [26], but this information is often obfuscated by the driver at the receiver, unless the transmitter explicitly includes this information in the beacon's payload (which is rarely done). As a result, most receivers cannot recover the advertising channel index on which a beacon was transmitted. Smartphones usually switch between all three advertising channels, resulting in RSS fluctuations. In [11], the authors proposed a method for identifying the advertising channel at the receiver by exploiting the pattern with which some smartphone models switch between the advertising channels.

2) *Multipath propagation:* Multipath propagation causes radio signals to arrive at the receiving antenna via multiple paths due to reflection, refraction, or scattering [37]. Signal components arriving through different paths can add up constructively or destructively, the latter resulting in multipath fading. The channel-dependent multipath fading of BLE signals was studied in [25]. Channels experience deep fades at different locations due to their different center frequencies. The effect of multipath fading was eliminated in a training phase by averaging the RSS in a window. In that case, window sizes of 0.5 s to 2 s mitigate fading effects for walking speed at a BLE packet reception rate of 25 Hz. However, in practice, such a high advertising rate is uncommon as it increases the energy consumption, so observation windows need to be longer to mitigate multipath fading.

The authors in [28] noticed RSS fluctuations on the order of 6 dB at the same transmitter–receiver (TX–RX) distance due to the presence of multipath and Wi-Fi interference. The authors in [29] also noticed fluctuations as large as 25 dBm over short periods of time, in particular for channels 38 and 39 due to channel-dependent fast fading.

In [34], the authors studied RSS fluctuations at various TX–RX locations and noticed that the average RSS is not always decreasing with distance, as predicted by path-loss models [27], but they observed that the average RSS at 2.5 m was consistently higher than the average RSS at 2 m, also when measurements were done with different BLE transmitters. They also observed signal fluctuations as high as 20 dB at constant TX–RX locations, due to human movement around the BLE transmitters.

3) *Orientation:* The way people are holding their mobile devices (e.g., inside front of back pockets, in hand, inside a bag, etc.) influences the relative orientation between transmitter and receiver antennas. These orientation changes can, in turn, cause RSS fluctuations. Fluctuations of up to 30 dB between maximum and minimum RSS at constant TX–RX distances were observed in [31] when RSS data acquired with different device orientations was aggregated.

The authors in [13] found that different device orientations can affect the RSS with differences of up to 3 dB at exactly the same TX–RX distance, and that an RSS at 3 m TX–RX distance can be higher (with few dBs) than the RSS at 1 m TX–RX distance, if different receiver orientations are used.

4) *Transmit power:* The RSS also depends on the transmission power, the RF front-end characteristics, and the antenna gain. Because these factors depend on the hardware or implemented firmware, the observed RSS from devices from

TABLE I
OVERVIEW OF STATE-OF-THE-ART: MEASUREMENT-BASED BLE RSS STUDIES AND DATASETS.

Reference	Year	BLE RSS provided in open access?	Measurement devices (transmitter side)	Studied effects related to RSS
[27]	2015	No	StickNFind (SNF) beacons	RSS fluctuations based on TX–RX distance.
[28]	2016	No	RN4020 PPDB board with RN4020 BLE chipsets	RSS fluctuations based on TX–RX distance and BLE channels 1, 10, 20, and 30 (non-advertising channels).
[29]	2018	No	iBKS105 BLE beacons	RSS fluctuations based on multiple carrier frequencies on advertising channels (i.e., channels 37,38, and 39); orientation effects are also discussed but directly in the context of positioning, not as effects on RSS.
[24]	2018	No	nRF52840 wireless System on Chip (SoC)	Advertising channels characterization in the BLE standard.
[13]	2018	No	A smartphone and Estimote BLE beacons	RSS fluctuations based on different TX power levels, device orientations, advertising intervals, LOS/NLOS cases, density of the devices.
[30]	2018	Yes	Gimbal SEries 10 BLE iBeacons	One-month measurement campaign generated from 10 BLE beacons transmitting signal and carried by people inside a university building, reporting a realistic scenario.
[31]	2019	No	Raspberry Pi 3 (Model B)	RSS fluctuations based on receiver orientation, transmit power, and TX–RX distance.
[32]	2019	Yes	Accent Systems' IBKS 105 with Nordic nRF51822 BLE chipset	None (database was used to study BLE-based positioning, but RSS fluctuations are not studied separately).
[11]	2020	No	8 Android mobile phones and one iPhone 6	RSS fluctuations based on multiple carrier frequencies on advertising channels and on the TX–RX distance.
[33]	2020	Yes	Nokia 8.1 with Android 10, HTC M9 with Android 7.0 Nougat	RSS fluctuations based on the TX–RX distance with respect to various transmitter orientations/placements on the body.
[34]	2020	No	3 Android mobile phone models	RSS fluctuations based on the TX–RX distance, on the presence of human bodies (and their orientation) around TX and RX devices, on a NLOS scenario due to wall presence, and on mobility of persons carrying the mobile phones.
[35]	2020	Yes	Raspberry Pi 3 (Model B)	Three scenarios of different room sizes and with the use of Zigbee, BLE, and WiFi are documented. The authors reported different techniques for precise and accurate location estimates, where K-Nearest Neighbor (KNN) was chosen as an optimal solution.
[20]	2020	Yes	iPhone 10, Ubetooth One, nRF52 eval board	RSSI dataset collected in different environments, with various device orientations and body placements.
[21]	2020	Yes	Raspberry Pi Zero W, Raspberry Pi 3, and Raspberry Pi 4	RSSI data collected at different distances via GAEN.
[22]	2020	Yes	Samsung Galaxy S7, iPhone 7, iPhone 11 Pro	Data set of measurements collected in the university environment, interchanging device pairs used as TX and RX.
[36]	2021	Yes	Android smartphones: Nokia 8.1 with Android 10, HTC M9 with Android 7.0 Nougat	Report of RSSI fluctuations with different device placements on human bodies; described the effect of window size on the accuracy of the estimates.
This article	-	Yes	Raspberry Pi 4 (Model B+, Cypress CYW43455)	RSS fluctuations based on LOS/NLOS scenarios, NLOS cases with different obstacles, receiver orientations, multiple carrier frequencies on BLE advertising channels, on-chip BLE and Wi-Fi interference in 2.4GHz band, hardware instabilities, and the TX–RX distance; also the test–retest reliability of measurements is addressed here.

different manufacturers can vary even when the environmental conditions are identical. This behavior was documented in [38] where, even though transmitters from different vendors had different TX powers, the RSS was within the same range. In [39] it was shown that the transmission power influences the localization accuracy and the authors proposed machine learning models to identify the individual TX power of the deployed beacons that maximize the localization accuracy.

One way to solve this issue is to compute RSS correction factors at the transmitter and the receiver [40]. The calibrated TX power can be measured for a particular model of transmitter at a known distance (e.g., 1 m for the iBeacon standard and 0 m for EddyStone) and sent in the payload of the advertising beacon. For instance, if a transmitter has a calibrated TX power of -45 dBm at 1 m, an RSS of -55 dBm will indicate that the receiver is at more than 1 m away from the transmitter, whereas for another device model -55 dB might be the calibrated TX power. Similarly, each receiver should have a correction coefficient that reflects the receiver efficiency, or with how much its RSS deviates from a reference value. Ideally, there should be a database with RSS correction factors for each mobile device. However, such a task is intractable because of the sheer number of mobile devices on the market. A 2015 report counted more than 24,000 Android devices made by almost 1300 companies [41]. Moreover, as we will show in Section IV-C, this does not account for RSS variations between devices from the same model. To the best of the authors' knowledge, the RSS variability within devices from the same vendor has not been documented in the literature.

5) *Non-line-of-sight between the devices*: RF signals propagate at a different speed through the air than through obstacles such as walls, furniture, or the human body. Therefore, obstructions between the transmitter and the receiver will typically cause fluctuations in the RSS. There are several research works [26], [42], [43] that investigated the effect of shadowing on the BLE RSS with applications in proximity detection or localization. [26] proposed artificial neural network (ANN) models for detecting human-body shadowing and compensating RSS values to improve distance measurements or localization based on the BLE RSS. In the best case, the ANN can correctly detect the obstacle more than 87% of the time. The method leverages measurements acquired on *individual* channels, so knowledge of the advertising channel is also required, as well as a training phase for the ANN.

In [43], the authors proposed a NLOS detection method based on the variance of the BLE RSS. The algorithm is able to detect when a concrete wall is blocking the direct path between the transmitter and the receiver with an accuracy of 76.25% based on a fixed threshold of the RSS variance, below which the signal is classified as being acquired in NLOS. The same method could not be applied on NLOS with plasterboards, since the standard deviation was inconsistent. The effect of several obstacles (wooden door, iron door, window, hand, paper) on the BLE RSS was studied in [38]. The mean RSS values obtained with these obstructions varied between -50 dBm to -90 dBm at a TX-RX distance of 2 m. The

strongest attenuation was caused when a hand covered the transmitter and when the LOS was blocked by an iron door. These results show that different NLOS obstacles can have a different impact on the RSS and that the topic should be further explored.

In [34], a NLOS case was analyzed with two types of walls between the TX and RX: a stud partition and a blockwork wall. No differences between LOS and NLOS scenarios were observed for the stud partition, while the blockwork wall introduced attenuations of up to 20 dB in the received signal strength compared to LOS case. The main conclusions in [34] are similar to the ones in our measurement-based analysis, that BLE signals have high fluctuations and their RSS does not necessarily follow classical path-loss models. Therefore, developing accurate BLE RSS-based proximity-detection methods remains a challenging topic.

Changes caused by the human body in wireless signal propagation in the 2.4 GHz band have also been documented in [44]–[46].

6) *Interference in ISM band at 2.4 GHz*: As the ISM band is heavily used by many wireless systems, fluctuations in the BLE RSS are also caused by RF interference, especially coming from shared antennas between Wi-Fi and BLE modules coexisting on the same chipset (as it is the case with most mobile phones). The authors in [28] noticed RSS fluctuations on the order of 6 dB at the same TX-RX distance due to multipath fading and interferences from Wi-Fi.

The authors in [38] performed an experiment in which a BLE TX was placed directly under a Wi-Fi access point (AP) and the RSS was recorded, in turns, when the AP was on and off. When the Wi-Fi AP was on, the reception rate dropped to 75% and the RSS decreased with 10 dB in 50% of the measurements compared to the case in which the AP was turned off. We further explore this topic in Section IV-G.

C. State-of-the-art Summary

The work in [34] can be seen as the closest to our work from the BLE RSS literature (as summarized also in Table I). However, our work focuses only on indoor scenarios in a more systematic approach, by duplicating BLE RSS measurements in two different locations (Tampere and Bucharest), by performing extensive and repetitive tests at distances relevant to contact-tracing apps (i.e., 1 to 3 m), and by investigating the effects of Wi-Fi interference and the three BLE advertising channels.

The main reason we focused on indoor scenarios is that outdoor proximity detection can be achieved with high-accuracy GNSS receivers. For indoor proximity detection, however, there are more viable candidates, out of which BLE is the most promising but also perhaps the most challenging one. In addition, in digital contact tracing apps, infectiousness levels are lower outdoors than indoors [47], [48].

Our paper offers a comprehensive survey of various causes of BLE RSS variability as well as of the related works in the literature. The state-of-the-art main studies on BLE RSS are summarized in Table I and the last row shows the contributions of this article at a glance.




Fig. 3. Raspberry Pi 4 Model B devices are used in our experiments as TX and RX.

Other works in similar spirit but for Wi-Fi, found variations across channels, time scales, interfaces used for 5GHz Wi-Fi [49], and across direction, device manufacturer, sampling period, presence of humans and of other radio devices [50].

For IEEE 802.15.4, that also uses 2.4GHz ISM band, but lower power, [51] finds that the main variability sources when measuring RSS are antenna orientation, hardware sample, and link asymmetry.

III. MEASUREMENT-BASED BLE DATA COLLECTION

In all our experiments, we used Raspberry Pi 4 Model B devices, as illustrated in Fig. 3. The internal 2.4 GHz antenna is located in the left upper corner, next to the Cypress CYW43455 combo Wi-Fi and BLE module. The devices have a 1.5 GHz 64-bit Quad-Core Cortex-A72 CPU in the middle of the Raspberry Pi. The Gigabit Ethernet, two USB 3.0 and two USB 2.0 ports are located on the right, which might cause signal degradation in some TX–RX orientations (which will be discussed in Section IV-F). One advantage of using this hardware is that we can configure the advertising channel and payload, an option that not many open-source smartphone applications offer.

We acquired a database of BLE RSS measurements between devices placed at 1, 2, and 3 m in several LOS and NLOS scenarios with obstructions caused by walls (with and without a whiteboard on it), human body, plexiglass panels, and doors, shown in Fig. 4. We conducted two measurement campaigns in parallel at UPB and TAU. The different locations enabled us to compare and validate measurements acquired with the same hardware models but in different settings.

We define a *measurement* as the process of collecting data in a specific manner. Measurements can be grouped in *recordings* (or *snapshot* measurements), when data is collected continuously from a start time to an end time, in a static setup and without modifying the devices in any way, and in *scenarios* (or *aggregated* measurements), which are collections of recordings according to a pre-defined criterion. For instance, a scenario can be a collection of recordings acquired in LOS, with a TX–RX distance of 1 m, on channel 37.

We configured the transmitter to send non-connectable un-directed advertisements (`ADV_NONCONN_IND`) with a period of 100ms, which satisfies the broadcasting interval

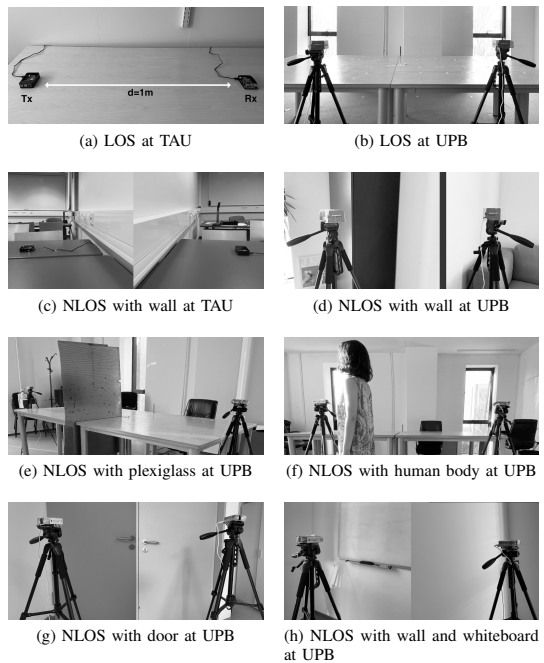


Fig. 4. The pictures of the receiver and transmitter in LOS and NLOS with wall acquired at UPB and TAU (Fig. 4a to 4d). Fig. 4e to 4h show NLOS scenarios with a plexiglass panel, human body, a door, and a wall with a whiteboard at UPB.

recommendation of 200 ms to 270 ms of the Bluetooth protocol for contact tracing developed by Apple and Google [52]. The same specification suggests a scanning period (at the receiver) of at least 5 min, although this is likely to vary depending on the application. For instance, in the GAEN API the scanning period was found to be between 2.5 and 4 min [53]. Since a higher scanning rate provides more RSS samples and the devices are not energy-constrained, we chose a scanning frequency of 1 Hz.

The recording time ranged from 3 minutes to 3 days. In some cases, we were interested in the stability of RSS measurements over a longer period of time, case in which the recording time spanned several days, whereas in other cases we were interested in the variability of RSS measurements at different locations with constant TX–RX distances, case in which shorter recording times of several minutes were more convenient. Fig. 4 shows examples of LOS and NLOS scenarios in which data was acquired at TAU and UPB.

IV. MEASUREMENT-BASED BLE RSS CHARACTERIZATION

This section provides an overview of the results acquired during our experiments and describes the challenges discussed in Section II-B.

In order to compare in a comprehensive manner the RSS distributions in different scenarios, throughout this section we




Fig. 5. Box plot used to describe and compare the RSS distribution.




Fig. 6. The effect of the time evolution on the RSS distribution in LOS on a single channel (37), recorded at UPB.

represent the data using standardized box plots such as the one in Fig. 5, as they give information at-a-glance about the mean, median and spread of the RSS. The box shows where most of the RSS values are found, namely the data from the first quartile (Q_1 or the 25th percentile) to the third quartile (Q_3 or the 75th percentile), also known as the interquartile range (IQR). The lines extending from the box are called whiskers and cover the range from $a \triangleq Q_1 - 1.5 * (Q_3 - Q_1)$ to $b \triangleq Q_3 + 1.5 * (Q_3 - Q_1)$ (corresponding to Tukey's original definition of box plots). The red vertical line inside the box plot denotes the median. In some plots, we also added via a diamond marker, the mean of the data. The circle markers to the right (can also appear to the left) of the whiskers are outliers. Occasionally, the outliers or the mean value are omitted in our plots to preserve a good readability of the plot. In some cases, the RSS is stable enough that the IQR contains only the median value and therefore the box is not shown.

A. The (In)stability of BLE RSS Measurements Over Time

We first investigate the stability of BLE RSS measurements in a particular setting over time. Fig. 6 shows the boxplots of the BLE RSS distribution in windows of 2 min up to 10 h acquired in LOS, on channel 37, at a distance of 1 m between the transmitter and the receiver. The measurements were acquired in a locked room during the weekend, so there was no human activity around the devices during the recording. Although the median RSS changes with up to 3 dB over the course of the recording, the mean RSS varies with less than 1 dB with different window lengths.

It can be seen from Fig. 6 that, if we are interested only in the mean RSS, then a recording time of 2 min is sufficient to obtain the mean RSS that best captures the characteristics of the particular setting in which measurements are acquired. If we are also interested in the shape of the distribution, a longer recording time of at least 30 min is necessary. In general, the RSS during each snapshot recording was stable over time with the exception of some random fluctuations that sometimes appeared at the beginning of a recording and which will be discussed in Section IV-G.

B. The (In)stability of BLE RSS Measurements Over Space or Test-Retest Reliability Studies

Next, we study the stability of the BLE RSS under LOS scenarios, at a fixed distance of 1 m between the same TX-RX pair, and using only the advertising channel 37 in order to eliminate frequency-dependent fluctuations. We acquired measurements at TAU and UPB, in different rooms or with the TX and RX placed in different spots in the same room, while maintaining a distance between the two devices of 1 m. Fig. 7 compares the RSS distribution in 15 recordings when taking a fixed number of 326 random measurements from each recording (the fixed number was selected based on the minimum length among all 15 recordings).

We expected to get similar RSS measurements in different snapshot recordings, given that the multipath fading is mitigated by averaging samples over several minutes. However, even after multiple test-retest measurements performed at UPB and TAU, results (see Fig. 7) indicate fluctuations of the median RSS of up to 40 dB between snapshot recordings even though the TX-RX distance was constant. Moreover, the median RSS can vary even in the same location between two recordings taken in different days, even though results in Section IV-A suggested that RSS measurements are very stable over time. For instance, recordings with indices 3, 4, and 5 were acquired at the exact same locations over multiple days but the mean RSS of recording number 4 is higher with 15 dB than the other two recordings. Such a large variability might be caused by the chipset warm-up after a reboot, interference in the ISM band, or other environmental factors such as the room temperature. Although we used the same model of devices for the measurements, the TAU data set from Fig. 7b had a smaller (but still significant) spread than the UPB data set from Fig. 7a, of 20 dB compared to 40 dB, respectively.

When aggregating data from multiple recordings, however, for at least 4 recordings the mean RSS converges to approximately -49 dBm and -55 dBm for UPB and TAU, respectively, as shown in Fig. 8. It is important to note that, although a relatively small number of recordings is necessary to capture the variability of the mean RSS between two devices across different locations, the shape of the distribution (and hence its spread) stabilizes only after 12–13 recordings.

C. The Impact of Hardware on the BLE RSS

We evaluated the impact of the hardware choice on the RSS when the same device model (Raspberry Pi 4 Model B) was used on both the transmitter and the receiver side. The devices were placed at the exact same location, with a fixed distance between them of 2 m, and the transmitter sent advertising beacons only on channel 37. We used in total four different Raspberry Pi boards, from exactly the same manufacturer and same model type, labeled RP11 to RP14 which integrate a Cypress CYW43455 BLE and Wi-Fi chipset.

Fig. 9 shows the RSS distribution of each pair of devices. Pair 1 consisted of the TX-RX pair RP11-RP12, pair 2 of RP11-RP13, pair 3 of RP11-RP14, pair 4 of RP12-RP14, and pair 5 of RP13-RP14. In other words, device




Fig. 7. Boxplots showing the RSS distribution in 15 recordings acquired between the same device pair in LOS, at 1 m, at UPB and TAU, using only channel 37.




Fig. 8. The impact of the number of snapshot recordings from a particular scenario on the RSS distribution. A total of 15 recordings were acquired in the same scenario (LOS at 1 m on channel 37). This figure presents the RSS distribution when an equal number of samples (326, corresponding to approximately 5 min) are taken from 1 to 15 recordings selected at random. The median, mean, and inter-quartile range (IQR) converge for more than 12–13 recordings.




Fig. 9. The impact of hardware choice on the RSS in recordings acquired with different device pairs placed at a distance of 2 m at exactly the same location. The median RSS varies with 5 dB, even though the devices have the same model.

pairs 1, 2, and 3 share the same transmitter, while device pairs 3, 4, and 5 share the same receiver. The median RSS varies with up to 5 dB even between devices from the same model. This experiment shows the difficulty of building a database that documents the transmitter and receiver efficiency of different brands of devices, since even devices that use the same hardware have RSS variations of several dB.

D. The Impact of the Advertising Channel Index on the BLE RSS

As mentioned earlier, BLE devices transmit beacons on channels 37, 38, and 39 which correspond to frequencies of

2.402, 2.426, and 2.48 GHz, respectively. Fig. 10 illustrates the impact of three advertising channels on the RSS, compared with a recording where all 3 advertising channels were used. The data was collected at the same location with the devices 2 m apart and on the same day within a short time interval. The same type of measurements were done in parallel at UPB and TAU. By default, beacons are transmitted on all three advertising channels. Therefore, a receiver cannot determine the channel of the transmitted packets, resulting in a larger variance of the samples and inaccurate distance estimates. At both UPB and TAU we noticed variations of at least 5 dB between measurements acquired on different channels. Other sources measured differences between BLE channels as high as 15 dB (Figure 2 in [54]).

E. The Impact of Transmitter-Receiver Distance on the BLE RSS

Under the LOS assumption (i.e., no obstacle between the a BLE transmitter and a receiver), one can start from the well-known free-space path-loss (FSL) model:

$$P_R = P_T - 20\log_{10}d - 20\log_{10}\left(\frac{4\pi f_c}{c}\right) + \eta, \quad (1)$$

where P_R is the received signal strength in dB scale, P_T is the apparent transmit power of the BLE transmitter computed at 1m away from the transmitter, d is the distance between the transmitter and the receiver (i.e., between the two persons under consideration in the digital contact-tracing app), f_c is the




Fig. 10. The channel impact on RSS values in a LOS scenario at 2 m distance, based on measurements acquired at UPB and TAU. This plot illustrates the RSS distributions (with an equal number of samples: 95 per snapshot) when receiving beacons on all three channels and on individual channels.

carrier frequency of the transmitted BLE signal (i.e., the carrier frequency on the used advertising channel or an average carrier frequency when several advertising channels are used), c is the speed of light (i.e., about $3 * 10^8$ m/s), and η is a noise factor encompassing the shadowing effects in the wireless channel, interference, and possible other noise sources. By virtue of the central-limit theorem, η can be assumed to be Gaussian distributed of variance σ^2 . We also assume that η is a zero-mean noise under LOS scenarios.

The FSL is rarely used as such in RSS modeling; instead, most authors prefer the one-slope path-loss model below for its simplicity [27], [55]–[57]:

$$P_R = P_{T_a} - 10n \log_{10} d + \eta, \quad (2)$$

where the apparent transmit power P_{T_a} factor includes also the frequency-dependent effects, in such a way that multi-frequency effects, as those generated by RSS measurements on multiple advertising channels can be lumped into a single parameter, and n is a positive number modeling the path-loss parameter. An n value below the FSL path-loss factor of 2 would signal the presence of some conductivity effects in the building walls as well as multipath-enhanced propagation (e.g., multipath adding constructively). The lower n is, the flatter the RSS curve is with the distance, and the harder would be to differentiate between close distances (e.g., between 1m and 2 m or between 2 m and 3 m). Typically, in model-driven RSS approaches (as opposed to data-driven approaches), the purpose is to estimate the best-fit parameters P_{T_a} and n of an underlying path-loss model. This is usually done via a least-square (LS) fit, where the unknown parameter vector $\mathbf{x} \triangleq [P_{T_a} \ n]$ is estimated via $\hat{\mathbf{x}}$ [55]:

$$\hat{\mathbf{x}} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}, \quad (3)$$

with $\mathbf{A} \in N_{meas} \times 2$ being a matrix with i -th row equal to $[1 - 10 \log_{10} d_i]$, $i = 1, \dots, N_{meas}$, and $\mathbf{b} \in N_{meas} \times 1$ being a vector with the i -th element equal to the received signal strength P_{R_i} observed in the i -th measurement at d_i distance between TX and RX. Above, N_{meas} is the number of measurements (or observations) used in the LS fitting, and encompassing various TX-RX distances d_i . The shadowing

variance $\hat{\sigma}_\eta^2$ is then computed as the error between the measurements and the reconstructed data, namely:

$$\hat{\sigma}_\eta^2 = \frac{1}{N_{meas}} \sum_{i=1}^{N_{meas}} \left(P_{R_i} - \hat{P}_{T_a} - 10\hat{n} \log_{10}(d_i) \right)^2. \quad (4)$$

When a NLOS obstacle such as a glass window, a wall, or the body of another person is present between the transmitter and receiver, we expect the NLOS apparent transmit power P_{T_a} to be smaller than the LOS P_{T_a} , as it should incorporate the additional absorption losses due to obstacles. However, repeated measurements are both TAU and UPB showed that this is not always the case.

Table II gives examples of the path-loss parameters estimated from aggregated measurements on all three BLE advertising channel, in four considered scenarios (two LOS and two NLOS, with two of them from TAU and two from UPB scenarios). In the NLOS scenarios, the obstruction was caused by a wall between the TX and the RX. In order to have a fair comparison also between long recordings, we extracted 326 samples from each available recording (which correspond to a recording time of around 5 min) and aggregated them. Several Monte Carlo runs showed very similar parameter-fit results from one run to another. For illustrative purposes, Table II shows the results based on one random run in each scenario.

The main conclusion is that there is not a one-size-fit-all model with constant $[\hat{P}_{T_a} \ \hat{n}]$ vector estimate, but that there are high fluctuations between the four shown scenarios, and therefore a model-driven approach for BLE RSS-based contact tracing will likely suffer from large errors. This is also reflected in the high shadowing standard deviations $\hat{\sigma}_\eta$ shown in Table II (around 6 dB for TAU data and around 9 dB for UPB data).

Indeed, other literature results have shown that the path-loss parameters used in different works vary widely. For instance, although a path-loss exponent between 2.4–2.6 is frequently recommended [26], in [43] the path-loss exponent was set to 1.8 for LOS scenarios and 2.2 for NLOS ones. In the survey part of [2], the path-loss exponents extracted from various research papers varied between 0.63–2.32 and TABLE II suggests that in some cases (NLOS, UPB) this value might be even lower.

Instead of model-driven contact tracing, data-driven approaches such as those based on large training data sets and machine learning solutions (e.g., in [26]) could be adopted, but they have high complexity and are impractical at large scales. Other works [58], [59] propose online path-loss estimation methods based on cooperating nodes in wireless sensor networks. However, tens of nodes are usually needed for an accurate estimation. Another solution is to have a gateway that collects the RSS of surrounding BLE beacons, tracks the fluctuations, and sends back RSS correction factors to individual nodes in real time [24]. However, such an approach is not suitable for a peer-to-peer and privacy-sensitive application like contact-tracing. Therefore, the challenges of finding the right approach (model-driven versus data-driven) and the right models (e.g., more

TABLE II
EXAMPLE OF PATH-LOSS PARAMETERS ESTIMATED FROM AGGREGATED MEASUREMENTS.

Environment	\hat{P}_{T_s} [dBm]	\hat{n} [-]	$\hat{\sigma}_\eta$ [dB]	Total number of measurements		
				at 1 m	at 2 m	at 3 m
LOS, TAU	-55.35	0.76	6.01	5868	2608	1630
NLOS, TAU	-49.98	3.31	5.70	3586	2934	1956
LOS, UPB	-44.07	1.60	8.73	5216	2608	2934
NLOS, UPB	-51.69	0.37	8.72	1630	1630	1630

sophisticated models than the simple single-slope path-loss model of Eq. (2)) are still important challenges to be solved by the research community dealing with BLE RSS-based contact tracing or proximity detection.

F. The Impact of Transmitter and Receiver Orientation on the BLE RSS

We considered the effect of the relative orientation between the transmitter and the receiver on the BLE RSS. We analyzed four poses depicted in Fig. 12, where the pose of the transmitter is fixed and the receiver is rotated clock-wise with 90° , 180° , and 270° with respect to the “front” orientation from Fig. 12a, resulting in the “left,” “right,” and “back” poses, respectively. The radiation pattern (Fig. 11) for the frequency of Bluetooth channel 37 shows a 2.7 dB standard deviation across all angles, but the maximum differences on each of the three planes is of 10.1, 13, and 14.1 dB.

Fig. 13 presents the RSS distribution in all poses, when the devices are placed at distances of 1, 2, and 3 m. The devices were placed on tripods which were kept fixed at the aforementioned distances, while only the receiver was rotated around its center axis for each pose. Each recording had a duration of approximately 10 min and was performed only on channel 37. First, we notice the same inconsistencies with the distance discussed in Section IV-E, in which the average RSS at 1 m distance is lower than the one at 2 and 3 m. Second, the RSS changes with the pose for a particular distance, although the receiver was not moved but only rotated around its axis and the transmitter’s position was the same in all recordings. There is no orientation which results in a higher RSS at all distances. However, the “back” pose has a lower median RSS than the other poses at all distances, most likely because in this pose, as can be seen from Fig. 12b, the metallic USB and Ethernet ports of the receiver board are in the LOS of the signal and attenuate it. While the median RSS in the “front,” “left,” and “right” poses varies with about 5 dB for the same distance, the median RSS in the “back” pose can be with even 20 dB lower than in the other poses.

G. Random Fluctuations Caused by BLE–Wi-Fi Combo Chipsets

The interference between Bluetooth and Wi-Fi is well documented in literature, and IEEE has recommendations




Fig. 11. Radiation pattern of Raspberry Pi 3B+ antenna plotted from anechoic chamber measurement data [60]. It is a PCB antenna designed by Proant AB present in many IoT devices operating in the 2.4GHz band.




Fig. 12. The four orientations of the receiver (device on the right) with respect to the transmitter (device on the left) we considered in our experiment: (a) front, (b) back, (c) left, and (d) right. In the back, left, and right poses the receiver was rotated clock-wise with, respectively, 180° , 90° , and 270° with respect to the front orientation.

[61] for the coexistence of the technologies operating in the ISM bands. [62] has shown experimentally that Bluetooth and ZigBee are affected by Wi-Fi, stating agreement with previous studies. One way to tackle the coexistence is to use specific algorithmic mitigations in the way each technology is used [63], but some might require updates to the standards. When both Wi-Fi and Bluetooth are implemented on the same chipset, as is the case with most smartphones, [64] determined through measurements that performance is degraded at the application layer.

For the purpose of contact tracing however, only RSS measurements and timestamps of the recordings are needed and the question is whether the BLE measurements are influenced by Wi-Fi activity on the same chipset. We turn the Wi-Fi on and off simultaneously at both the transmitter and the receiver every one hour and record the RSS. During the time when the Wi-Fi is on, synthetic Wi-Fi traffic is generated with 112 kb/s. Fig. 14 shows that, on average, the mean RSS when the Wi-Fi is on is 2.5 dB lower than during the time the Wi-Fi is off. There is also a small difference in the standard deviation: when the Wi-Fi is on, the standard deviation of the RSS is 1.1 dB compared to 0.83 dB when the Wi-Fi is off. Although Fig. 14 presents the results for 6 hours only, for better visualization, the pattern remained consistent over two days, during which there was no human activity around the devices.

Wi-Fi scanning might occasionally cause even larger differences in the BLE RSS than 2.5 dB. We sometimes




Fig. 13. The impact of orientation on the BLE RSS. The front, left, right, and back orientations are shown in Fig. 12.




Fig. 14. The impact of Wi-Fi and BLE combo chipsets: the Wi-Fi is turned on and off every one hour at the indicated markers. On average, the mean RSS with the Wi-Fi off is with 2.5 dB higher than with the Wi-Fi off.

noticed spurious measurements occurring only when the Wi-Fi was on, usually at the beginning of a snapshot recording, as shown in Fig. 15 around minute 100, when the signal fluctuated for several minutes between -40 , -55 , and -90 dB. The recordings with the settings of Wi-Fi on and off were acquired during different times of the day; however, the environment was static with no people moving inside the room. We acquired results which show the mean RSS values to be with 6.8 dB higher with Wi-Fi off than with Wi-Fi switched on. Overall, Fig. 15 illustrates the instabilities in single recordings which might be caused by coexistence of different signals within the 2.4 GHz frequency. A similar pattern was observed also in [65], where Figures 6 and 7 reveal a 20 dB difference in BLE readings when Wi-Fi scanning is active with a Samsung Galaxy S4 smartphone.

H. On the Difficulty of LOS/NLOS detection

In this section, we investigate the effect of different types of obstructions on the BLE RSS. At both UPB and TAU we acquired measurements in LOS and NLOS with wall shadowing. In addition, at UPB, we tried more types of obstructions: wall and whiteboard, door, human body, and plexiglass panel. All the setups are shown in Fig. 4, where for NLOS measurements we varied only the distance between the devices, while in LOS we also tried different locations.

We consider a “scenario” a set of snapshot measurements acquired at the same distance, on the same channel, in the same LOS/NLOS setting. TABLE III presents the mean and the




Fig. 15. The impact of Wi-Fi switched on and off, in LOS at 1m distance. When the Wi-Fi is off, the mean RSS is higher with 6.8 dB than when the Wi-Fi is on. When the Wi-Fi is on, we also notice RSS fluctuations of up to 30 dB around minute 100.

standard deviation of the RSS computed in different scenarios from measurements acquired at UPB. Because recordings in the same scenarios have lengths from 5 min to several days and we do not want longer recordings to bias the statistics, when there are multiple recordings in the same scenario we chose an equal number of measurements at random from each recording from that scenario and computed the mean and standard deviation using only the subset of samples. Usually, recordings within the same scenario were acquired at different locations to capture the variability of the RSS across space for the same TX–RX distance. As a result, when there are multiple recordings in a scenario, the standard deviation of the RSS is higher than in single-recording scenarios. The number of recordings (“Nr. rec.”) and the number of samples in each recording (“Nr. samples per rec.”) are specified in TABLE III for each scenario, as well as the advertising channel(s) on which measurements were acquired.

Comparing the statistics in LOS and NLOS when all advertising channels are used, we note that the mean RSS varies within a range of 20 dB, indicating that the attenuation introduced by an obstacle depends on the type of obstacle. The mean RSS in LOS is usually higher than the one in NLOS but not always—the mean RSS in the “NLOS door” recording is higher than the one in LOS. Note also that the human body causes a higher standard deviation than the other obstructions. This can be seen more clearly in Fig. 16, which shows the distributions of selected NLOS measurements from TABLE III and one recording acquired in LOS on all channels. The large spread can be caused by slight movements of the body which, by nature, cannot be perfectly immobile (breathing alone causes a slight movement of the body). These characteristics can make the human body more easily detectable than other obstructions, as previous works showed [26]. However, other obstructions might be more difficult to detect. For instance, the mean RSS in the “NLOS plexiglass” case is similar to the mean RSS in LOS on individual channels, while the highest mean RSS was obtained in the “NLOS door” case.

The inconsistency can be also caused by the fact that most of the NLOS statistics were computed based on a single recording and, as we saw in Section IV-B, single recordings can deviate

from statistics computed on aggregated data with more than 10 dB. Therefore, next we compare LOS and NLOS with a wall distributions aggregated from all channels, at distances of 1, 2, and 3 m, acquired independently at UPB and TAU, shown in Fig. 17. The distributions are plotted based on the same data used in Section IV-E to estimate the path-loss parameters from TABLE II. Each distribution was computed based on 5 to 18 recordings based on 326 measurements selected at random from each recording. Based on the results in Section IV-B, the mean computed based on 5 recordings should be within several dB of the “stable” mean, but the standard deviation can still fluctuate for less than 12–13 recordings.

Although the distributions in Fig. 17 mostly behave as expected, i.e. the mean RSS should decrease with the distance and the mean RSS should be lower in NLOS than in LOS at the same distance, there are exceptions. The average RSS in NLOS is higher than the one in LOS at 3 m for the UPB data set and at 1 m for the TAU data set. Also, the NLOS distributions have higher or equal spread than LOS ones in most of the cases, even though the NLOS data sets contained less recordings than the LOS ones. This result contradicts observations in [43], where NLOS obstructions caused by walls were identified when the standard deviation of RSS measurements in a window was *lower* than a fixed threshold. Although UPB and TAU data sets were acquired using the same model of Raspberry Pis, measurements acquired at TAU had a smaller spread than those from UPB even in LOS, which points once more to the instability of RSS measurements.

In proximity-detection or RSS-based localization applications, obstructions will most of the time lead to inaccurate distance or location estimates. Therefore, multiple solutions have been proposed to correct RSS-based ranges by detecting the NLOS condition [26], [42], [43] with the caveat that such solutions might not generalize easily, as our measurements show, or that large data sets might be necessary to extract features that improve classification. In contact-tracing applications, such instabilities can lead to false alarms or failures in detecting potentially unsafe interactions. For instance, since human body shadowing sharply attenuates the signal, the distance predicted by a standard path-loss model can be larger than in reality, so people might not be notified of risky encounters. On the contrary, if the RSS reported when devices (or people) are separated by walls is larger or equal than the average RSS in LOS, an alert might be raised even if people staying in different rooms are safe from each other. Therefore, LOS/NLOS detection is still a highly relevant topic with room for improvement. Hybrid solutions that combine BLE with UWB, cameras, or other sensors might increase the reliability of NLOS detection.

V. DISCUSSION

One of the unexpected results of our measurement campaigns—a result which has also not been emphasized enough until now in the current literature—is the fact that snapshot BLE RSS measurements are highly unstable and fluctuating, and only by lumping together enough measurements (i.e., by using aggregated data), the results

TABLE III
THE MEAN AND STANDARD DEVIATION OF THE BLE RSS IN DIFFERENT SCENARIOS AT A TX–RX DISTANCE OF 2 m.

Scenario	Ch.	Nr. rec.	Nr. samples per rec.	RSS	
				Mean [dBm]	Std. [dB]
LOS	37	7	326	-50.5	7.6
LOS	38	1	3215	-46.0	5.6
LOS	39	1	25 607	-54.2	1.7
LOS	all	5	736	-45.6	7.3
NLOS wall	37	3	1726	-55.4	6.0
NLOS wall	all	1	1824	-58.2	2.0
NLOS human	all	1	495	-60.0	4.3
NLOS plexiglass	all	1	1824	-50.7	2.9




Fig. 16. Comparison of selected RSS distributions acquired on all advertising channels at 2 m, in LOS, NLOS with a plexiglass panel, NLOS with a wall and a whiteboard, and NLOS with human body shadowing.




Fig. 17. Comparison of RSS distributions based on data acquired in LOS and NLOS with a wall at UPB and TAU at distances of 1, 2, and 3 m, irrespective of the advertisement channel.

seem to converge, to some extent, to the classical path-loss models (e.g., average RSS decreasing with transmitter-receiver distances, average RSS under LOS scenarios stronger than the average RSS under NLOS scenarios). Nevertheless, for fast proximity-detection or contact-tracing algorithms, when the observation window can be as small as 15 min, aggregated RSS data may be unavailable, and estimations based on what we called snapshot recordings can suffer from significant errors due to high RSS fluctuations. We also provided guidelines for building data sets that best represent the conditions in a particular scenario.

VI. CONCLUSIONS, AND OPEN ISSUES

This paper presented a detailed analysis of BLE RSS fluctuations based on an extensive measurement campaign performed in tandem in Tampere, Finland, and Bucharest, Romania. We documented in detail the main sources of high fluctuations (or instabilities) of BLE RSS measurements occurring, surprisingly, in static scenarios and diverging from the classical path-loss models, e.g., as given in Eqs. (1) and (2). We defined controlled scenarios, such as fixing the transmitter and receiver BLE models, fixing the BLE advertising channel to have transmission on a single carrier frequency, turning the Wi-Fi transmitter off in chipsets sharing the 2.4 GHz antenna between BLE and Wi-Fi, and fixing the transmitter-receiver distance.

We emphasized several challenges that still remain to be addressed by the research community when standalone BLE RSS measurements are used for contact tracing, proximity detection, or positioning purposes, namely: the challenges of NLOS scenarios with stronger average (and median) RSS than LOS scenarios at the same distance, the challenge of increased RSS fluctuations (or variance) when the measurements are acquired on multiple BLE advertising channels (as it is customary in contact-tracing applications) or with different receiver-transmitter orientations (which again are highly variable, as users can keep their mobile devices in various positions: in hand, inside bags, inside front or back pockets, etc.).

A possible solution to overcome the instability of snapshot BLE RSS recordings is, for example, the hybridization of BLE RSS measurements with other sensors, such as vision sensors (to enable LOS/NLOS detection) or time-of-arrival UWB sensors (to enhance the range estimation). However, this will increase the energy consumption of end-user devices, so the trade-off between proximity detection accuracy and energy consumption must also be considered. Collecting data from additional sensors can also potentially decrease user privacy. Another possibility would be to collect large training databases in hotspot areas (e.g., shopping centers, commuting halls, etc.), which could facilitate a baseline statistical modeling based on both snapshot and aggregated training data, and to further use machine-learning approaches to derive data-driven estimators instead of the model-driven estimators which rely on path-loss modeling.

The main goal of this paper is to shed additional light on the challenges encountered in BLE-based contact tracing and to raise awareness among the research community that several challenges related to BLE RSS ranging and positioning are still to be solved. One solution based on our measurements is to use enough aggregated data, as, by virtue of the central-limit theorem, this seems to remove the outliers and to converge towards known path-loss models. Such a solution could be sufficient for positioning purposes when training databases can be based on large amounts of aggregated data, but it may still be unfeasible for contact-tracing solutions in need working with snapshot data. Another solution could envisage more sophisticated path-loss modeling, such as by taking waveguide effects [66] into account or using stochastic ray-

tracing modeling [67].

The measurement data will be made open-access at the research community on Zenodo², in order to enable the reproducibility of the research and to provide benchmark data for further investigations on BLE RSS-based contact tracing. Future work also includes collecting data from more devices, including various types of mobile phones, and looking into more detail at the yet-unsolved research question of whether NLOS situations can be separated with high accuracy from LOS situations and under which conditions.

ACKNOWLEDGMENT

The authors gratefully acknowledge funding from European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints, <http://www.a-wear.eu/>); work was also supported in part by a grant from the Romanian National Authority for Scientific Research and Innovation, UEFISCDI project PN-III-P2-2.1-PED-2019-5413.

REFERENCES

- [1] V. Shubina, S. Holcer, M. Gould, and E. S. Lohan, "Survey of Decentralized Solutions with Mobile Devices for User Location Tracking, Proximity Detection, and Contact Tracing in the COVID-19 Era," *Data*, vol. 5, no. 4, p. 87, 2020.
- [2] V. Shubina, A. Ometov, A. Basiri, and E. S. Lohan, "Effectiveness modelling of digital contact-tracing solutions for tackling the COVID-19 pandemic," *Journal of Navigation*, p. 1–37, 2021.
- [3] C. Nguyen and et al., "A Comprehensive Survey of Enabling and Emerging Technologies for Social Distancing—Part I: Fundamentals and Enabling Technologies," *IEEE Access*, vol. 8, pp. 153479–153507, 2020.
- [4] C. Nguyen and et al., "A Comprehensive Survey of Enabling and Emerging Technologies for Social Distancing—Part II: Emerging Technologies and Open Issues," *IEEE Access*, vol. 8, pp. 154209–154236, 2020.
- [5] E. Hernández-Orallo, P. Manzoni, C. T. Calafate, and J.-C. Cano, "Evaluating How Smartphone Contact Tracing Technology Can Reduce the Spread of Infectious Diseases: The Case of COVID-19," *IEEE Access*, vol. 8, pp. 99083–99097, 2020.
- [6] I. Braithwaite, T. Callender, M. Bullock, and R. W. Aldridge, "Automated and partly automated contact tracing: a systematic review to inform the control of COVID-19," *The Lancet Global Health*, 8 2020.
- [7] N. Ahmed, R. A. Michelin, W. Xue, S. Ruj, R. Malaney, S. S. Kanhere, A. Seneviratne, W. Hu, H. Janicke, and S. K. Jha, "A Survey of COVID-19 Contact Tracing Apps," *IEEE Access*, vol. 8, pp. 134577–134601, 2020.
- [8] 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*, vol. 24, pp. 1–12, 2017.
- [9] S. Tomic, M. Beko, R. Dinis, M. Tuba, and N. Bacanin, "Bayesian methodology for target tracking using combined RSS and AoA measurements," *Physical Communication*, vol. 25, pp. 158–166, 2017.
- [10] F. Della Rosa, T. Paakki, J. Nurmi, and M. Pelosi, "Exploiting RSS measurements among neighbouring devices: A matter of trust," in *International Conference on Indoor Positioning and Indoor Navigation*, pp. 1–8, IEEE, 2013.
- [11] C. Gentner, D. Günther, and P. H. Kindt, "Identifying the ble advertising channel for reliable distance estimation on smartphones," 2020.
- [12] A. Kara and H. L. Bertoni, "Blockage/shadowing and polarization measurements at 2.45 ghz for interference evaluation between bluetooth and ieee 802.11 wlan," in *IEEE Antennas and Propagation Society International Symposium. 2001 Digest. Held in conjunction with: USNC/URSI National Radio Science Meeting (Cat. No.01CH37229)*, vol. 3, pp. 376–379 vol.3, 2001.

²<https://doi.org/10.5281/zenodo.4643668>

- [13] U. M. Qureshi, Z. Umair, Y. Duan, and G. P. Hancke, "Analysis of bluetooth low energy (ble) based indoor localization system with multiple transmission power levels," in *2018 IEEE 27th International Symposium on Industrial Electronics (ISIE)*, pp. 1302–1307, 2018.
- [14] Y. Li, Z. He, Y. Li, Z. Gao, R. Chen, and N. El-Sheimy, "Enhanced Wireless Localization Based on Orientation-Compensation Model and Differential Received Signal Strength," *IEEE Sensors Journal*, vol. 19, no. 11, pp. 4201–4210, 2019.
- [15] F. Della Rosa, T. Paakki, J. Nurmi, M. Pelosi, and G. D. Rosa, "Hand-grip impact on range-based cooperative positioning," in *2014 11th International Symposium on Wireless Communications Systems (ISWCS)*, pp. 728–732, 2014.
- [16] P. Handayani, L. Mubarakah, and G. Hendrantoro, "Pathloss and shadowing characteristics in indoor environment at 2.4 ghz band," in *2015 International Seminar on Intelligent Technology and Its Applications (ISITIA)*, pp. 423–428, 2015.
- [17] C. Xiang, P. Yang, C. Tian, L. Zhang, H. Lin, F. Xiao, M. Zhang, and Y. Liu, "Carm: Crowd-sensing accurate outdoor rssi maps with error-prone smartphone measurements," *IEEE Transactions on Mobile Computing*, vol. 15, no. 11, pp. 2669–2681, 2016.
- [18] Q. Wang, D. W. Matolak, and B. Ai, "Shadowing characterization for 5-ghz vehicle-to-vehicle channels," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 3, pp. 1855–1866, 2018.
- [19] M. Sasaki, T. Nakahira, K. Wakao, and T. Moriyama, "Human blockage loss characteristics of 5 ghz wi-fi band in a crowded stadium," *IEEE Antennas and Wireless Propagation Letters*, pp. 1–1, 2021.
- [20] Massachusetts Institute of Technology (MIT), "PACT Datasets and Evaluation." <https://github.com/mitll/BLE-RSSI-Variou-Static-Configurations>, Accessed: 24 June 2021.
- [21] Chang Li and Lu Shi, "Exposure Notification (Contact Tracing) System on Raspberry Pi." <https://github.com/ecechl/Exposure-Notification-on-RPi>, Accessed: 24 June 2021.
- [22] Lukas Trommer and Niclas Kühnapfel, "Test of Covid-19 Bluetooth LE localization precision." <https://github.com/niclasku/ble-distance-measurements>, Accessed: 24 June 2021.
- [23] "A technical roadmap for the UK's contact tracing app functionality." <https://www.turing.ac.uk/blog/technical-roadmap-uks-contact-tracing-app-functionality>, Accessed: 31 March 2021.
- [24] A. Nikoukar, M. Abboud, B. Samadi, M. Güneş, and B. Dezfouli, "Empirical analysis and modeling of Bluetooth low-energy (BLE) advertisement channels," in *2018 17th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, pp. 1–6, IEEE, 2018.
- [25] R. Faragher and R. Harle, "Location fingerprinting with bluetooth low energy beacons," *IEEE journal on Selected Areas in Communications*, vol. 33, no. 11, pp. 2418–2428, 2015.
- [26] S. Naghdi and K. O'Keefe, "Detecting and correcting for human obstacles in ble trilateration using artificial intelligence," *Sensors*, vol. 20, no. 5, p. 1350, 2020.
- [27] E. S. Lohan, J. Talvitie, P. Figueiredo e Silva, H. Nurminen, S. Ali-Löytty, and R. Piché, "Received signal strength models for wlan and ble-based indoor positioning in multi-floor buildings," in *2015 International Conference on Localization and GNSS (ICL-GNSS)*, pp. 1–6, 2015.
- [28] J. Neburka, Z. Tlamsa, V. Benes, L. Polak, O. Kaller, L. Bolecek, J. Sebesta, and T. Kratochvil, "Study of the performance of rssi based bluetooth smart indoor positioning," in *2016 26th International Conference Radioelektronika (RADIOELEKTRONIKA)*, pp. 121–125, 2016.
- [29] G. De Blasio, A. Quesada-Arencibia, C. R. García, J. C. Rodríguez-Rodríguez, and R. Moreno-Díaz, "A protocol-channel-based indoor positioning performance study for bluetooth low energy," *IEEE Access*, vol. 6, pp. 33440–33450, 2018.
- [30] D. Sikeridis, I. Papapanagiotou, and M. Devetsikiotis, "BLEBeacon: A real-subject trial dataset from mobile Bluetooth low energy beacons," *arXiv preprint arXiv:1802.08782*, 2018.
- [31] N. Mohsin, S. Payandeh, D. Ho, and J. P. Gelinas, "Study of activity tracking through bluetooth low energy-based network," *Journal of Sensors*, vol. 2019, p. 6876925, Feb 2019.
- [32] G. M. Mendoza-Silva, M. Matey-Sanz, J. Torres-Sospedra, and J. Huerta, "Ble rssi measurements dataset for research on accurate indoor positioning," *Data*, vol. 4, no. 1, 2019.
- [33] P. C. Ng, P. Spachos, and K. Plataniotis, "Covid-19 and your smartphone: Ble-based smart contact tracing," *arXiv preprint arXiv:2005.13754*, 2020.
- [34] D. J. Leith and S. Farrell, "Coronavirus contact tracing: Evaluating the potential of using bluetooth received signal strength for proximity detection," 2020.
- [35] S. Sadowski, P. Spachos, and K. N. Plataniotis, "Memoryless Techniques and Wireless Technologies for Indoor Localization with the Internet of Things," *IEEE Internet of Things Journal*, vol. 7, no. 11, pp. 10996–11005, 2020.
- [36] P. C. Ng, P. Spachos, and K. N. Plataniotis, "COVID-19 and your smartphone: BLE-based smart contact tracing," *IEEE Systems Journal*, 2021.
- [37] A. F. Molisch, *Wireless communications*, vol. 34. John Wiley & Sons, 2012.
- [38] J. Paek, J. Ko, and H. Shin, "A measurement study of ble beacon and geometric adjustment scheme for indoor location-based mobile applications," *Mobile Information Systems*, vol. 2016, 2016.
- [39] M. Castillo-Cara, J. Lovón-Melgarejo, G. Bravo-Rocca, L. Orozco-Barbosa, and I. García-Varea, "An empirical study of the transmission power setting for bluetooth-based indoor localization mechanisms," *Sensors*, vol. 17, no. 6, p. 1318, 2017.
- [40] "Exposure notifications BLE calibration calculation." <https://developers.google.com/android/exposure-notifications/ble-attenuation-computation#rssi-tx-power>, Accessed: 31 March 2021.
- [41] "Android fragmentation visualized." https://www.opensignal.com/sites/opensignal-com/files/data/reports/global/data-2015-08/2015_08_fragmentation_report.pdf, Accessed: 31 March 2021.
- [42] G. De Blasio, A. Quesada-Arencibia, C. R. García, J. M. Molina-Gil, and C. Caballero-Gil, "Study on an indoor positioning system for harsh environments based on wi-fi and bluetooth low energy," *Sensors*, vol. 17, no. 6, p. 1299, 2017.
- [43] A. A. Juri, T. Arslan, and F. Wang, "Obstruction-aware bluetooth low energy indoor positioning," in *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, pp. 2254–2261, 2016.
- [44] M. Youssef, M. Mah, and A. Agrawala, "Challenges: device-free passive localization for wireless environments," in *Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, pp. 222–229, 2007.
- [45] K. Woyach, D. Puccinelli, and M. Haenggi, "Sensorless sensing in wireless networks: Implementation and measurements," in *2006 4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, pp. 1–8, IEEE, 2006.
- [46] N. Fet, M. Handte, and P. J. Marrón, "A model for wlan signal attenuation of the human body," in *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pp. 499–508, 2013.
- [47] S. Galmiche, T. Charmet, L. Schaeffer, J. Paireau, R. Grant, O. Chény, C. V. Platen, A. Maurizot, C. Blanc, A. Dinis, S. Martin, F. Omar, C. David, A. Septons, S. Cauchemez, F. Carrat, A. Mailles, D. Levy-Bruhl, and A. Fontanet, "Exposures associated with SARS-CoV-2 infection in france: A nationwide online case-control study," *The Lancet Regional Health - Europe*, vol. 7, p. 100148, aug 2021.
- [48] T. C. Bulfone, M. Malekinejad, G. W. Rutherford, and N. Razani, "Outdoor transmission of SARS-CoV-2 and other respiratory viruses: A systematic review," *The Journal of Infectious Diseases*, vol. 223, pp. 550–561, nov 2020.
- [49] D. Niculescu, "Interference map for 802.11 networks," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pp. 339–350, 2007.
- [50] Y. Chapre, P. Mohapatra, S. Jha, and A. Seneviratne, "Received signal strength indicator and its analysis in a typical wlan system (short paper)," in *38th Annual IEEE Conference on Local Computer Networks*, pp. 304–307, IEEE, 2013.
- [51] D. Lymberopoulos, Q. Lindsey, and A. Savvides, "An empirical characterization of radio signal strength variability in 3-d ieee 802.15. 4 networks using monopole antennas," in *European Workshop on Wireless Sensor Networks*, pp. 326–341, Springer, 2006.
- [52] "Exposure notification bluetooth specification, version 1.2." <https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ExposureNotification-BluetoothSpecificationv1.2.pdf>, Accessed: 31 March 2021.
- [53] D. J. Leith and S. Farrell, "Gae'n due diligence: Verifying the google/apple covid exposure notification api," *CoronaDef21, Proceedings of NDSS '21*, vol. 2021, 2020.
- [54] J. Powar, C. Gao, and R. Harle, "Assessing the impact of multi-channel ble beacons on fingerprint-based positioning," in *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, IEEE, 2017.
- [55] S. Shrestha, J. Talvitie, and E. S. Lohan, "Deconvolution-based indoor localization with wlan signals and unknown access point locations," in

- 2013 *International Conference on Localization and GNSS (ICL-GNSS)*, pp. 1–6, 2013.
- [56] A. Zanella, “Best practice in rssi measurements and ranging,” *IEEE Communications Surveys Tutorials*, vol. 18, no. 4, pp. 2662–2686, 2016.
- [57] Z. He, Y. Li, L. Pei, R. Chen, and N. El-Sheimy, “Calibrating multi-channel RSS observations for localization using gaussian process,” *IEEE Wireless Communications Letters*, vol. 8, no. 4, pp. 1116–1119, 2019.
- [58] A. Bel, J. L. Vicario, and G. Seco-Granados, “Localization algorithm with on-line path loss estimation and node selection,” *Sensors*, vol. 11, no. 7, pp. 6905–6925, 2011.
- [59] G. Mao, B. D. Anderson, and B. Fidan, “Path loss exponent estimation for wireless sensor network localization,” *Computer Networks*, vol. 51, no. 10, pp. 2467–2483, 2007.
- [60] Antenna Test Lab Company, “Raspberry Pi Model 3B+ Antenna Evaluation.” <https://antennatestlab.com/antenna-examples/raspberry-pi-model-3b-antenna-evaluation-gain-pattern>. Accessed: 31 March 2021.
- [61] L. M. S. Committee *et al.*, “Coexistence of wireless personal area networks with other wireless devices operating in unlicensed frequency bands,” *IEEE Computer Society*, vol. 425, p. 429, 2003.
- [62] R. G. Garroppo, L. Gazzarrini, S. Giordano, and L. Tavanti, “Experimental assessment of the coexistence of wi-fi, zigbee, and bluetooth devices,” in *2011 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, pp. 1–9, IEEE, 2011.
- [63] N. Golmie, N. Chevrollier, and O. Rebala, “Bluetooth and wlan coexistence: challenges and solutions,” *IEEE Wireless Communications*, vol. 10, no. 6, pp. 22–29, 2003.
- [64] R. Friedman, A. Kogan, and Y. Krivolapov, “On power and throughput tradeoffs of wifi and bluetooth in smartphones,” *IEEE Transactions on Mobile Computing*, vol. 12, no. 7, pp. 1363–1376, 2012.
- [65] R. Faragher and R. Harle, “An analysis of the accuracy of bluetooth low energy for indoor positioning applications,” in *Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2014)*, pp. 201–210, 2014.
- [66] K. Park, J. Lee, S. Hyun, and S. Kim, “Analysis of path loss properties in indoor hallway with waveguide channel model,” in *2019 IEEE VTS Asia Pacific Wireless Communications Symposium (APWCS)*, pp. 1–5, 2019.
- [67] T. Zwick, C. Fischer, D. Didascalou, and W. Wiesbeck, “A stochastic spatial channel model based on wave-propagation modeling,” *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 1, pp. 6–15, 2000.

PUBLICATION 6

FlexTDOA: Robust and Scalable Time-Difference of Arrival Localization Using Ultra-Wideband Devices

George-Cristian Pătru, Laura Flueratoru, Iuliu Vasilescu, Dragoş Niculescu,
and Daniel Rosner

IEEE Access 11 (2023), pp. 28610–28627

DOI: 10.1109/ACCESS.2023.3259320

Publication reprinted with the permission of the copyright holders.

FlexTDOA: Robust and Scalable Time-Difference of Arrival Localization Using Ultra-Wideband Devices

GEORGE-CRISTIAN PĂTRU¹, LAURA FLUERATORU^{1,2}, IULIU VASILESCU³, DRAGOȘ NICULESCU¹ AND DANIEL ROSNER¹

¹Department of Computer Science, University Politehnica of Bucharest, 060042 Bucharest, Romania

²Electrical Engineering Unit, Tampere University, 33720 Tampere, Finland

³TIA RESEARCH, 820169 Tulcea, Romania

Corresponding author: George-Cristian Pătru (e-mail: cristi.patru@upb.ro)

The authors would like to acknowledge the support of NXP Romania towards supporting the PhD studies of George-Cristian Pătru, as well as providing valuable insights regarding trends for IoT solutions.

The authors gratefully acknowledge funding from European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska Curie grant agreement No. 813278 (A-WEAR: A network for dynamic wearable applications with privacy constraints).

The results presented in this article were obtained with the support of the Ministry of European Investments and Projects through the Human Capital Sectoral Operational Program 2014-2020, Contract no. 62461/03.06.2022, SMIS code 153735.

ABSTRACT In this paper, we propose FlexTDOA, an indoor localization method using ultra-wideband (UWB) radios, and we demonstrate its performance in a functional system. Our method uses time-difference of arrival (TDOA) localization so that the user device remains passive and is able to compute its location simply by listening to the communication between the fixed anchors, ensuring the scalability of the system. The anchors communicate using a custom and flexible time-division multiple-access (TDMA) scheme in which time is divided in slots. In each time slot, one anchor interrogates one or more anchors which respond in the same slot. The anchors do not need to have their clocks synchronized. We implemented FlexTDOA on in-house designed hardware using a commercial UWB module. We evaluate the localization accuracy of FlexTDOA with different system parameters such as the number of responses, the order of responses, and the number of anchors. We simulate and evaluate the effect of the physical speed of the tag on the choice of optimum system parameters. We also compare FlexTDOA against the classic TDOA approach and range-based localization in a deployment of ten anchors and one tag, both with and without obstructions. Results show that FlexTDOA achieves the highest localization accuracy in most of the scenarios, with up to 38% reduction in the localization error compared to the classic approach.

INDEX TERMS ultra-wideband, indoor localization, time-difference of arrival, Internet of Things (IoT)

I. INTRODUCTION

Ultra-wideband (UWB) technology has experienced a revival during the past few years, mainly for its high-accuracy ranging and localization capabilities. It is estimated that more than 1 billion UWB devices will be shipped by 2025, and that over the next 5–10 years all smartphones will have UWB capabilities [1]. With the growing number of users, localization systems will face high scalability requirements to satisfy network demands with acceptable location update rates [2].

UWB-based localization systems usually consist of a mobile node that needs to be localized, called *tag*, and several fixed nodes with known locations, called *anchors*, which

communicate with the tag and aid the localization process. Range-based localization is arguably the most popular localization technique since it provides the highest accuracy [6]. In range-based localization, the location is computed based on distances between each anchor and the tag using multilateration. To avoid synchronizing the transmitter (TX) and the receiver (RX) [3], [4], at least two message exchanges between the tag and each anchor are needed to compute *one* distance, technique known as TWR. Because of the high number of messages, range-based localization (or multilateration) scales poorly with the growing number of tags and anchors.

Time-difference of arrival (TDOA) is an alternative localization method which uses the difference between the arrival

times of two packets (usually, exchanged by the tag and two anchors) [5]. For each difference, the tag can be located on a hyperbola focused at the anchors (in a 2D coordinate system). By computing the TDOA for more anchor pairs, the tag's location can be found at the intersection of multiple hyperbolas [6]. One TDOA variant frequently called downlink (DL) TDOA [7] has gained popularity for its scalability properties. In this setup, only the anchors transmit periodic messages, while the tag records their arrival times and localizes itself. In DL TDOA, the tag can remain *passive*, i.e., does not need to transmit any uplink messages. Using passive tags, since there is no need for the tags to share the channel, DL TDOA can scale to an unlimited number of users [8]–[11].

The main drawback of TDOA localization is that the anchors need to be synchronized, usually by estimating the clock offsets of each anchor relative to a reference clock [11]. A convenient way to obtain the reference clock is to designate a reference anchor that periodically broadcasts a synchronization beacon [12]. The rest of the anchors respond to this beacon and track their clock offsets with respect to the reference anchor. The disadvantage of this approach is that, if the link between the reference anchor and the tag is obstructed, the tag will receive the synchronization beacon with a delay and *all* subsequent TDOAs pertaining to that beacon will be corrupted. This method is therefore not reliable in the presence of obstructions.

In this paper, we propose an alternative TDMA scheduling scheme for TDOA localization called FlexTDOA. In FlexTDOA, there is no single reference anchor. Instead, all the anchors in the system can be configured to take turns in transmitting the synchronization beacon. Similarly, the order of the anchors that respond to the beacon changes in a round-robin manner. Therefore, depending on the needs of the system, less anchors than the maximum available can respond to a beacon, which reduces clock drift errors caused by the delay between the first and the last response while allowing all anchors to participate in the localization process. FlexTDOA therefore exploits the full channel diversity of the environment, is not subject to single-link failures, and can maintain small errors even in large networks.

We implemented FlexTDOA in a localization system based on the Qorvo DW3000 UWB chipset [13]. We compare the proposed system against the classic TDOA approach and the standard range-based multilateration algorithm in a deployment of ten anchors and one tag in an office environment, in both line-of-sight (LOS) and non-line-of-sight (NLOS) conditions. We also evaluate the impact of several parameters on the ranging and localization accuracy, such as: the number of responses for each synchronization beacon for different system update rates, the number of anchors in the system, and the impact of changing the initiator and/or the order of responses.

Besides providing an increased robustness to harsh conditions, FlexTDOA distinguishes itself from previous approaches by not using tracking filters to estimate clock parameters, such as in [10], [11]. Instead, each receiver uses

carrier frequency offset (CFO) estimation to locally correct the relative time skew between its clock and the transmitter's clock. This allows administrators to easily add new anchors to the system without increasing its complexity. Moreover, the flexible scheduling scheme facilitates the deployment of the localization system in larger spaces where the anchors might be split over multiple rooms and only a subset of anchors should respond to an initiating message. The proposed scheme also preserves the location privacy of the user, since the user device remains passive in the localization process [7]. Therefore, no information of the user presence or their location is leaked to the infrastructure.

We make all the measurements available open-source¹, in order to facilitate the evaluation of other network parameters than those covered in the paper. The dataset can also be used to prototype NLOS error-mitigation techniques in the future.

The rest of the paper is organized as follows. Section III introduces the basics of the localization methods used throughout the paper and the particularities of the proposed scheduling scheme. Section IV presents the setup of the localization system, such as the hardware used for evaluation, the environment, and the anchor placement. Section V evaluates the most important parameters of the localization system, while Section VI compares range-based localization, the classic TDOA, and FlexTDOA. Section II reviews related work and highlights differences from previous approaches. Finally, Section VII sums up our contributions.

II. RELATED WORK

In the following, we will review the most important previous works on TDOA localization, with a focus on *DL* TDOA schemes which offer the best multi-user scalability.

A. SCALABLE UWB LOCALIZATION

In [10], a DL TDOA localization system which implements a clock synchronization protocol with a reference anchor is proposed. The authors mention that the system does not scale to large anchor networks. In a setup of eight anchors, the system obtained a 2D localization root-mean squared error (RMSE) of 14 cm and a 3D RMSE of 28 cm. In a comparable setup of seven anchors in LOS, FlexTDOA obtained a 2D RMSE of 16.2 cm and a 3D RMSE of 23.5 cm (averaged over all considered locations), so comparable to the ones in [10]. In [11], a similar approach to [10] is proposed, in which the pairwise clock error is tracked using a Kalman filter that can handle fluctuating reception periods.

In [14], the authors propose a DL TDOA scheme in which the anchors respond only to the previously-transmitted message instead of responding to a single synchronization beacon, as in our case. The mean and maximum localization errors obtained with a configuration of four anchors were 31 cm and 81 cm, respectively.

Although named concurrent ranging, the works in [15], [16] essentially implement the classic DL TDOA scheme.

¹<https://zenodo.org/record/7619764>

However, here the focus is on processing all the responses within a single reception period by exploiting the multipath information from the channel impulse response of the signal.

In [17], a TDOA localization system implemented using UWB devices called *ATLAS* is introduced. Although the system also uses an initiating anchor for synchronization, it is not clear if the tag is completely passive since it is mentioned that only whitelisted tags are processed at a localization server. Therefore, *ATLAS* relies on centralized processing to correct the TDOAs, whereas in *FlexTDOA* the localization is offloaded to the user device, preserving its location privacy. Iterations of *ATLAS* have been introduced in [18], [19]; however, in these works, the tag is active.

A localization system named *VULOC* that follows the principles of DL TDOA has recently been proposed in [20]. However, in *VULOC*, the initiator sends one additional message after all of the anchors have responded, which means that the tag has to listen to one extra message compared to *FlexTDOA*. Perhaps the most significant difference between *VULOC* and *FlexTDOA* is that we also propose a flexible, highly-configurable TDMA scheme for anchor transmissions, whereas in [20] it is mentioned that *VULOC* does not need a scheduling protocol because tags are passive. However, we argue that in large-scale building deployments there is a need to easily add or remove anchors from the system and to schedule their transmissions. As a result, we could easily evaluate the performance of *FlexTDOA* with up to ten anchors, whereas *VULOC* was evaluated only with five anchors. While *VULOC* also uses changing initiators, it does not evaluate the impact brought by the added channel diversity, as we do. In a laboratory setting, *VULOC* obtained a median error of 15.5 cm and a 90 % error of 23.6 cm. In a similar experiment in which the tag was placed in the center of the room, *FlexTDOA* obtained a median error of 15.4 cm and a 90 % error of 22.2 cm, so comparable to the errors achieved by *VULOC*.

Although it does not implement a DL TDOA scheme, the work in [21] proposes a scalable UL TDOA localization scheme called *TALLA*. The high-precision synchronization necessary for TDOA localization is maintained by a server, which can compute the clock parameters to synchronize to any reference anchor in the system. While this approach provides great flexibility in the case of large-scale deployments, it poses more privacy concerns since the network has information about the location of all tags in the system. In our approach, since the tag localizes itself, the network does not have any information about the users' locations.

Another important contribution of our work is that we evaluated the performance of a DL TDOA localization system in NLOS conditions *experimentally*, since most previous works either consider only LOS scenarios or base their observations on simulated data. In [20], the proposed TDOA system is also evaluated in NLOS conditions and an anchor selection method based on an empirically-chosen confidence parameter is proposed. However, in [20], the purpose of the evaluation is to demonstrate the effectiveness of the anchor

selection method, whereas we quantified the effect of the channel diversity brought by the scheduling method proposed in *FlexTDOA*. In [22], the authors propose a sensor placement strategy for cluttered environments that is validated through experimental data. A UL TDOA localization system that takes into account NLOS conditions has been proposed and evaluated experimentally in [23]. In [24], the authors propose an algorithm to select anchor pairs in a UL TDOA by taking into account errors caused by NLOS propagation.

In [25], the authors propose a framework designed for scalable indoor localization and implement it using UWB radios. The work in [25] is focused on the software, which enables cooperation between both fixed and mobile nodes in order to achieve seamless localization. In comparison, we focus on the specific TDOA localization algorithm, which enables scalable and accurate localization.

B. CLOCK OFFSET CORRECTION

In our TDOA scheme, we avoid tracking the clock parameters using Kalman filters like in previous works [10], [11]. Instead, we correct the relative clock offset between two devices directly at the receiver using the CFO estimation feature of the DW3000 chipset. The method has been described in [26] and the systematic error has been derived for single-sided two-way ranging (SS-TWR), A-TDOA, and SS-TWR with A-TDOA extension. The method has been evaluated experimentally but only for TWR schemes. A similar CFO correction is evaluated for a TDOA scheme in [27]. However, the proposed TDOA scheme is based on the alternative double-sided TWR (AltDS-TWR) method, in which the tag is active, which is different from the DL TDOA schemes evaluated in our work. A CFO-assisted synchronization algorithm for TDOA has been proposed in [28], in which the TOA at each receiver is compensated using the estimated CFO w.r.t. the reference (master) node. Although the TDOA scheme from [28] is different from the DL TDOA scheme used in our work, the correction method is similar. In addition, we evaluate the feasibility of the correction method through experiments with commercial UWB systems, whereas in the cited work the method was evaluated in simulations and using software-defined radios.

III. BACKGROUND

In this section, we introduce principles related to the proposed localization methods. For more in-depth details on UWB ranging and localization, we refer the reader to the papers [29], [30]. In Section III-A and III-B we explain how distance measurements and, respectively, TDOA measurements are obtained using UWB devices. In Section III-C, we describe the two approaches we use to solve the system of equations in order to estimate the user's location: either least-squares minimization between the measured and the calculated ranges or an extended Kalman filter (EKF). In Section III-D, we describe the scheduling scheme implemented in *FlexTDOA*.




Fig. 1: Localization based on TWR or on DL TDOA (the time periods are not to scale).

A. TWR

Range-based localization uses distances between a mobile target, called tag, and anchors with known locations to compute the location of the tag. The tag is found at the intersection of circles (in a 2D space) or spheres (in a 3D space) with a radius equal to the anchor–tag distances and centered at the anchors.

We use the following notations:

- $\vec{X}_{A_1}, \vec{X}_{A_2}, \dots, \vec{X}_{A_N}$ are the locations of anchors A_1, A_2, \dots, A_N , respectively;
- $\vec{X}_T(x, y, z)$ is the location of the tag;
- $d_{\vec{X}\vec{Y}}$ is the *true* distance between the nodes with locations $\vec{X}, \vec{Y} \in \{\vec{X}_T, \vec{X}_{A_1}, \dots, \vec{X}_{A_N}\}$;
- \tilde{d}_{XY} is the *measured* distance between the nodes $X, Y \in \{T, A_1, \dots, A_N\}$.

The true distance between the tag and anchor A_i is the Euclidean distance between their locations:

$$d_{\vec{X}_T \vec{X}_{A_i}} = \|\vec{X}_T - \vec{X}_{A_i}\|. \quad (1)$$

The measured distance between a tag and an anchor A_i can be written as:

$$\tilde{d}_{TA_i} = d_{\vec{X}_T \vec{X}_{A_i}} + \omega_{TA_i}, \quad i = 1, \dots, N, \quad (2)$$

where ω_{TA_i} is the measurement noise, which is modeled as a zero-mean Gaussian random variable with variance σ_i^2 .

To avoid synchronizing the anchors and the tag, the distances are usually obtained using TWR by exchanging at least two messages between the tag and each anchor [29]. We implemented the SS-TWR variant which uses two message exchanges between each anchor and the tag, illustrated in Fig. 1a. One distance is obtained as:

$$\tilde{d}_{TA_i} = c \cdot \frac{\Delta T_i^T - \Delta T_i}{2}, \quad (3)$$

where c is the speed of light. ΔT_i is the time between the arrival of the tag’s request at anchor A_i and the anchor’s transmission of the response message, as measured by the anchor. Similarly, ΔT_i^T is the time between the tag’s transmission of the request and the arrival of the anchor’s response, as measured by the tag.

Because each device measures its own processing time, there is no need to synchronize the clocks of neither the anchors nor of the anchors and the tag. However, due to the

fact that ΔT_i^T is measured by the tag while ΔT_i is measured on the anchor, the relative clock skew of the two nodes compounded with the significant value of ΔT_i can introduce large errors in the measurement. To counter this effect, the DS-TWR was proposed [31], in which the tag transmits a third message which is used by the anchor to measure the clock skew. Modern radios, such as the one we are using, can directly measure the relative clock skew by analyzing the carrier frequency of the received packet. Research shows [32] that DS-TWR and SS-TWR attain almost the same precision with clock skew estimation.

SS-TWR based localization (which we will alternatively call TWR localization) is attractive because it enables centimeter-level localization and does not need any synchronization between the devices. However, it does not scale well when increasing the number of anchors and tags, since it needs pair-wise message exchanges between each anchor and each tag in the system. DS-TWR is even less scalable than SS-TWR because it needs an extra message for the ranging. Moreover, because of the asymmetry, in DS-TWR the distance is computed by the responder (anchor), not the initiator (tag). Therefore, for navigation, the distance would need to be sent back to the user in another message. Given these drawbacks, we will use SS-TWR as a comparison baseline for FlexTDOA.

B. TDOA

An alternative localization technique uses the time *difference* between the arrival of two messages either at one device or at multiple clock-synchronized devices [4]. We define the true *distance* difference between anchors A_i and A_j relative to the tag as:

$$d_{\vec{X}_T \vec{X}_{A_i} \vec{X}_{A_j}} = d_{\vec{X}_T \vec{X}_{A_j}} - d_{\vec{X}_T \vec{X}_{A_i}} \quad (4)$$

$$= \|\vec{X}_T - \vec{X}_{A_j}\| - \|\vec{X}_T - \vec{X}_{A_i}\|. \quad (5)$$

The *measured* distance difference between anchors A_i and A_j relative to the tag is:

$$\tilde{d}_{TA_i A_j} = d_{\vec{X}_T \vec{X}_{A_i} \vec{X}_{A_j}} + \omega_{TA_j} - \omega_{TA_i}, \quad (6)$$

where, similarly to the TWR case, ω_{TA_k} is the measurement noise between the tag and anchor $A_k, k \in \{1, \dots, N\}$, modeled as zero-mean Gaussian noise with variance σ_k^2 .

In the noiseless case, any $N - 1$ distance difference measurements that form a minimum spanning subtree are sufficient for TDOA localization [33]. However, more redundant measurements can be used to improve the resilience to noise in a realistic setup. In TDOA localization, the tag’s position is found at the intersection of hyperbolae. This makes TDOA localization more sensitive to noise than range-based (or TOA) localization [6].

The *time* difference of arrival can be derived from Eq. (6) by dividing the distance difference by the speed of light:

$$t_{TA_i A_j} = \frac{d_{\vec{X}_T \vec{X}_{A_i} \vec{X}_{A_j}}}{c} = \frac{d_{\vec{X}_T \vec{X}_{A_j}} - d_{\vec{X}_T \vec{X}_{A_i}}}{c}. \quad (7)$$

The time differences can be computed either by a *passive* tag when the anchors transmit simultaneously their messages or by the anchors (or a central entity) when an *active* tag transmits a broadcast message which is received by all anchors. We will, respectively, call the two methods downlink (DL) and uplink (UL) TDOA.

To avoid synchronizing the anchors, we use a DL TDOA variant with reference and responding anchors previously used in [10], [11], [15], [16]. The scheme is illustrated in Fig. 1b. Anchor A_0 is the initiator and transmits a broadcast message, received by the tag at time T_0 . Anchors A_1 to A_N receive the message and then wait a period ΔT_i which includes the processing time and a delay necessary to avoid overlapping transmissions from successive anchors. The tag receives the responses at times T_1 to T_N .

For the general case in which A_i is the initiator and A_j is the responder, let us denote by $\Delta T_{ij} \triangleq T_j - T_i$ the difference between the time at which the tag receives the response (T_j) and the time at which the tag receives the request (T_i). In order to obtain only the TDOA from Eq. (7), we need to subtract the processing time ΔT_j and the TOF between anchors A_i and A_j (denoted by t_{ij}) from the timestamp difference ΔT_{ij} :

$$\tilde{t}_{TA_iA_j} = T_j - T_i - \Delta T_j - t_{ij}, \quad (8)$$

where $\tilde{t}_{TA_iA_j}$ is the *estimated* TDOA between the tag and the anchors A_i and A_j . The TOF t_{ij} is usually known because the anchors are placed at fixed, known locations.

Because ΔT_j is measured by A_j but subtracted from the timestamp difference ΔT_{ij} measured by the tag, the TDOA will contain an error due to the relative clock skew between anchor A_j and the tag. Because in TDOA localization multiple anchors respond to the same synchronization message, the processing times are longer for this scheme than in TWR localization. It is therefore crucial to correct the relative clock skew errors in TDOA localization [15]. To eliminate the errors, we leverage the capability of the DW3000 chipset to estimate the carrier frequency offset (CFO) between the local receiver and the remote transmitter [26]. The interval measured by the tag will contain an error due to the tag running at a different frequency from an “ideal” nominal frequency [13]:

$$\Delta \hat{T}_{ij} = k_T \Delta T_{ij} = (1 - e_T) \Delta T_{ij}, \quad (9)$$

where k_T is the multiplicative and e_T the additive error of the tag’s clock. Similarly, the processing time measured by the anchor A_j will contain an error:

$$\Delta \hat{T}_j = k_j \Delta T_j = (1 - e_j) \Delta T_j, \quad (10)$$

where k_j and e_j are the multiplicative, respectively, the additive errors of the anchor’s clock.

The relative CFO between the clocks of the tag and A_j as measured by the tag is [26]:

$$\kappa_{Tj} = \frac{k_T}{k_j} = 1 - e_{Tj}, \quad (11)$$

where κ_{Tj} and e_{Tj} are the multiplicative, respectively, the additive *relative* clock frequency offsets. The DW3000 chipset estimates e_{Tj} (expressed in ppm), which we use to correct the processing time²:

$$\Delta T_j^{corr} = \Delta T_j (1 - e_{Tj}). \quad (12)$$

C. LOCALIZATION ALGORITHMS

So far, we have discussed the basic principles to obtain the ranges or the range differences between the anchors and the tag. In order to estimate the user’s location, we need to solve a system of equations based on Eq. (1) and (5). We implemented two localization algorithms, each capable of operating with either TWR or TDOA data, each suiting different needs.

The first algorithm, AlgMin, solves the localization problem for a series of consecutive measurements using squared error minimization. This algorithm does not track the user’s location nor does it smooth the location estimates, and it is therefore suitable to evaluate the impact of several parameters (e.g., the number of responses or anchors) on the localization accuracy.

The second algorithm, AlgEKF, solves the localization problem using an EKF, by incrementally updating the location with each additional available measurement. This approach is advantageous because we do not need to wait for the minimum number of measurements (four in the case of TWR localization and five for TDOA localization) in order to update the tag’s location. However, it smooths the location estimates and hides the impact of noisy measurements. Therefore, we use it only when we compare several setups that generate a different number of equations per time slot in Section V-C.

Both algorithms start with the known anchor positions $\vec{X}_{A_1}, \vec{X}_{A_2}, \dots, \vec{X}_{A_N}$ and estimate the tag’s location \vec{X}_T . In our experiments, the anchors’ positions are determined using a self-localization algorithm, which is a variant of AlgEKF using TWR measurements between each pair of anchors, run over a period of several minutes. The locations of the anchors are determined once, at the beginning of the experiments, and kept fixed thereafter.

1) AlgMin

We use least square error minimization between the measured and the calculated distances. Therefore, for the TWR measurements:

$$\hat{\vec{X}}_T = \arg \min_{\vec{X}} \sum_i (d_{\vec{X}\vec{X}_{A_i}} - \tilde{d}_{TA_i})^2, \quad (13)$$

²The conventions for the additive clock offset are opposite for the DW1000 and DW3000 chipsets. For the DW1000 chipset, which uses the conventions from [26], if the additive clock offset is positive, then the receiver’s clock is running at a faster rate than the transmitter clock [34], while for DW3000 the reverse is true [35]. In this paper, we used the conventions for the DW3000 chipset.

where $d_{\vec{X}\vec{X}_{A_i}}$ is the calculated Euclidean distance between the 3D locations \vec{X} and \vec{X}_{A_i} and \tilde{d}_{TA_i} is the TWR measured distance between the tag T and anchor A_i (Fig. 1a).

For TDOA measurements, the equivalent minimization problem is as follows:

$$\hat{\vec{X}}_T = \arg \min_{\vec{X}} \sum_{i,j} (d_{\vec{X}\vec{X}_{A_j}} - d_{\vec{X}\vec{X}_{A_i}} - \tilde{d}_{TA_iA_j})^2, \quad (14)$$

where $\tilde{d}_{TA_iA_j} = c \cdot \tilde{t}_{TA_iA_j}$ and $\tilde{t}_{TA_iA_j}$ is the TDOA measurement performed by the tag T while listening to the two-way communication between nodes A_i and A_j (Fig. 1b). The TDOA measurement $\tilde{t}_{TA_iA_j}$ is computed using Eq. (8), where t_{ij} is obtained from the payload of the packets transmitted by the anchors.

The algorithm needs a good initialization to avoid converging to a local minimum of the function that is optimized [36]. In practice, we first initialize the algorithm using the starting position of the tag and for subsequent initializations we use the position of the tag estimated by AlgMin in the previous iteration.

2) AlgEKF

The location of the moving tag is computed iteratively using an EKF. Again, we implemented two filters one for the TWR and one for the TDOA. In both cases, the state of the filter is constituted by the position of the tag. The EKF assumes a system that can be described by the following general equations:

$$\begin{aligned} \vec{X}_k &= f(\vec{X}_{k-1}, U_k) + w_k \\ z_k &= h_k(\vec{X}_k) + v_k, \end{aligned} \quad (15)$$

where:

- \vec{X}_k is the position of the tag estimated at moment k ;
- $f(\vec{X}_{k-1}, U_k)$ is the motion model for the tag dependent on input command U_k and previous state;
- w_k is Gaussian noise $\mathcal{N}(0, Q_k)$ due to the uncertainty in the motion model, where Q_k is the covariance matrix of the motion model;
- z_k is the vector of measurements (either TWR or TDOA);
- $h_k(\vec{x})$ is a function that computes the expected value of the measurements given the state (position) \vec{X} of the tag;
- $v_k \sim \mathcal{N}(0, R_k)$ is Gaussian noise modeling the uncertainty in the measurements, where R_k is the covariance matrix quantifying the uncertainty in our measurements.

In our case, we do not assume a known motion model, so $f(\vec{X}_{k-1}, U_k) = \vec{X}_{k-1}$, corresponding to a static tag. We choose Q_k , the uncertainty in the model, large enough to cover the motion of the tag. In IV-B, we detail the choice of values for the covariance matrices using in the EKF. The choice of using a static model was deliberate, in order to evaluate the raw performance of our UWB localization method. In a real system, any knowledge about the actual

movement of the tag can be input into f to improve the accuracy of the system.

The function h and the vector z depend on the type of measurement performed. In the case of TWR:

$$\underbrace{\begin{pmatrix} \tilde{d}_{TA_i} \\ \tilde{d}_{TA_j} \\ \dots \end{pmatrix}}_{z_k} = \underbrace{\begin{pmatrix} d_{\vec{X}_k\vec{X}_{A_i}} \\ d_{\vec{X}_k\vec{X}_{A_j}} \\ \dots \end{pmatrix}}_{h_k}. \quad (16)$$

In the case of TDOA:

$$\underbrace{\begin{pmatrix} \tilde{d}_{TA_{i_1}A_{j_1}} \\ \tilde{d}_{TA_{i_2}A_{j_2}} \\ \dots \end{pmatrix}}_{z_k} = \underbrace{\begin{pmatrix} d_{\vec{X}_k\vec{X}_{A_{j_1}}} - d_{\vec{X}_k\vec{X}_{A_{i_1}}} \\ d_{\vec{X}_k\vec{X}_{A_{j_2}}} - d_{\vec{X}_k\vec{X}_{A_{i_2}}} \\ \dots \end{pmatrix}}_{h_k}. \quad (17)$$

In both cases, we use the classical prediction and update steps of the EKF filter to compute the tag's current position \vec{X}_k , starting from the tag's previous position estimate \vec{X}_{k-1} and the new measurements:

$$\vec{X}'_k = \vec{X}_{k-1} \quad (18)$$

$$P'_k = P_{k-1} + Q_k$$

$$K_k = P'_k H_k^T (H_k P'_k H_k^T + R_k)^{-1} \quad (19)$$

$$\vec{X}_k = \vec{X}'_k + K_k (z_k - h_k(\vec{X}'_k))$$

$$P_k = (I - K_k H_k) P'_k,$$

where P_k is the covariance matrix quantifying the uncertainty in our estimation and H_k is the Jacobian of the function h . The rest of the notations are defined above. Compared to the general EKF case for a system of equations, ours is simpler given that f is the identity function and its Jacobian is the unit matrix I .

The EKF update step can be done with a single measurement at a time (z_k is a single distance) or with multiple measurements at a time. To minimize the latency, we update the tag's position with every incoming measurement.

For the self-localization of the anchors, we use AlgEKF with TWR measurements. The algorithm works similarly to the tag's localization, with the following differences: (1) the state variable \vec{X}_k contains the 3D location of all anchors, (2) we use a much smaller covariance Q_k to account for the fact that the anchors are fixed, and (3) we perform an update using multiple ranges at a time (instead of updating the location for every incoming measurement). For simplicity, we align the resulted coordinate frame such that anchor A_0 has the location $(0, 0, 0)$, anchor A_1 is on the O_y axis, anchor A_2 is located in the O_{xy} plane with a positive x , and A_3 has a negative z coordinate.

D. SCHEDULING

In the "classic" TDOA approach, there is a single designated reference anchor which broadcasts the synchronization




Fig. 2: TDMA schedule used for both TDOA and TWR localization. Each slot in a frame belongs to a node which is the initiator in that slot and decides which K nodes to interrogate (depending on the current scheme).

message. The rest of the anchors respond to the broadcast in a predefined order. Instead, we propose, implement, and evaluate a *flexible* TDOA scheduling scheme in which all the anchors in the system can play the role of the initiator and the order of responses can also change.

We propose and implement a time-division multiple access (TDMA), scheme shown in Fig. 2, which can be configured for either TWR or TDOA localization. At this point, we do not differentiate between anchors and tags and instead consider all of them equally-participating nodes. The distinction will be made according to the implemented localization method.

The TDMA scheme is organized in *time slots*, which are comprised of a broadcast message sent by an initiating node, which we will call a *request*, and K responses from other nodes, where $K < N$ and N is the number of nodes in the system. Each response will provide a TWR measurement between the initiator node and the responding node and a TDOA measurement for each of the other nodes listening to the exchange. Inside a time slot, each transmission by a node occurs in a *subslot* with duration $t_{subslot}$. At the beginning of a time slot there is a guard time, followed by the request of the initiating node. The request includes the number of nodes that will respond, their ID, and the order of their response. All the listening nodes in the system process the request. If the initiator requested a response from the listening node in the subslot with index $k \in \{1, \dots, K\}$, the node will wait a period of $(k - 1) \times t_{subslot}^{resp}$ and then answer. During the last part of a time slot, the initiator processes the responses.

The time slots are organized in frames (Fig. 2). Each frame contains M time slots, each of them assigned to one of the nodes (anchors or tags). For instance, if slot i is assigned to node j , then, in all frames, node j will be the initiator in slot i , and will decide which K nodes to interrogate based on the currently-selected scheme. More than one slot can be assigned to the same node and there may be nodes that have no slot assigned (e.g., a tag that is passively listening and only using TDOA measurements for localization).

The network starts with the default programmed parameters. The nodes first listen for a few seconds before transmitting. From any received message, they can determine the cur-




Fig. 3: Localization packet format used by the anchors and the tag.

rent slot and synchronize with the TDMA schedule. All the TDMA parameters are decided by the system administrator and can be changed at runtime: M , K , the slot assignment, the list of nodes that are allowed to be interrogated (for example, to avoid the interrogation of a passive tag). To change the parameters, the system administrator connects via USB to any of the nodes, uploads the new configuration into the node, and requests the node to broadcast it. The node broadcasts the information using a special packet format. All the nodes that hear the new configuration will rebroadcast it for several seconds. In a real system, we could envision an algorithm that makes the decisions instead of the administrator.

All localization messages exchanged by the anchors or the tag have the format shown in Fig. 3. It consists of: the message type (request or response), the source ID, the list of interrogated anchors (in case of a request), and a payload in which the sender can include a previously-acquired TWR measurement. Each packet contains the current slot ID, which is a 32-bit counter incremented continuously. This counter is used by the nodes to determine the current slot in the frame (modulo M) and to stay synchronized with the TDMA schedule.

In the current implementation, the node IDs, K , and M are 8-bit packet fields and variables, which implies a maximum of 256 devices and 255 slots in a frame. A larger network can be accommodated by choosing larger packet fields. At the moment, a slot can only be assigned to one node. For a network spread over a larger area, this should be changed so that nodes that are sufficiently far apart (i.e., not in radio range) can reuse the same slot.

To configure the TDMA scheme to perform TWR localization, the tag will be set as initiator in all slots and the anchors will be the responders. In one time slot with K responses, the tag obtains K raw distance measurements which are input to the multilateration system to estimate the tag's location.

To perform DL TDOA localization, only the anchors will be initiators, interrogating other anchors, while the tag will be a passive listener. Depending on how we choose the initiators and the responders, we can derive four variants of TDOA localization:

- *Fixed initiator, fixed responders* (FI-FR), or the “classic” TDOA, with a designated reference anchor (or initiator), where the other anchors in the system respond in a fixed order according to their index. In this case, all slots in a frame are assigned to a single anchor (the initiator).




Fig. 4: Custom build UWB Node: a completed hardware and software UWB node which is battery powered and capable to expose ranging information over multiple serial communication interfaces.

In every slot, the initiator interrogates the same list of (responding) anchors.

- *Fixed initiator, changing responders* (FI-CR) with a designated reference anchor, where the responding anchors change in a round-robin (RR) manner. In this case, in the first slot, the initiator will pick K anchors to interrogate (from the allowed list), in the next slot the next K , and so on, wrapping around.
- *Changing initiator, fixed responders* (CI-FR), in which the initiator changes every time slot in a RR manner and the rest of the anchors respond in a (fixed) ascending order of their index. In this case, there are multiple initiators (anchors that have slots assigned) and each of them interrogate the same fixed list of responders.
- *Changing initiator, changing responders* (CI-CR), in which both the initiator and the responder order change every time slot in a RR manner. This is the scheme used in FlexTDOA. Multiple initiators interrogate K other nodes at a time from the pool of allowed responders.

We note that the TDMA scheme also allows the anchors to localize themselves. The initiator measures the distance between itself and each responder (other anchors) using TWR. These distances are then transmitted as piggyback payloads of future messages (requests or responses) of the initiator. As a result, simply by listening to the message exchanges, the localization engine can compute the location of each anchor in the local coordinate system using the self-localization algorithm described in Section III-C.

IV. EVALUATION SYSTEM

In this section, we present the localization system used to evaluate the TWR and TDOA algorithms. In Section IV-A, we describe the hardware used; in Section IV-B, we present the settings used for the UWB radio, scheduling algorithm, and EKF. In Section IV-C, we describe the environment in which we acquired the measurements and the placement of the anchors.

A. HARDWARE

For the experimental evaluation, we designed and fabricated our own UWB node, shown in Fig. 4. We used the Qorvo DWM3000 wireless transceiver [13], which implements the IEEE 802.15.4 standard [37].

At the core of the UWB node is an Arm Cortex-M4 based STM32F429ZIT6 microcontroller with 2 MB of Flash memory, 256 KB of SRAM memory and a frequency of 168 MHz [38]. All the ranging and scheduling algorithms are running on the on-board MCU and the ranging results are transmitted over an USB 2.0 port. However, the board can be configured to transmit the information over any other interface, like I²C, SPI or UART. The entire software stack was written in-house.

The UWB node is powered by a single Li-ion rechargeable battery with a capacity of 6Wh, providing over 15h of autonomy. This allowed us to easily place the nodes independently of available power sources.

The power management is implemented using an MCP73830 IC, which is a 1 A Single-Cell Li-Ion battery charge management controller. This allows easy charging over the USB port. The MCU and the radio are powered through a 500 mA LDO voltage regulator, NCV8705.

B. SYSTEM SETTINGS

We configured the UWB transceiver to operate on channel 5 (6.5 GHz) with a preamble length of 128 symbols, a 6.8 Mb /s data rate, and a pulse repetition frequency of 64 MHz.

The duration of one time slot in the TDMA scheme shown in Fig. 2 is computed as:

$$t_{TS} = t_{\text{guard}} + t_{\text{subslot}}^{\text{req}} + t_{\text{process}}^{\text{req}} + K \cdot t_{\text{subslot}}^{\text{resp}} + t_{\text{process}}^{\text{resp}} \quad (20)$$

where:

- $t_{\text{guard}} = 250 \mu\text{s}$ is the guard time at the beginning of the time slot necessary to wait for the responders to go into the receive mode;
- $t_{\text{subslot}}^{\text{req}} = 2000 \mu\text{s}$ is the period it takes the initiator to send a request;
- $t_{\text{process}}^{\text{req}} = 250 \mu\text{s}$ is the period during which the initiator enters the receive mode and the responders process the request message;
- K is the number of responses;
- $t_{\text{subslot}}^{\text{resp}} = 250 \mu\text{s}$ is the duration of one response in a subslot;
- $t_{\text{process}}^{\text{resp}} = K \cdot 600 \mu\text{s}$ is the time allocated for the initiator to process the responses.

The subslots have different durations for sending a synchronization beacon or a response because of additional time needed to prepare request message sent it over UWB radio and send measurement data over serial port. Table 1 shows the duration of one time slot for each number of responders.

For both EKF filters (based on TWR and TDOA measurements) we chose a variance of the model uncertainty of $\sigma_Q^2 = 100 \text{ cm}^2$, which accounts for the motion of the tag between measurements and assumes a maximum speed of the

Table 1: Duration of one time slot (t_{TS}) for each number of responders K .

K	1	2	3	4	5	6	7	8	9
t_{TS} (ms)	3.35	4.20	5.05	5.90	6.75	7.60	8.45	9.30	10.15




Fig. 5: Office setup.

tag of 10 cm/s. We chose a variance of $\sigma_R^2 = 10 \text{ cm}^2$ for the measurement noise, which was based on the measurement noise we obtained during experiments.

For the EKF filter used for the self-localization of the anchors, we used $\sigma_Q^2 = 1 \text{ cm}^2$ because the anchors are static. The location of the anchors is determined once, at the beginning of the experiments and kept fixed thereafter.

C. ENVIRONMENT AND ANCHOR PLACEMENT

We evaluate the localization systems in the office shown in Fig. 5. The 3D anchor placement is shown more clearly in Fig. 6. Five of the anchors (A_0 to A_3 and A_9) are fixed on the ceiling using metallic bars, while the rest of the anchors are either placed on the ground (A_4) or on tables (A_5 to A_8). The location of the anchors is determined using the self-localization algorithm described in Section III-C. We validated the resulting locations using a laser level and a laser rangefinder, both with mm-level precision.

To accurately measure the ground truth (GT) of the tag, we have built a custom electronic linear actuator shown in Fig. 7. We used a 140 cm-long aluminum rail, a stepper motor, a timing belt, an aluminum trolley carrying the tag, and a driver connected via USB to the computer. The speed of actuator is 10 cm/s and the positioning resolution is 0.1 mm.

The actuator continuously sends the current position over the USB serial port. To measure the GT of the node in our positioning system, we placed the actuator at a known location relative to the anchor A_0 , which is the origin of the coordinate system. Unless stated, the actuator moves only along the x axis of the coordinate system, between the x axis coordinates of approx. (110, 270) cm, as shown in Fig. 6. We have two experiments in which the actuator is placed in a 3D orientation.

In all of the experiments, we precisely align the actuator

with the local coordinate system using a laser rangefinder and a laser level.

V. EVALUATION OF SYSTEM PARAMETERS

In this section, we evaluate the impact of several factors on the localization accuracy: the order of response of an anchor, the number of responses in a time slot for the maximum update rate of the localization system and for lower update rates, and the number of anchors available. We evaluate these parameters for localization algorithms that use distance measurements (obtained using TWR), which we call “TWR localization,” and for the proposed FlexTDOA system, called simply “TDOA localization.” The goal of the comparison between TWR and TDOA localization is to evaluate the impact of system parameters of both distance and TDOA measurements.

Unless explicitly mentioned, we use the AlgMin algorithm described in Section III-C to estimate the user’s location. The localization error is computed as the Euclidean distance between the 3D GT location and the estimated location:

$$e = \sqrt{(x - \hat{x})^2 + (y - \hat{y})^2 + (z - \hat{z})^2}, \quad (21)$$

where (x, y, z) and $(\hat{x}, \hat{y}, \hat{z})$ are the Cartesian coordinates of the GT, respectively, the estimated location.

We use boxplots to illustrate the error distributions. In a boxplot, such as the ones in Fig. 8, the box is drawn from the first to the third quartiles (or, respectively, the 25th and the 75th percentiles), which is also known as the interquartile range (IQR). Boxplots drawn for samples that can take negative and positive values (e.g., the distance and TDOA errors) have whiskers that extend from the 5th to the 95th percentiles. For strictly positive errors (for instance, the localization errors which are computed as the Euclidean distance between the estimated and the ground truth locations), the whiskers extend from the 0th to the 95th percentiles. The reasoning is that, when we plot the distribution of *absolute* errors, we are interested in the minimum value of the error. We omit plotting the outliers for simplicity. We will frequently report the 95th percentile, which we will alternatively call the 95% error (or P_{95}) for short, which represents the value below which 95% of the errors are found.

A. ORDER OF RESPONSE

First, we investigate how the TWR or TDOA *measurement* error changes depending on the order of the response in a time slot. In Section III-B, we mentioned that it is important to correct the processing times to account for the relative clock skew between a transmitter and a receiver. With a longer time period elapsed between the initiator’s request and an anchor’s response, the error in a TWR or a TDOA measurement will increase because the additive relative clock offset ϵ_{T_j} will be multiplied by a longer processing time ΔT_j in Eq. (12). Although we compensate for these errors using CFO estimation, the correction is imperfect. Therefore, we expect to see larger TWR and TDOA errors for higher response indexes in a time slot.




Fig. 6: Setup of the anchors and the tag in the (a) xy and (b) xz planes. The anchors are denoted by A_0 to A_9 . We evaluate the localization accuracy at four positions of the linear actuator, along which the tag moves, denoted by P1 to P4.




Fig. 7: Ground truth linear actuator: An aluminum trolley that carries the tag and returns over USB the position of the tag relatively to the zero point of the actuator.

To evaluate the magnitude of the errors, we perform an experiment in which the tag is kept unmoved, in order to avoid any accuracy loss due to the movement of the tag. We configure the system to compute either the distance (using TWR) between each anchor and the tag or the TDOA between the tag and each pair of anchors. We use the maximum number of anchors ($N = 10$) and of responses ($K = 9$ for TDOA and $K = 10$ for TWR). To compute the distances, the tag is always the initiator and the anchors respond to its broadcast. The order in which the anchors respond changes every time slot using round-robin (RR) scheduling. To compute the TDOA, the tag only listens to the messages and both the initiating anchor and the responding nodes rotate every timeslot in a RR manner (CI-CR scheduling). We run

each experiment for 3 min and compute the TWR and TDOA errors for each possible index of response.

Fig. 8 shows the TWR and TDOA error distributions for each order of response. Table 2 shows the statistics of the errors in the first and last response: the mean, standard deviation (σ), IQR, and whisker spread ($P_{95} - P_5$). The whisker spread of TDOA errors increases with 14.7 cm between the last and the first in the list of responses. TWR errors are less affected by the order of response than TDOA errors: the whisker spread of TWR errors increases with only 4.1 cm between the last and first order of response. The results are in line with the theory: with a longer waiting time between the initiator's message and the response, the error due to the relative clock skew increases and the measurements are corrupted by noise. Over many measurements, the mean error remains small, in absolute terms, regardless of order of response. What increases significantly is the noise in each measurement. This suggests that, for a relatively static tag, a larger number of responses and averaging would increase the accuracy (since it uses the time more efficiently), while for a fast moving tag, where averaging is not possible without a motion model, a smaller number of responses would perform better. We further explore this trade-off in Section V-C.

B. NUMBER OF RESPONSES

Since the TWR and TDOA *measurement* error increases with the order of response in the time slot, we investigate to what extent the *localization* accuracy changes with the number of responses in a time slot. In schemes with a fixed order of responses (FI-FR and CI-FR), the number of responding anchors is given by the number of anchors in the system (so that all anchors can participate in the localization process). In schemes with a changing order of response, however, we can decrease the number of responses to be smaller than the number of anchors in the system and still have all the anchors




Fig. 8: The distribution of (a) TWR and (b) TDOA errors (expressed in cm using the speed of light) against the order of response aggregated over all anchors.




Fig. 9: Localization error of TWR and TDOA localization depending on the number of responses in a slot. The error increases for more responses due to the longer period between the first and the last response, which increases the effect of clock drift estimation error.

Table 2: Statistics of raw measurement errors in the first and last (9th) response: mean, standard deviation (σ), IQR, and difference between the 95th (P_{95}) and 5th (P_5) percentiles.

Measurement	Response index	Mean (cm)	σ (cm)	IQR (cm)	$P_{95} - P_5$ (cm)
TWR	1	-1.5	4.8	5.7	18.1
	9	-1.2	6.9	8.9	22.1
TDOA	1	0.7	9.0	11.8	29.3
	9	1.3	13.5	17.9	44.0

participate, only in different time slots.

For this evaluation, we keep the same setup as in Section V-A, so using $N = 10$ anchors, but we vary the number of responses $K \in \{1, \dots, 9\}$ and let the tag move on the trolley.

Fig. 9 shows the localization error of TWR and TDOA localization with a varying number of responses. As expected, the localization error is the smallest for the minimum number of responses. However, the increase in the mean and IQR of the localization error with a higher number of responses is almost negligible: less than 3 cm between the maximum and the minimum number of responses for both TDOA and TWR localization. The 95% TDOA localization error has a slightly

larger increase than the IQR with nine vs. one responses, of 5.2 cm.

The results seem counter-intuitive: more responses in a slot yield more measurements per unit of time, which should decrease the error. This is not the case because of several reasons. (1) Even in the worst case (one response in a slot), the system generates about 300 measurements per second. This is enough to approach the maximum theoretical performance given the relatively slow speed of the tag of 10 cm/s. So, for a slow tag, we get a better performance with few precise measurements than with many noisy measurements. (2) Even with a lower number of responses in a slot, all anchors remain engaged as initiators and pick their responders round-robin from all other anchors, thus preserving the diversity. We further investigate the effect of K , under additional constraints, in the next section.

C. NUMBER OF RESPONSES FOR DIFFERENT TAG SPEEDS

Although using only one response per time slot yields the smallest error spread, this configuration has at least two disadvantages. First, over a fixed time period, the number of TDOA measurements decreases with the number of responses per time slot, because of the overhead added by




Fig. 10: (a) The number of TWR/TDOA measurements (N_{meas}) obtained per second vs. the number of responses (K) in a slot and (b) the ratio between N_{meas} and the number of messages (N_{msg}) obtained per second vs. K .

the initiator’s request. This trend is illustrated in Fig. 10a, which shows the number of distance or TDOA measurements (N_{meas}) obtained per second for $K \in \{1, \dots, 9\}$ responses. With $K = 9$, the number of measurements per second is approx. $3 \times$ larger than with $K = 1$. Second, because of the same reason, the energy consumed by the tag to receive a certain number of TDOAs increases as the number of responses decreases. We can compute the ratio between the number of measurements and the number of exchanged messages (N_{msg}) over the same time period, which is an indicator of the efficiency of the tag. This ratio (denoted by R_m) is illustrated in Fig. 10b. With $K = 4$, the efficiency is $1.6 \times$ higher than with $K = 1$, but the slope declines as K increases.

In a real system, the system administrator will face the question of choosing K to optimize the location accuracy under the specific conditions: available power to the system, available air time for transmissions, and maximum speed of the tag. To give an insight on the trade-offs involved, we evaluate how the 3D localization error changes with K when we keep constant over the same time period either the total transmission time (T_{TX}) or the number of exchanged messages (N_{msg}).

Additionally, to measure the effect on errors of a tag that is moving faster than our ground truth trolley, we deliberately slow down our system by introducing some idle time. To achieve all this, we group one or more time slots plus some the necessary idle time in a frame which has a fixed length. We call this the repetition period (Fig. 11). In our experiments, it takes the values $T_{rep} \in \{0.02, 0.5, 1\}$ s. By increasing the repetition period with additional idle time, we simulate the scenario in which the tag is listening for localization messages at a lower rate or, equivalently, the situation in which the tag is moving at a higher speed.

In both experiments, we used $N = 10$ anchors and varied the number of responses $K \in \{1, 4, 9\}$. We therefore want to find the repetition time (or tag speed) for which more TDOA measurements compensate for the clock drift error incurred by a higher number of responses either when we have a fixed time budget (T_{TX}) or a fixed energy budget (N_{msg}). We consider that the number of received messages is proportional

Table 3: Setup for experiments with approximately the same transmission time (T_{TX}).

K	N_{TS}	N_{meas}	T_{TX} (ms)
9	1	9	10.15
4	2	8	11.80
1	3	3	10.05

to the energy consumed by the tag.

For the evaluation, we perform the localization using the AlgEKF algorithm from Section III-C, which updates the location for every incoming measurement. We prefer using the AlgEKF over the AlgMin because the latter needs a minimum of four TDOA equations for one location update. If that were the case, would have to update each experiment at different rates, which can bias the results.

1) Same transmission time

We first evaluate how the 3D localization error changes when the total transmission time (T_{TX}) is constant and the number of responses varies. In all three experiments, during each repetition period, we have approximately 10 ms of transmission time (1 slot with 9 responses, 2 slots with 4 responses, or 3 slots with 1 response). The rest is idle time. Given our tag’s actual speed of 10 cm/s and the repetition period of 20 ms, 0.5 s, and 1 s, we essentially simulate a tag speed of 20 cm/s, 5 m/s, and 10 m/s, respectively.

Fig. 12 shows the error distributions for all combinations of number of responses and number of TDOA measurements per repetition period ((K, N_{meas})) and for all repetition periods (T_{rep}). When T_{rep} is the lowest, the highest localization accuracy is obtained for the minimum number of responses, $K = 1$, as in Section V-B. However, as T_{rep} increases, so does the average localization error for $K = 1$. At $T_{rep} = 1$ s, for $K = 1$, the mean error is 10 cm higher and the spread is almost double compared to $K = 9$.

The error is so high for $K = 1$ because the number of TDOAs obtained every second is lower than the minimum number of TDOAs needed to obtain a location. Because the EKF updates the tag’s location for every incoming measurement, we do not need to wait for the minimum number of TDOAs. However, the location estimate suffers as a result.

It is interesting to note that, for $T_{rep} = 0.5$ s, the localization error is lower for $K = 4$ and $N_{meas} = 8$ than for $K = N_{meas} = 9$. This means that one extra TDOA cannot compensate for the higher clock drift errors of $K = 9$ responses. However, at $T_{rep} = 1$ s the update rate becomes low enough such that the mean 3D error is approximately equal for $K = 4$ and $K = 9$. In this case, the error spread is actually smaller for the highest number of responses. Therefore, at a high tag speed (low update rate), it is preferable to use a high number of responses.




Fig. 11: To evaluate the optimal K under various simulated situations (time constraints, energy constraints, and tag speed), we create localization frames of fixed duration T_{rep} , containing N_{TS} time slots and the required idle time. Each time slot contains one request and K responses. We do two experiments: (1) we vary N_{TS} and K while keeping the total air time ($T_{TX} = N_{TS} \cdot t_{TS}$) constant, thus simulating time constraints, and (2) we vary N_{TS} and K but we keep the total number of messages exchanged ($N_{msg} = N_{TS} \cdot (K + 1)$) constant, thus simulating energy constraints. In both cases, we vary the idle time to simulate a tag moving at various speeds.




Fig. 12: Error distributions for the same TX time (T_{TX}). By increasing T_{rep} , we simulate a higher speed of the tag.

2) Same number of messages

We consider that the energy consumed by the tag is proportional to the number of received messages³ denoted by N_{msg} . Therefore, we want to find out which configuration is optimal for a fixed number of messages (or energy consumption). Note that, although the number of messages is fixed, the number of TDOAs that can be extracted from these messages increases with the number of responses in one time slot.

We keep the number of messages (N_{msg}) fixed during each repetition period and we vary the repetition time as in the previous experiment. Fig. 13 shows the error distributions for a fixed number of exchanged messages. Similar to the previous case, for a high update rate, $K = 1$ is the optimal number of responses. However, as T_{rep} increases, it is more beneficial to have more TDOAs than to minimize the clock drift error. At $T_{rep} = 0.5$ s, the mean error and the error spread in all three configurations are approximately equal. Beyond this T_{rep} value, it pays off to maximize the number of responses per time slot.

In a real deployment, the number of responses cannot be adjusted based on each user's needs, since the mobile devices are passive and the localization system may not even be aware of their presence. Therefore, we should aim to find the

³Although the transmission time T_{TX} is higher for the "Same N_{msg} " configuration than for the others in Fig. 11, the transmission time also includes the time during which the responding anchors process the initiator's request and the guard time between the anchors' responses. Because this time is known, we put the tag in the idle mode between the reception of two successive responses, such that the time during which the tag is in the receive mode is shorter than T_{TX} .

Table 4: Setup for experiments with the same number of messages (N_{msg}).

K	N_{TS}	N_{meas}	N_{msg}
9	1	9	10
4	2	8	10
1	5	5	10




Fig. 13: Error distributions for the same number of transmitted messages (N_{msg}).

optimal number of responses that covers a wide range of user movement patterns, update rates, and accuracy constraints. Based on our data, we suggest that using up to $K = 9$ responses is preferable to using the minimum number of responses for several reasons. First, at the maximum rate, the decrease in the localization accuracy is only 3 cm in the mean and 8 cm in the spread for $K = 9$ compared to $K = 1$. However, at $T_{rep} = 1$ s, the mean localization error is 10 cm lower and the IQR twice as small for $K = 9$ compared to $K = 1$. Therefore, it is preferable to use the maximum number of responses in most cases.

D. NUMBER OF ANCHORS

During the previous experiments, we kept the highest number of anchors available ($N = 10$). However, depending on the deployment, it might be unfeasible to have that many anchors available in the tracking area. Therefore, we evaluate the 3D localization error (computed as in Eq. (21)) when varying the number of anchors participating in the localization between 5 and 10.

As illustrated in the anchor setup from Fig. 6, the first four




Fig. 14: The 3D localization error of range-based localization depending on the number of devices.

anchors (which are used in the smallest configuration) are placed close the ceiling and the fifth anchor on the ground, to ensure enough spread on the z axis. The first four anchors are placed on the perimeter of the tracking area and the fifth anchor approximately in its center. For the configurations with more anchors, we add one anchor in the order of its index (e.g., the setup with 6 anchors uses the anchors A_0 to A_5). In order to eliminate the influence of the number of responses on the localization error, we set the number of responses to $K = 4$ (which is the minimum number of responses for the minimum number of anchors $N = 5$).

Fig. 14 shows the 3D localization error for TWR and TDOA localization when varying the number of anchors. The general trend is that the localization error decreases for more anchors. However, the decrease is much smaller for TWR than for TDOA localization. For TDOA localization, the mean localization error with 10 anchors is 1.7 cm smaller than with 5 anchors, while the 95% error decreases with 2.1 cm. For TWR localization, the gains in the localization accuracy are even smaller.

There is an anomaly in TDOA localization, in which the errors are the highest for $N = 6$ anchors, and not for the minimum number of anchors, as one would expect. These errors are caused by the unfavorable geometry of this configuration, since anchors A_4 and A_5 are placed quite close to each other. This configuration results in large localization errors when the tag is furthest away from these anchors. We also noticed that TDOA localization is more sensitive to a poor anchor placement than TWR localization. This means that the deployment should be done according to a simulation of the expected localization errors for different configurations. Strategies for optimal anchor placement are outside the scope of the paper, but we refer the reader to [22], [39] for analyses of sensor placement in TDOA localization.

The improvement brought by using more anchors might seem modest, since we could perhaps expect the localization accuracy to improve more dramatically with more anchors. However, the improvement is bounded by the accuracy of the *technology*, which is reflected in the distance/TDOA errors from Fig. 8. Since the measurements for different numbers

of anchors were acquired in ideal conditions, i.e., with LOS between each pair of nodes, adding more devices cannot improve the localization accuracy beyond the capabilities of the technology itself. However, having more anchors than the minimum necessary is more beneficial in challenging conditions, where part of the nodes are in NLOS with each other. We will evaluate this scenario in Section VI-B.

VI. COMPARISON OF LOCALIZATION METHODS

In this section, we compare the localization accuracy of the considered localization methods. In Section VI-A, we first compare the four variants of TDOA localization presented in Section III-D: FI-FR (or the classic TDOA), FI-CR, CI-FR, and CI-CR (or FlexTDOA). The goal is to evaluate the improvement brought by changing only the initiator, only the list of responders, or both. In Section VI-B, we compare only the classic TDOA, FlexTDOA, and TWR localization in a NLOS scenario. Throughout this section, we use the AlgMin algorithm to estimate the user's location.

A. FIXED VS. CHANGING INITIATOR AND/OR RESPONSE ORDER

In the classic (FI-FR) TDOA scheme, there is a designated reference anchor (analogous to our initiator) which broadcasts a message. All the anchors in the system respond in a predefined order with a certain time delay in order to avoid overlapping answers at the receiver. This approach is not ideal because it does not fully exploit the channel diversity of all anchor pairs. Moreover, if the tag does not have a good link to the reference anchor, all the measured TDOAs will be biased or noisy.

As we mentioned in Section III-D, we can derive three other TDOA schemes from the classic approach (denoted by FI-FR): with a fixed initiator but changing the order in which anchors respond (FI-CR), changing the initiating anchor every time slot but keeping the order of the responding anchors fixed (CI-FR), and changing both the initiator and the order of responses every time slot (CI-CR, implemented in FlexTDOA). In the FI schemes, anchor A_1 will be the initiator. The FI-FR experiences the lowest and




Fig. 15: Comparison between FI-FR (classic TDOA), FI-CR, CI-FR, and CI-CR (FlexTDOA) in LOS at location P1.

CI-CR the highest channel diversity. Therefore, we expect these schemes to have the worst and, respectively, the best localization accuracy, assuming that all anchors have equally good propagation conditions to the tag. However, in case the link between some anchors and the tag leads to higher errors, the CI-CR will be at a disadvantage because those anchors can become initiators and corrupt all the TDOAs in one time slot.

In this part, we evaluate to what extent the channel diversity improves the localization accuracy in LOS conditions. We evaluate the localization errors for $(N, K) \in \{(5, 4), (7, 6), (10, 9)\}^4$ at three positions of the rail on which the tag moves, denoted by P1, P2, and P3 in Fig. 6. Position P1 is in the center of the room, parallel to the XY plane, where we should have the highest accuracy. Position P2 and P3 are inclined relative to the XY plane, so that we can evaluate the errors at multiple tag heights.

Each experiment, i.e., for each location, for every method and for each (N, K) combination, lasts one minute and a half, which gives us approx. 7000 location estimates obtained with AlgMin. When aggregating the measurements over the 3 locations of the rail, we obtain approx. 21,000 location estimates in LOS for each method and (N, K) combination. The number of actual TDOA measurements in each experiment depends on the (N, K) configuration; for instance, for $(N, K) = (10, 9)$, we obtain approx. 80,000 TDOAs in one run of the experiment.

Fig. 15 compares the four TDOA variants at location P1. The FI schemes have the highest localization errors for five and seven anchors, but the lowest for ten anchors. One explanation for the reverse trend in the case of ten anchors is the fact that, in this case, some of the anchors on the ground or on the tables will become initiators and they do not have ideal propagation conditions to all other anchors. Therefore, in this scenario, the localization errors will be slightly higher for the CI schemes than for the FI schemes, where the initiating anchor is placed at an ideal location.

We observe that changing the order of responders does not bring a significant improvement to the localization accuracy. Instead, the *initiator* plays a crucial role in the localization

⁴In each case, $K = N - 1$ so that, even for a fixed order of responses, all anchors get to participate in the localization process.

Table 5: Median (P_{50}) and 95th percentile (P_{95}) of the 3D localization error for all combinations of changing/fixed initiators and order of responses. The errors are presented for the positions P1, P2, and P3, shown in Fig. 6.

Pos.	Method	N = 5, K = 4		N = 7, K = 6		N = 10, K = 9	
		P_{50} (cm)	P_{95} (cm)	P_{50} (cm)	P_{95} (cm)	P_{50} (cm)	P_{95} (cm)
P1	FI-FR	13.6	30.2	14.1	28.5	7.9	16.1
	CI-CR	12.3	26.6	10.8	23.0	10.1	19.4
P2	FI-FR	23.0	71.0	20.8	59.9	17.6	46.9
	CI-CR	19.1	69.4	15.2	41.4	13.8	30.6
P3	FI-FR	29.0	79.3	18.4	67.3	11.4	29.1
	CI-CR	25.2	78.1	19.1	62.3	15.6	38.2




Fig. 16: Distributions of the localization errors of classic TDOA (FI-FR) and FlexTDOA (CI-CR), in LOS, aggregated over all the evaluated positions (P1, P2, P3).

process because all the TDOAs in a time slot are computed with respect to its time frame. Any error in timestamping the initiator's message at the tag will affect all the TDOAs in that time slot. For this reason, from now on, we will consider only the FI-FR and CI-CR schemes, which we will alternatively call the classic TDOA and FlexTDOA, respectively. FlexTDOA is our proposed method for improving TDOA localization.

Table 5 shows the localization errors of the classic TDOA (FI-FR) and FlexTDOA (CI-CR) at the three considered locations: P1, P2, and P3. Fig. 16 shows the distributions of the localization errors for the same methods, aggregated over all considered locations. In LOS, where the channel between all anchors and the tag should be "ideal" (i.e., without obstructions), the FlexTDOA yields a similar or slightly better accuracy compared to the classic TDOA.

B. NLOS PROPAGATION

Obstacles between the nodes of a localization system are common in real-life scenarios. In this part, we evaluate the NLOS performance of the three localization approaches considered so far: based on TWR, the classic TDOA, and FlexTDOA.

We performed measurements at two positions of the rail on which the tag moves. The positions are denoted by P1 and P4




Fig. 17: Photos of the setups used to obtain NLOS measurements. The setup in Fig. 17a, which corresponds to the actuator position P1 from Fig. 6, includes one aluminum panel placed as an obstruction between anchor A_1 and the tag. The setup in Fig. 17b, which corresponds to the actuator position P4 from Fig. 6, includes two aluminum panels placed as obstructions. The rightmost aluminum panel blocks the LOS between the tag and A_1 and partially A_5 . The leftmost aluminum panel blocks A_0 and A_9 . There is significant interference due to multipath propagation for A_4 , A_7 and A_8 .

Table 6: Median (P_{50}) and 95th percentile (P_{95}) of the 3D localization error in the FI-CR and CI-CR TDOA schemes, in NLOS, in positions P1 and P4 (from Fig. 6), using the obstructions shown in Fig. 17a and Fig. 17b, respectively.

Pos.	Method	$N = 5$		$N = 7$		$N = 10$	
		P_{50} (cm)	P_{95} (cm)	P_{50} (cm)	P_{95} (cm)	P_{50} (cm)	P_{95} (cm)
P1	TWR	15.5	55.4	11.8	45.5	8.1	29.4
	Classic TDOA	16.5	45.1	17.3	48.1	14.5	48.8
	FlexTDOA	15.5	44.6	11.8	35.4	10.6	32.4
P4	TWR	26.9	64.4	24.0	60.7	23.7	44.4
	Classic TDOA	55.5	167.7	32.5	89.6	32.6	106.9
	FlexTDOA	40.3	176.5	22.5	79.3	22.3	59.4

in Fig. 6. At P1, we placed a panel covered in aluminum foil between the anchor A_1 and the tag, shown in Fig. 17a. At P4, we placed two such panels, shown in Fig. 17b. The panel on the left can obstruct multiple anchors depending on the tag's position on the actuator. The panel on the right obstructs the direct path to anchor A_1 at all times.

We perform the experiments for $N \in \{5, 7, 10\}$ anchors. For TDOA localization, we use $K = N - 1$ responses. For TWR, we use $K = N$ responses. For each algorithm and each (N, K) combination, we have approx. 14,000 location estimates in NLOS (7,000 at each position of the rail).

At the moment, we do not implement any NLOS mitigation procedure and we take into account all the measurements.

Table 6 shows the median and the 95th percentile (P_{95}) 3D localization errors of each method at the positions P1 and P4. Fig. 18 shows the distribution of 3D localization errors aggregated over both locations. FlexTDOA achieves lower errors than the classic TDOA in all NLOS scenarios. With five anchors, FlexTDOA has only a modest improvement of 5–7 cm in the median and P_{95} errors compared to the classic




Fig. 18: Distributions of localization errors in NLOS using TWR, classic TDOA (FI-FR), and FlexTDOA (CI-CR), aggregated over both NLOS scenarios (at location P1, with one obstruction, and at location P4, with two obstructions).

TDOA. However, the improvement is more evident for seven and ten anchors, where FlexTDOA reduces the P_{95} error by 19% and 38%, respectively, compared to the classic TDOA.

As expected, TWR has the smallest errors in all scenarios. This happens because, in TWR-based localization, an obstacle placed between an anchor and the tag introduces an error only in the distance between them. Therefore, the bias in the location estimate using TWR can be reduced if the rest of the distance measurements have small errors.

On the other hand, in TDOA localization, if the obstacle is between the initiating anchor and the tag, it incurs an error in *all* the TDOAs from that time slot. This is where FlexTDOA is more advantageous than the classic TDOA: by changing the initiating anchor, we ensure enough channel diversity to improve the robustness of the location estimate if the initiating anchor is obstructed.

Even though TWR-based localization provides the most accurate location estimates, it scales poorly with the increasing number of tags. For $N = 9$ anchors, the minimum duration of a time slot needed to obtain measurements between all anchors and the tag is approx. 10 ms (see Table 3). Therefore, with TWR, at most 100 tags could be localized per second. In fact, the number of tags would likely be smaller given that such a scheme would need a multiple access protocol to synchronize their access to the channel. TDOA-based localization, on the other hand, can scale to an unlimited number of tags. Using the same example, each tag could obtain 100 measurements per second independent of the number of tags in the area.

VII. CONCLUSION

In this paper, we presented a new flexible TDMA scheduling scheme for TDOA localization that fully exploits the channel diversity in the environment. We compared FlexTDOA, the proposed method, against the classic TDOA implementation with a fixed reference anchor and responder list and against range-based localization in a deployment of up to ten anchors and one tag in an office environment.

FlexTDOA achieves lower localization errors than the classic TDOA in most scenarios, with and without obstructions. In LOS, the improvement in the median accuracy brought by FlexTDOA compared to the classic TDOA is modest (2–3 cm) because the initiator in the classic TDOA already has a good link to the tag. However, the robustness brought by the increased diversity is evident in NLOS. In NLOS, FlexTDOA reduces 95th percentile of the localization error with up to 38% compared to the classic TDOA. Overall, FlexTDOA achieves a median localization error of 13–17 cm in LOS and 15–22 cm when one or more anchors are in NLOS with the tag (the error depends on how many anchors are used).

While TWR localization yields the highest accuracy among all methods, it has poor scalability with a growing number of anchors and responders. In contrast, FlexTDOA can scale to an unlimited number of tags.

In the future, we will scale up the proposed system to a multi-room or building environment. This will require several issues to be addressed: pairs of anchors that are not in communication range, system calibration (self-localization) for the sparsely connected network, and an efficient TDMA scheme which reuses slots for out-of-range nodes.

REFERENCES

- [1] FiRa™, “Unleashing the potential of UWB: Regulatory considerations,” Aug. 2022.
- [2] F. Zafari, A. Gkelias, and K. K. Leung, “A survey of indoor localization systems and technologies,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2568–2599, 2019.
- [3] M. Ridolfi, A. Kaya, R. Berkvens, M. Weyn, W. Joseph, and E. D. Poorter, “Self-calibration and collaborative localization for UWB positioning systems: a survey and future research directions,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 4, pp. 1–27, 2021.
- [4] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Z. Win, “Ranging with ultrawide bandwidth signals in multipath environments,” *Proceedings of the IEEE*, vol. 97, no. 2, pp. 404–426, 2009.
- [5] F. Gustafsson and F. Gunnarsson, “Positioning using time-difference of arrival measurements,” in 2003 *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2003. *Proceedings (ICASSP’03)*, vol. 6, pp. VI–553, IEEE, 2003.
- [6] A. Urruela, J. Sala, and J. Riba, “Average performance analysis of circular and hyperbolic geolocation,” *IEEE Transactions on Vehicular Technology*, vol. 55, no. 1, pp. 52–66, 2006.
- [7] L. Santoro, M. Nardello, D. Brunelli, and D. Fontanelli, “Scale up to infinity: The UWB indoor global positioning system,” in 2021 *IEEE International Symposium on Robotic and Sensors Environments (ROSE)*, pp. 1–8, IEEE, 2021.
- [8] J. Tiemann, F. Eckermann, and C. Wietfeld, “Multi-user interference and wireless clock synchronization in TDOA-based UWB localization,” in 2016 *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–6, IEEE, 2016.
- [9] J. Friedrich, J. Tiemann, and C. Wietfeld, “Accurate multi-zone UWB TDOA localization utilizing cascaded wireless clock synchronization,” in 2021 *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, IEEE, 2021.
- [10] A. Ledergerber, M. Hamer, and R. D’Andrea, “A robot self-localization system using one-way ultra-wideband communication,” in 2015 *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3131–3137, IEEE, 2015.
- [11] M. Hamer and R. D’Andrea, “Self-calibrating ultra-wideband network supporting multi-robot localization,” *IEEE Access*, vol. 6, pp. 22292–22304, 2018.
- [12] J. Tiemann and C. Wietfeld, “Scalability, real-time capabilities, and energy efficiency in ultra-wideband localization,” *IEEE Transactions on Industrial Informatics*, vol. 15, no. 12, pp. 6313–6321, 2019.
- [13] Qorvo, “Dwm3000 6.5 & 8.0 ghz ultra-wideband (uwb) module.” <https://www.qorvo.com/products/p/DWM3000>, Accessed on: Nov. 15, 2022.
- [14] M. Pelka and H. Hellbrück, “S-TDoA—sequential time difference of arrival—A scalable and synchronization free approach for positioning,” in 2016 *IEEE Wireless Communications and Networking Conference*, pp. 1–6, IEEE, 2016.
- [15] 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 2019 *18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pp. 61–72, IEEE, 2019.
- [16] P. Corbalán, G. P. Picco, and S. Palipana, “Chorus: UWB concurrent transmissions for GPS-like passive localization of countless targets,” in 2019 *18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pp. 133–144, IEEE, 2019.
- [17] J. Tiemann, F. Eckermann, and C. Wietfeld, “Atlas—an open-source tdoa-based ultra-wideband localization system,” in 2016 *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–6, IEEE, 2016.
- [18] J. Tiemann and C. Wietfeld, “Scalable and precise multi-UAV indoor navigation using TDOA-based UWB localization,” in 2017 *international conference on indoor positioning and indoor navigation (IPIN)*, pp. 1–7, IEEE, 2017.
- [19] J. Tiemann, Y. Elmasry, L. Koring, and C. Wietfeld, “ATLAS FaST: Fast and simple scheduled TDOA for reliable ultra-wideband localization,” in 2019 *International Conference on Robotics and Automation (ICRA)*, pp. 2554–2560, IEEE, 2019.
- [20] J. Yang, B. Dong, and J. Wang, “VULoc: Accurate UWB localization for countless targets without synchronization,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 3, pp. 1–25, 2022.
- [21] D. Vecchia, P. Corbalán, T. Istomin, and G. P. Picco, “TALLA: Large-scale TDoA localization with ultra-wideband radios,” in 2019 *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, IEEE, 2019.
- [22] W. Zhao, A. Goudar, and A. P. Schoellig, “Finding the right place: Sensor placement for UWB time difference of arrival localization in cluttered indoor environments,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6075–6082, 2022.
- [23] A. Prorok, P. Tomé, and A. Martinoli, “Accommodation of NLOS for ultra-wideband TDOA localization in single-and multi-robot systems,” in 2011 *international conference on indoor positioning and indoor navigation*, pp. 1–9, IEEE, 2011.
- [24] B. Van Herbruggen, J. Fontaine, and E. De Poorter, “Anchor pair selection for error correction in time difference of arrival (TDoA) ultra wideband (UWB) positioning systems,” in 2021 *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, IEEE, 2021.
- [25] B. Teague, Z. Liu, F. Meyer, A. Conti, and M. Z. Win, “Network localization and navigation with scalable inference and efficient operation,” *IEEE Transactions on Mobile Computing*, vol. 21, no. 6, pp. 2072–2087, 2020.
- [26] I. Dotlic, A. Connell, and M. McLaughlin, “Ranging methods utilizing carrier frequency offset estimation,” in 2018 *15th Workshop on Positioning, Navigation and Communications (WPNC)*, pp. 1–6, IEEE, 2018.
- [27] J. Sidorenko, V. Schatz, N. Scherer-Negenborn, M. Arens, and U. Hugentobler, “Error corrections for ultrawideband ranging,” *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 11, pp. 9037–9047, 2020.
- [28] Z. Ebadi, C. Hannotier, H. Steendam, F. Horlin, and F. Quitin, “An over-the-air CFO-assisted synchronization algorithm for TDOA-based localization systems,” in 2020 *IEEE 92nd Vehicular Technology Conference (VTC2020-Fall)*, pp. 1–5, IEEE, 2020.
- [29] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu, “Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks,” *IEEE signal processing magazine*, vol. 22, no. 4, pp. 70–84, 2005.
- [30] 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*, vol. 48, no. 4, pp. 679–689, 2000.
- [31] 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)*, pp. 1–4, IEEE, 2016.

[32] J. Tiemann, J. Friedrich, and C. Wietfeld, "Experimental evaluation of iee 802.15.4z ubw ranging performance under interference," *Sensors*, vol. 22, no. 4, p. 1643, 2022.

[33] J. Smith and J. Abel, "Closed-form least-squares source location estimation from range-difference measurements," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 35, no. 12, pp. 1661–1669, 1987.

[34] "DW1000 device driver application programming interface (API) guide: Using API functions to configure and program the DW1000 transceiver. Version 2.7." Decawave Ltd. 2021.

[35] "DW3xxx device driver application programming interface (API) guide: Using API functions to configure and program the DW3000 and QM33120 UWB transceiver. Version 2.2." Decawave Ltd. 2021.

[36] F. Gustafsson and F. Gunnarsson, "Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements," *IEEE Signal processing magazine*, vol. 22, no. 4, pp. 41–53, 2005.

[37] "IEEE standard for low-rate wireless networks," *IEEE Std 802.15.4-2020 (Revision of IEEE Std 802.15.4-2015)*, pp. 1–800, 2020.

[38] STMicroelectronics, "Stm32f429zi - high-performance advanced line, arm cortex-m4 core with dsp and fpu, 2 mbytes of flash memory, 180 mhz cpu, art accelerator, chrom-artaccelerator, fmc with sdram, tft." <https://www.st.com/en/microcontrollers-microprocessors/stm32f429zi.html>, Accessed on: Nov. 15, 2022.

[39] J. T. Isaacs, D. J. Klein, and J. P. Hespanha, "Optimal sensor placement for time difference of arrival localization," in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pp. 7878–7884, IEEE, 2009.



IULIU VASILESCU originary from Tulcea, Romania obtained his PhD in computer science and robotics from Massachusetts Institute of Technology (MIT) in Boston, MA in 2009. He worked on underwater autonomous robots, underwater optical communications, underwater imaging and underwater acoustic communication and localization. He pioneered the use of underwater robots in synergy with underwater sensor networks.



DRAGOŞ NICULESCU obtained a PhD in Computer Science from Rutgers University (New Jersey) in 2004, with a thesis on sensor networks routing and positioning. He spent five years as a researcher at NEC Laboratories America in Princeton, NJ, working on simulation and implementation of mesh networks, VoIP, and WiFi-related protocols. At University Politehnica of Bucharest he is currently teaching courses in Mobile Computing and Services for Mobile Networking; also researching mobile protocols, UWB, and 802.11 networking.



GEORGE-CRISTIAN PĂTRU is currently a PhD candidate at the Automatic Control and Computer Science Faculty, University Politehnica of Bucharest, with a focus on IIoT and Robotics. Cristian has gained experience over the last four years in hardware development while working with several hardware-centric start-ups; in his role as technical mentor, he helped them to quickly build and scale live prototypes. Such prototypes include LoRaWAN based IIoT sensors, custom-built sensors with high autonomy, IIoT-based monitoring in smart buildings, specialized wearable devices, autonomous drones and terrestrial robots.

built sensors with high autonomy, IIoT-based monitoring in smart buildings, specialized wearable devices, autonomous drones and terrestrial robots.



DANIEL ROSNER is an Associate Professor in the Automatic Control and Computer Science Faculty, University Politehnica of Bucharest, and has earned his PhD in Computer Science and Engineering, on Assisted living technologies for ubiquitous health, with a focus on IIoT devices. He has ample experience in developing IIoT and IIoT (Industrial Internet of Things) hardware and software systems while working with start-ups aiming to develop their prototypes or to validate their technology assumptions. His latest research projects featured IIoT infrastructure deployment solutions, energy monitoring for smart homes and Embedded Software Solutions for automotive grade microcontrollers.



LAURA FLUORATORU obtained her MSc in Electrical Engineering from ETH Zürich, Switzerland, in 2019. She is currently pursuing a double PhD degree at University Politehnica of Bucharest, Romania, and at Tampere University, Finland, as a Marie Skłodowska-Curie Fellow in the A-WEAR project. During her studies, she gained experience in both industry and research from internships at EPF Lausanne, Schindler Group, and NXP Semiconductors. Her research interests include indoor localization, ultra-wideband technology, wireless and mobile communications, embedded systems, and machine learning.

interests include indoor localization, ultra-wideband technology, wireless and mobile communications, embedded systems, and machine learning.

