



HAL
open science

Unlocking wireless sensing potential in Wi-Fi and IoT networks

Mohamed Naoufal Mahfoudi

► **To cite this version:**

Mohamed Naoufal Mahfoudi. Unlocking wireless sensing potential in Wi-Fi and IoT networks. Networking and Internet Architecture [cs.NI]. COMUE Université Côte d'Azur (2015 - 2019), 2019. English. NNT : 2019AZUR4063 . tel-02431424v2

HAL Id: tel-02431424

<https://hal.science/tel-02431424v2>

Submitted on 14 May 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT

Libérer le potentiel de détection sans fil dans les réseaux Wi-Fi et IoT

Mohamed Naoufal MAHFOUDI

Inria

Présentée en vue de l'obtention
du grade de Docteur en Informatique
d'Université Côte d'Azur
Dirigée par : Walid DABBOUS
/ Robert STARAJ
Soutenue le : 1^{er} Octobre 2019

Devant le jury composé de:
Andrzej DUDA Professor, Grenoble INP-Ensimag
Thomas NOËL, Professor, University of Strasbourg
Philippe MARTINS, Professor, Télécom ParisTech
Guillaume URVOY-KELLER, Professor, I3S/UCA
Walid DABBOUS, DR, Inria Sophia Antipolis
Robert STARAJ, Professor, LEAT/UCA

Libérer le potentiel de détection sans fil dans les réseaux Wi-Fi et IoT

Jury:

Président du jury :

Guillaume URVOY-KELLER Professor I3S/UCA

Rapporteurs:

Andrzej DUDA Professor Grenoble INP-Ensimag
Thomas NOËL Professor University of Strasbourg

Examineurs:

Philippe MARTINS Professor Télécom ParisTech
Guillaume URVOY-KELLER Professor I3S/UCA

Directeurs:

Walid DABBOUS DR Inria Sophia Antipolis
Robert STARAJ Professor LEAT/UCA

Libérer le potentiel de détection sans fil dans les réseaux Wi-Fi et IoT

Résumé: La détection sans fil a évolué depuis la découverte de la détection radar en 1886. L'analyse des réflexions électromagnétiques d'objets a ouvert la voie à un large éventail d'applications allant de la localisation de cibles à longue distance pour la navigation civile et militaire à la surveillance du vent et des précipitations, les prévisions météo à la détection de vitesse pour la sécurité routière. Cependant, pendant très longtemps, la détection sans fil a rarement été utilisée pour des applications centrées sur l'homme en raison de limitations techniques, d'impraticabilité ou de coût. L'introduction des réseaux sans fil a suscité un nouvel intérêt pour le développement de nouveaux services de détection sans fil en raison de leur souplesse et de leur polyvalence. L'intégration de ces fonctionnalités contribuerait à résoudre certains problèmes de société importants. La localisation, la détection de mouvements et la surveillance des signes vitaux ont un grand potentiel pour promouvoir le vieillissement en bonne santé, la sécurité publique et le commerce. La détection sans contact offre un degré de liberté appréciable, permettant de surveiller à distance les personnes âgées isolées sans entraver leur vie quotidienne. Elle pourrait aider les services de sécurité publique à dénombrer les foules et à détecter les survivants à l'intérieur des bâtiments en cas d'urgence. Les commerces de détail et les établissements publics tireraient parti d'une localisation active et passive pour offrir une expérience améliorée à leurs visiteurs et faciliter leurs efforts logistiques.

Cette thèse aborde le problème de l'exploitation des réseaux sans fil commerciaux pour les applications de détection :

L'un des défis de la surveillance sans fil consiste à détecter l'orientation d'une personne avec précision. Tandis que d'autres travaux fournissent des solutions à granularité grossière pour résoudre de tels problèmes, nous utilisons les techniques de radar MIMO pour fournir un système d'estimation d'orientation précis pour les infrastructures Wi-Fi. Pour être plus précis, nous analysons les informations de phase des signaux reçus sur le réseau d'antennes afin de calculer le cap d'un terminal Wi-Fi.

Un deuxième défi consiste à fournir un système de positionnement précis aux systèmes LPWAN afin de maintenir la cohérence des informations des capteurs déployés. Les solutions actuelles sont complexes, coûteuses ou consomment beaucoup d'énergie. Pour résoudre ce problème, nous introduisons les fonctions MIMO dans les systèmes LoRa LPWAN afin de permettre une localisation précise avec des coûts de démarrage limités. Nous activons l'estimation de l'angle d'arrivée en utilisant une deuxième antenne sur la passerelle LoRaWAN. Nous prouvons également l'utilité de ces informations pour augmenter l'efficacité des communications sans fil.

Un troisième défi pour la localisation sans fil est l'inefficacité des approches actuelles basées sur un modèle en cas de conditions de non-visibilité et la rigidité des approches basées sur les données en cas de changements d'environnement de propagation. Pour relever ce défi, nous proposons une nouvelle solution de localisation passive pilotée par les données afin de remédier aux limitations des techniques de localisation basées sur un modèle.

Pour donner vie à de tels systèmes et leur donner une chance d'impact sur notre quotidien, nous devons promouvoir la réutilisabilité et la reproductibilité. Pour cela,

nous essayons de relever le défi de la reproductibilité dans les réseaux sans fil en analysant l'état actuel, en réalisant une étude de cas et en présentant les enseignements qui en découlent.

Mots-clés: Localisation active, Localisation passive, Sans fil, Wi-Fi, OFDM, Information d'état du canal (CSI), Orientation, LoRaWAN, LoRa, Filtrage spatial, Combinaison cohérente, Apprentissage profond, Adaptation de domaine, Reproductibilité

Unlocking Wireless Sensing Potential in Wi-Fi and IoT Networks

Abstract: Wireless sensing has evolved since the discovery of radio wave echo detection and radar in 1886. Analyzing electromagnetic reflections from objects opened the way for a wide range of applications spanning from locating long-range targets for navigation and military to monitoring wind and precipitation for weather-forecasting to velocity detection for public safety. However, for the longest time, its usefulness was seldom for human-centric applications because of technical limitations, impracticality or costliness. Introducing wireless networks awakened a newfound interest in developing new wireless sensing services for their seamlessness and versatility. Integrating such functionalities would contribute to resolving some prominent societal issues.

Localization, motion detection, and vital signs monitoring have great potential for promoting healthy aging, public safety, and retail. Contactless sensing offers an appreciable degree of freedom, enabling remote monitoring of the isolated elderly without hampering their daily lives. It could assist public safety services for crowd counting and detection of survivors inside buildings during emergencies. Retail and public facilities would benefit from passive and active localization to offer an enhanced experience to their visitors and to help their logistical efforts.

This thesis addresses the problem of leveraging commercial off-the-shelf wireless networks for sensing applications:

One challenge for wireless monitoring is to detect the attitude of a person accurately. While other works provide coarse-grained solutions for resolving such issues, we use MIMO radar techniques to provide an accurate orientation estimation system for Wi-Fi infrastructures. To be more precise, we analyze the phase information of signals received on the antenna array to compute the heading of a Wi-Fi terminal.

A second challenge is to provide an accurate positioning system for LPWAN systems to maintain the information consistency of deployed sensors. Current solutions are complex, costly, or not energy-efficient. To address this problem, we introduce MIMO capabilities to LoRa LPWAN systems that provide accurate localization with limited startup costs. We enable the angle of arrival estimation by leveraging a second antenna on the LoRaWAN gateway. We also prove the usefulness of such information for wireless communication efficiency.

A third challenge for wireless localization is the inefficiency of current model-based approaches in case of non-line-of-sight conditions and the rigidity of data-driven approaches in case of propagation environment changes. To address this challenge, we propose a new data-driven solution for passive localization to address the limitations of model-based localization techniques.

To give life to such systems and provide them with a chance of impacting our everyday lives, we should promote reusability and reproducibility. For that, we focus on the challenge of reproducibility in wireless networking by surveying the current state, performing a case study, and presenting the engendered lessons.

Keywords: Active Localization, Passive Localization, Wireless, WiFi, OFDM, Channel state information (CSI), Orientation, LoRaWAN, LoRa, Spatial Filtering, Coherent Combining, Deep learning, Data-driven, Domain adaptation, Reproducibility

Acknowledgements

First and foremost I want to thank my advisors Walid Dabbous, Robert Staraj, and Thierry Turetli. It has been an honor to be their PhD student. They made me discover how amazing research can be. I appreciate their help and funding to make my Ph.D. fruitful and stimulating. Even during difficult times Walid and Thierry were present, understanding and supportive. I would like to also thank Thierry Parmentelat, Fabien Ferrero and Leonardo Lizzi for their involvement in my research journey.

I am also immensely grateful for being part of Diana. It has been at the origin of friendship and helpful advice. I had the privilege of calling Damien Saucez a friend, he contributed tremendously to widening my research perspective. I am grateful for meeting and working with amazing people: Ramon Dos Reis Fontes, Mario Zancanaro, Farzaneh Pakzad, Yassir Mrabet, Gayatri Sivadoss, Othmane Bensouda Korachi, Anas Errahali, Ghada Moualla, Vitalii Poliakov, Karyna Gogunska, Abdelhakim Akodadi, Yanis Boussad, and Muhammad Jawad Khokhar. I would like to also thank Chadi Barakat for his good advice and his thoughtfulness, and Arnaud Legout for his helpfulness and support.

I am also grateful to Xinyu Zhang (Professor at UCSD) for hosting me in his team for 6 months for conducting research on data-driven localization discussed in this dissertation. During my stay at UCSD, I got to expand my perspective of how research can contribute to humanity, and what it takes to be part of the research journey. I had some amazing research partners who became my friends during this visit: Changhan Ge and Yifan Huang. This visit would not have been possible without the support of my advisors and their help.

The reproducibility studies discussed in this dissertation would not have been possible without Damien Saucez, Olivier Bonaventure, Luigi Iannone, Matthias Flittner, Matthias Wählisch, Vaibhav Bajpai, Alex Afanasyev and all the other participants of the Sigcomm Reproducibility workshop 2017. As for the LoRA contributions, Gayatri Sivadoss, Othmane Bensouda Korachi, Anas Errahali, and Antonello Florio played a key role in making this project a reality.

My time at Inria was enjoyable thanks to my friends (old and New), those I discovered during my masters and PhD: Dimitra Politaki, Othmane Belmoukadam, Thierry Spettebrouet, Alina Tugalukova. A special thank you to Qariani Ayoub my once-in-a-lifetime friend.

Lastly, I would like to thank my family for all their love and their unwavering support. For my parents who taught me the value of compassion and resilience and supported me in all my pursuits. For my sisters, Lamya and Sara, who were very supportive and source of inspiration.

Thank you.

Contents

Abstract	iii
Acknowledgements	v
1 Introduction	1
1.1 Motivation	1
1.2 Contributions	2
2 Estimating Orientation with COTS Wi-Fi Devices	5
2.1 Introduction	5
2.2 MIMO-CSI Primer	6
2.2.1 Signal Phase and emitter’s position estimation	7
2.3 The ORION system	9
2.3.1 A look into the system model	9
2.3.2 COTS Wi-Fi and AoD Estimation	11
2.3.3 Phase correction for AoA/ AoD estimation	12
2.3.4 Channel correlation and Spatial Diversity	15
2.3.5 Estimating orientation	16
2.4 Evaluation	17
2.4.1 Implementation	17
2.4.2 Experimental Results	20
2.5 Related Work	22
2.6 Conclusion	22
3 Joint localization and range extension for LPWAN	23
3.1 Introduction	23
3.2 LoRa primer	24
3.2.1 Physical Frame Format	25
3.3 LP-WAN MIMO Gateway for Range extension	26
3.3.1 Coherent Combining for Uplink Communication	27
3.3.2 Beamforming for Downlink Communication	27
3.4 LP-WAN MIMO Gateway for AoA-based localization	29
3.5 SDR implementation	31
3.6 System Evaluation	32
3.6.1 Position estimation precision	33
Evaluation of the positioning system in Indoor Environments.	33
Evaluation of the positioning system in Outdoor Environments.	34
3.6.2 Coherent Combining	36
Coherent Combining in Outdoor Environments	37

	Coherent Combining in Indoor Environments	37
3.6.3	Beamforming	38
	Beamforming in Indoor Environments	38
	Beamforming in Outdoor Environments	38
3.7	Related Works	40
3.8	Conclusions	41
4	Robust Passive Localization Using Wi-Fi CSI: A Data-Driven Paradigm	43
4.1	Introduction	43
4.2	Related Work	46
4.3	The RFLoc Dataset	47
4.4	Data-Driven Passive Localization in NLoS and Heterogeneous Environments	49
	4.4.1 Position Inference module	49
	4.4.2 Environment Adaptation Module	53
	4.4.3 Implementation and End-to-End Training	56
4.5	DPLoc Evaluation	57
	4.5.1 Effectiveness of the Inference Module	57
	4.5.2 Effectiveness of Environment Adaptation	59
4.6	Conclusion	61
5	Reproducibility for Wireless Networking Experiments	63
5.1	Introduction	63
5.2	Why is Reproducibility is so hard to achieve?	64
5.3	Survey of Reproducibility in the Community	65
	5.3.1 Method	65
	5.3.2 Results	65
	5.3.3 Discussion	68
5.4	Reproducibility for Wireless Experiments:	69
	5.4.1 Reproducing a wireless experiment: A case study	69
	OpenRF	70
	Reproducibility challenges	70
	Observation and Discussion	71
	5.4.2 Towards Reproducible Wireless Experiments Using R2lab	73
	Experimental Methodology	73
	R2lab, testbed for reproducible wireless experiments	73
	Experimental material in a <i>git</i> repository	78
	5.4.3 A step further towards runnable papers using R2lab:	79
5.5	Thoughts On Improving Reproducibility	82
	5.5.1 Author incentives	82
	5.5.2 Reviewer incentives	83
	5.5.3 Review Form Design	83
	5.5.4 Sharing Artifacts	84
5.6	Conclusion	84

Appendices	85
5.A OpenRF reproducibility: Detailed verification process	85
5.B Artifact Evaluation Review Form	86
6 Conclusion	91
6.1 Future Work	92

List of Figures

2.1	Packet in case of two spatial streams[102].	6
2.2	Signal impinging on an M -element uniform linear antenna array: Each antenna receives a copy of the signal with a phase shift $(k - 1)dsin(\theta)$ relatively to the reference element (First antenna), with d , the antenna spacing, in the order of half a wavelength.	8
2.3	MIMO system with N -element send and M -element receive coplanar ULA with omnidirectional antennas.	9
2.4	Signal radiation from an N -element uniform linear antenna array: The k -th antenna transmits a signal with a phase shift of $e^{-j2\pi(k-1)\frac{d}{\lambda}\sin(\phi)}$ relatively to the reference antenna.	12
2.5	Partial view of the transceiver Block with a focus on the phase ambiguity generated by the signal synthesizers.	14
2.6	Partial view of the MIMO Transmission Block as described and specified in the IEEE 802.11n standard.	15
2.7	Signal phase in radian of each spatial stream as seen by the first receiving antenna with and without phase correction for CSD.	16
2.8	ORION System Design for orientation Estimation.	17
2.9	Inria Sophia Antipolis' faradized anechoic chamber (R2lab) with 37 Wi-Fi nodes equipped with 3-elements ULA and Intel 5300 cards. The antennas are placed under the nodes (blue boxes) to have clear line-of-sight (LoS) communication between nodes.	18
2.10	Omni-directional antenna array rotated through a computer controlled module (red circle).	18
2.11	Experimental cumulative density function of orientation error according to absolute reference.	19
2.12	Box-plot of Orientation Error using SAGE and MUSIC.	19
2.13	Box-plot of Orientation Error using SAGE and MUSIC w/ outliers.	20
2.14	Office room where the experimentation is conducted.	20
2.15	Experimental cumulative density function of orientation error for long distances according to absolute reference.	21
3.1	LoRa Gateways each equipped with a uniform linear array apply spatial filtering by adopting a more focused radiation pattern (beamforming) for Transmitting Down-link packets to clients instead of the usual omnidirectional radiation pattern.	24
3.2	LoRA demodulation process.	25
3.3	Spectrum showing the different SF, where the higher the SF the longer the time-on-air.	25

3.4	LoRA Physical frame format composed of sequential Chirp Signals. . .	26
3.5	Upon receiving a signal at the LoRa gateway's antenna array, the different signal copies are combined constructively according to their signal-to-noise ratio.	28
3.6	Distorted Chirps in the LoRa signals impinging on the antenna array are combined by aligning, scaling and summing the different copies producing a signal with a higher SNR.	28
3.7	LoRa signal beamforming by applying complex weights to each RF chain. This allows combining the signals constructively towards the direction of the target.	29
3.8	Signal impinging on an N-element uniform linear antenna array: Each antenna receives a copy of the signal with a phase shift $k.d.\sin(\theta)$ relatively to the reference element (First antenna), with d , the antenna spacing, in the order of half a wavelength.	30
3.9	FFT for all SFs	33
3.10	USRP b210 equipped with a 2-element Uniform Linear Array for estimating the AoA and spatial filtering.	33
3.11	Experiment testbed showing the end-device locations (blue circles) and the gateway locations (red squares). The testbed is designed to evaluate Snipe in a wide variety of deployment scenarios indoors. Snipe has been tested when both the gateway and the end-device are along a corridor, office and meeting room.	34
3.12	ECDF of the AoA estimation error in different environments.	35
3.13	Experiment testbed showing the end-device locations (blue circles) and the gateway locations (red squares). The testbed is designed to evaluate Snipe in outdoor LoS condition with large building, cars and vegetation.	35
3.14	Outdoor Experiment testbed deployed on Inria campus showing the end-device locations (blue circles) and the gateway locations (red squares). The testbed is designed to evaluate Snipe in an outdoor environment rich with multipath and in NLoS propagation conditions.	37
3.15	Empirical CDF power gain using coherent combining in a meeting room and corridor.	38
3.16	Empirical CDF power gain when using beamforming in a meeting room and corridor.	39
4.1	(a) LoS and (b) NLoS in passive localization.	44
4.2	The DPLoc dataset is built using Quantenna Wi-Fi radio for CSI collection and the Kinect v2 for location labeling which provides both depth maps and hip coordinate as location label for each CSI sample.	48
4.3	Floor plans of experiment environments.	49
4.4	Network Design of the Inference Module	50
4.5	Network design for environment adaptation.	53
4.6	NLoS setup between a Tx and a subject.	57
4.7	Yoga room environment with obstacles.	57
4.8	Location error CDF when changing the layout in the same environment (yoga room).	58

4.9	Location error CDF when using adaptation in a different LoS environments.	58
4.10	Estimation error with and without retraining on generated corridor data.	59
4.11	Location error CDF using the retrained inference w/ and w/o adaptation generator.	60
4.12	Location error CDF when enabling the environment adaptation module.	60
4.13	Location error of Widar 2.0.	61
5.1	Beamforming System Design	72
5.3	Wi-Fi dual-band antennas for a node	75
5.4	Pseudospectrum for AoA (left) and AoD (right) before (blue) and after (red) phase correction (from [93]).	76
5.5	A glimpse at the Jupyter notebook for R2lab radiomap click the images to run on mybinder.org	80

List of Tables

3.1	The estimated Angle of arrivals and the incurred errors according to the distance from the gateway and the elevations for LoS Experiments.	36
3.2	The estimated Angles of arrival and the errors according to the distance from the gateway for NLoS Experiments.	36
3.3	The power gain using coherent combining according to the distance from the gateway and the elevations for LoS Experiments.	37
3.4	The power gain using beamforming according to the distance from the gateway and the elevations for LoS Experiments.	39
3.5	The power gain using beamforming according to the distance from the gateway for NLoS Experiments.	40
4.1	Structure of the RFLoc dataset.	51
4.2	Convolutional layers used in our location inference model.	51
4.3	Split and estimation error on same validation set.	58
4.4	Estimation error in different environments in NLoS.	59
4.5	Training time delay with different splitting methods.	59
4.6	Estimation error when changing the layout in the same general environment (yoga room).	59
5.1	Summary of artifact survey, compared to the overall number of published papers per conference.	66
5.2	Summary of artifact nature. Please note: Some artifacts are counted in multiple rows (if applicable). But only once per column.	67
5.3	Available SDR devices	75
5.4	Controllable parameters from the notebook	81

List of Acronyms

ACM	Association for Computing Machinery
AI	Artificial Intelligence
AoA	Angle of Arrival
AoD	Angle of Departure
AP	Access Point
BW	Band width
CCI	Co-Channel Interference
CFO	Carrier Frequency Offset
CNN	Convolutional Neural Network
COTS	Commercial Off-The-Shelf
CRC	Cyclic Redundancy Check
CSD	Cyclic Shift Diversity
CSI	Channel State Information
DFT	Discrete Fourier Transform
DSSS	Direct-Sequence Spread Spectrum
EM	Expectation- Maximization
FFT	Fast Fourier Transform
FHSS	Frequency- Hopping Spread Spectrum
FMCW	Frequency-Modulated Continuous-Wave
GAN	Generative Adversarial Network
GPS	Global Positioning System
HT	High Throughput
IBSS	Independent Basic Service Sets
MIMO	Multiple-Input Multiple-Output
IMU	Inertial Measurement Unit
IoT	Internet of Things
IQ	In-phase and Quadrature
ISM	Industrial, Scientific, and Medical
LoS	Line-of-Sight
LPWAN	Low-Power Wide-Area Network
LS	Least Square
LSTM	Long Short-Term Memory
LTF	Long Training Field
LTS	Long Term Support
MAC	Media Access Control
MAP	Maximum A-Posteriori
MLE	Maximum Likelihood Estimation
MMSE	Minimum Mean Square Error

MUSIC	M Ultiple S ignal Classification
NIC	N etwork Interface Controller
NLoS	N on- L ine- o f- S ight
OFDM	O rthogonal Frequency- D ivision M ultiplexing
PLL	P hase- L ocked L oop
PSD	P ower S pectral D ensity
RF	R adio F requency
RSS	R eceived S ignal S trength
RSS	R eceived S ignal S trength I ndication
SAGE	S pace A lternating G eneralized E xpectation- M aximization
SF	S preading F actor
SFO	S ampling F requency O ffset
SLAM	S imultaneous L ocalization A nd M apping
SM	S patial M ultiplexing
SNR	S ignal- t o- N oise R atio
STS	S pace- T ime S treams
TDoA	T ime D ifference o f A rrival
ToF	T ime- o f- F light
ULA	U niform L inear A rrays
USRP	U niversal S oftware R adio P eripheral
VR	V irtual R eality

*Dedicated to my beloved parents,
Abaacrouche Mina and Mahfoudi Mustapha.*

Chapter 1

Introduction

1.1 Motivation

Our society is morphing with the new industrial revolution. Prompted by the emergence of artificial intelligence (AI), big data, and the internet of things (IoT), this shift aims at enhancing the individual ability and raising the standard of living. This newly defined human-technology relationship imposes a more in-depth understanding of the individual and her needs. In this sense, remote sensing could provide seamless solutions transparent to the user, thus avoiding clutter and anxiety-inducing gadgets.

Remote sensing is not a new topic; It dates back to the 1840s, where some balloonists had been taking pictures of the ground using the freshly invented cameras. The principle is still the same where a sensor collects radiation from impinging electromagnetic waves from either direct or reflected sources. This information acquisition thrives by its non-invasive nature, making it an ideal choice for human-related applications. Given this attractive advantage, remote sensing has a significant role to play in this industrial revolution.

However, for the longest time, these tools were for niche markets and required specialized hardware. Fortunately, new wireless networks (5G, IoT networks, and Wi-Fi) which lead this modern technological era could take on advanced remote sensing applications on top of providing a robust communication medium for the billions of connected devices. Thanks to their versatility and omnipresence, they could open a broad spectrum of applications unexplored until a few years ago or only accessible through specialized hardware.

This generation of wireless networks could help address critical societal challenges. For instance, several countries are witnessing the worrisome aging of their population. This demographic change raises the question of elderly care with isolation being one of the most pressing matters; It increases seniors' health-related risks dramatically and supervising their wellbeing becomes even more challenging considering the limited human resources. Using these systems for elderly remote monitoring would ease the burden on the caregiver who can remotely and continuously assess the elder's health while preserving her autonomy, dignity, and privacy.

These solutions could also accompany retail professionals in their shift towards adopting the retail 4.0 model. Store owners are interested in studying the customers' behavior and provide her a custom-tailored experience. Engaging clients in shallow surveys on their level of satisfaction and interests is a hardly informative diagnosis tool. This new trend promises an in-depth behavioral examination of the shopper's

habits by tracking her position, path, orientation, and gait. Retailers could recreate the customer's mental patterns accurately, offer personalized services to each client, and in fine improve customer retention.

They could play an essential role in public safety. Crowd counting would assist law enforcement during large events by planning their resource allocation according to the attendees' distribution without infringing on their privacy. During disasters, these technologies would allow to detect and locate lives inside buildings, allowing a surgical intervention of rescue teams while avoiding unnecessary risks.

Finally, they will be the go-to solution for tracking the position of the billions of IoT devices, averting the need for dedicated personnel to establish and maintain accurate cartography. This solution is especially attractive for personal IoT networks where cost-effectiveness and energy efficiency are determining adoption factors.

For most of these uses cases, camera and computer vision are well-explored and efficient under some conditions. However, they are not applicable under occlusion, limited range, or privacy restrictions, whereas radio frequency (RF)-based passive and active localization become serious contenders.

Over the past two decades, the wireless sensing community has dedicated an extensive effort to investigating active and passive RF localization using radio devices [151]. Wi-Fi has always played a substantial role in driving novel sensing applications because of its omnipresence and cost-effectiveness. Therefore, Wi-Fi was hailed as the solution for indoor localization and wireless sensing for its potential seamless integration into the existing infrastructure. However, introducing IEEE 802.11n wireless cards supporting multiple-input multiple-output (MIMO) spurred an even greater interest, as researchers had access to a new fine-grained primitive, the channel state information (CSI). With this gained spatial diversity and CSI, it set us to unlock more sensing potential in commercial off-the-shelf (COTS) wireless networks.

In particular, given the context of MIMO enabled devices, are we able to propose accurate localization and orientation estimation systems?

1.2 Contributions

In this thesis, we seek to address some essential localization challenges for COTS wireless networks. (1) We enriched systems with original orientation estimation capabilities for Wi-Fi, which goes beyond coarse-grained solutions. (2) We enable MIMO capabilities for Low-power wide-area networks (LPWAN) to provide accurate localization using spatial diversity. (3) We introduce a new paradigm for passive localization to alleviate the limitations of model-based localization techniques. (4) We touch upon reproducibility issues in the wireless networking community to help draw a picture of the current state and promote the reusability of community provided tools. This work makes these specific contributions to RF-based sensing:

- **Orientation estimation** Integrating MIMO functionalities in W-Fi led the way to adopt antenna array techniques for fine-grained localization using commodity hardware. This approach has also spurred interest to reach beyond localization and now allows us to consider estimating the device's orientation in space, which once required other sources of information. Wi-Fi's popularity

and the availability of metrics related to channel propagation (CSI) make it a candidate readily available for experimentation. In chapter 2, we propose the ORION system to estimate the orientation (heading and yaw) of a MIMO Wi-Fi-equipped object, relying on a joint estimation of the angle of arrival and the angle of departure. Although several phase inconsistencies plague the CSI phase data, we show that an appropriate phase compensation strategy improves estimation accuracy.

- **Active Localization for LPWAN** Chapter 3 introduces Snipe, a novel system offering joint localization and range extensions for LPWANs. Although LPWAN systems such as long range (LoRa) achieve high communication range with low energy consumption, they suffer from fading in obstructed environments with dense multipath components, and their localization system is sub-par in terms of accuracy. In this work, we leverage MIMO techniques to achieve a higher signal-to-noise ratio at both the end device and the gateway while providing an opportunistic accurate radar-based system for localization with limited additional cost. The proposed system offloads the position estimation processing to the gateway and thus eases the energy cost incurred when using the Global Positioning System (GPS) devices. It eliminates the need for fine-grained synchronization between gateways, which is necessary for other localization techniques such as time difference of arrival (TDoA).
- **Robust Passive Localization** Passive localization technology can leverage RF reflection signals to identify a target's position. Existing passive localization solutions use closed-form models to resolve location based on signal path geometries. However, these models either rely on wideband radar or fail in non-line-of-sight (NLoS) conditions. In Chapter 4, we argue that a data-driven, deep learning-based approach may overcome such limitations, enabling robust and accurate passive localization when using low-profile Wi-Fi devices. We identify the significant challenges and establish RFLoc an open-source data infrastructure to encourage researchers to approach this vision. Based on the 1.1 million location-labeled Wi-Fi channel samples in RFLoc, we design a deep learning framework that uses a generative model to eliminate the cross-environment labeling overhead. Our experiments show that this framework can achieve sub-meter localization accuracy, even in unseen NLoS environments with multiple targets.
- **Reproducibility for Wireless Networking Experiments** Experimentation is an essential step for a realistic evaluation of wireless network protocols and testing and comparing the performance of wireless systems, like for any other scientific area, requires the ability to reproduce experimental results. The reproducibility of artifacts encourages enhanced reusability of tools, allowing a more efficient research project cycle. In Chapter 5, based on our experience in reproducing wireless sensing projects, we draw the principal aspects of an artifact crucial for easing the reproducibility. We describe to the current state of the reproducibility community by conducting a community survey and provide a recount of the steps we carried to promote artifact reusability in the

special interest group on data communication (SIGCOMM) organized by the association for computing machinery (ACM).

Chapter 2

Estimating Orientation with COTS Wi-Fi Devices

2.1 Introduction

With the popularity of Wi-Fi, indoor localization would be well-integrated into its existing infrastructure and therefore would avoid new dedicated systems. The shift to MIMO techniques through the adoption of the 802.11n standard further reinforced this idea. Fingerprinting [21] and other methods based on received signal strength indication (RSSI) measurements [24] cleared the way for more sophisticated localization approaches based on radar tracking techniques [16]. Hence, MIMO extended the set of localization techniques a single-input single-output (SISO) Wi-Fi infrastructure can propose. The infrastructure can now locate a terminal with decimeter level precision, either by using time-of-flight based ranging [12, 23], or angle of arrival (AoA) estimation techniques [20, 17]. With MIMO and orthogonal frequency-division multiplexing (OFDM) present on recent Wi-Fi chipsets, we are set to unlock more potential than simple localization, knowing that radar tracking techniques allow estimating, for instance, the orientation of a target [19, 22]. Although inertial measurement units (IMU), which consist of accelerometers, gyroscopes and also magnetometers, provide orientation information for devices, the gyroscope only provides the derivative of the yaw and magnetometers suffer from a perturbation in measuring the heading in indoor environments [10]. Therefore, an important question to answer is: Can we propose an accurate alternative to fingerprinting and IMU-based orientation estimation using radar tracking techniques on commodity Wi-Fi infrastructures? In this chapter, we intend to allow off-the-shelf MIMO Wi-Fi access points to estimate the orientation of a MIMO-enabled terminal using antenna array signal processing, by jointly estimating both the AoA and the angle of departure (AoD). The idea is to propose a deployable system in every Wi-Fi platform without modifying the equipment. The benefit would range from providing a reference for referencing IMU calibration, to indoor clients orientation tracking as they will all have the same heading reference unlike magnetometers measurements, and enabling simultaneous localization and mapping (SLAM) techniques [18], which require position and orientation measurements. Designing such a system in commodity Wi-Fi equipment implies tackling several design challenges:

1. CSI, i.e., the output matrix used for antenna array signal processing, suffers




FIGURE 2.1: Packet in case of two spatial streams[102].

from phase shifts. We should compensate these errors to achieve higher accuracy, and reproducibility of AoA estimation results.

2. Estimating a signal's AoD relies on the measurement of the phase difference between transmitted synchronized signals. The system should propose an approach for estimating the AoD by both allowing a synchronized transmission on multiple RF chains and being robust to phase inconsistencies associated with the chosen transmission scheme.
3. Estimating the orientation requires accurate and precise AoA and AoD values. Therefore, the system should be robust enough to estimation uncertainty.

Thus, our system design pays special attention to detecting and correcting phase inconsistencies, and to reducing measurement outliers. Hence, our contributions are along the same lines. First, we propose a phase correction technique for calibrating the system based on a single initial measurement reference, to provide a meaningful, accurate, and reproducible estimation of the AoA. Second, we present a novel mechanism for estimating the AoD of the signal by a Wi-Fi access point, using a common MIMO technique called spatial multiplexing (SM). Using SM introduces some phase inconsistencies we correct before launching the estimation. To the best of our knowledge, ORION[93] is the first system to propose the AoD estimation on commodity Wi-Fi devices. Finally, we propose an approach for estimating the orientation of the terminal. We apply a joint estimation of AoA and AoD to enhance estimation accuracy. ORION relies on signal processing done in a remote server to avoid hardware modifications of the access points and facilitate seamless adoption.

We organize the chapter as follows: we start by giving a brief primer on CSI and angle estimation in Section 2, then we lay down the system design in Section 2.3. We present the evaluation in Section 2.4, where we assess the system performance in an office room. Finally, we expose the related work in Section 2.5 and conclude in Section 2.6.

2.2 MIMO-CSI Primer

Wireless communication systems traded single antenna for a multi-antenna configuration for achieving higher spectral efficiency. However, exploiting antenna arrays for the purpose above requires detailed knowledge of the propagation channel, which we can acquire through the analysis of the response of the propagation environment to an emitted signal. More precisely, this channel response, called channel state information, is determined by the propagation channel (reflection, refraction,

scattering, fading, path loss). We characterize the multipath phenomenon present in the propagation environment by the following channel transfer function $H(f)$:

$$H(f) = \sum_{n=0}^{N-1} \rho_n e^{j\phi_n} e^{-j2\pi f \tau_n} \quad (2.1)$$

where τ_n represents the delay of the n -th path, N the total number of paths, f the frequency, and $\rho_n e^{j\phi_n}$ the complex gain of each path. Knowing we are exploiting a MIMO-OFDM system, we will have a value of $H(f)$ for each subcarrier corresponding to a receiving antenna and a transmitted signal. Interestingly, the channel estimation mechanism at the receiver can capture this multipath phenomenon (CSI) to a certain extent.

Generally, the CSI feedback from a receiver helps to achieve a more reliable communication by allowing a transmitter to target a receiver in particular, subsequently avoiding a spectrally inefficient omnidirectional transmission. By using a phased array, the transmitter shapes its radiation pattern [34] to reach more efficiently the intended receiver and avoid causing interference to other clients.

With 802.11n, we acquire the CSI by specifying particular pilot signals and training fields encoded in space (antennas) and time [105] [126], which helps to compute the CSI matrix for a MIMO system. As detailed in the IEEE 802.11n amendment [73], a specific training field is used within high throughput (HT) packets to estimate the CSI per receiving antenna and transmitted signal; this packet format was introduced for exploiting the advantages of MIMO in terms of spectral efficiency and thus achieving higher throughput and reliability. The training field is called high throughput long training field (HT-LTF), whose symbols are encoded in time and space for each one of the subcarriers, knowing that MIMO is usually used with OFDM in Wi-Fi systems. In the IEEE 802.11n amendment, two types of HT packets (Mixed mode and greenfield) are specified, and both are holding the HT-LTF training field. Figure 2.1 showcases the greenfield packets for two signal streams. As we mentioned earlier, the pattern used in HT-LTF is encoded over time and space and is known by the receiver. The number of symbols is determined by the number of subcarriers used throughout the communication. For example, with 20 MHz bandwidth, there will be 56 symbols, or with 40 MHz we will get 114 symbols. The transmitted symbols are distributed across subcarriers; this channel estimation implements block type pilot [44]. The number of transmitted HT-LTF fields depends on the number of transmitted streams, as shown in Figure 2.1. To estimate the CSI matrix corresponding to the transmitted streams and receiving antenna, the receiver can apply for instance a least square (LS) or a minimum mean square error (MMSE) or a maximum a-posteriori (MAP) estimator.

2.2.1 Signal Phase and emitter's position estimation

Multiple antennas on Wi-Fi devices not only offer better link reliability but also allows the estimation of the origin of a signal through measuring the signal's phase and amplitude from different positions (spatially separated antennas). The measurement of the AoA of a signal, which equates estimating its point of origin, is based on the observation that the signal of interest impinging on the receiver's antennas is




FIGURE 2.2: Signal impinging on an M -element uniform linear antenna array: Each antenna receives a copy of the signal with a phase shift $(k - 1)d\sin(\theta)$ relatively to the reference element (First antenna), with d , the antenna spacing, in the order of half a wavelength.

detected at slightly different time delays, see Figure 3.8. Knowing that the measured baseband signal is represented in the frequency domain, we use the relative phase between the signals received at the antenna array to estimate the AoA. This phase shift is defined by the distance between the antenna elements and the AoA of the signal. This distance is kept in the order of half of the wavelength λ .

The received signal and the relative phase can be expressed by the following model:

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_M \end{pmatrix} = \begin{pmatrix} a_1(\theta_0) & a_1(\theta_1) & \cdots & a_1(\theta_{L-1}) \\ a_2(\theta_0) & a_2(\theta_1) & \cdots & a_2(\theta_{L-1}) \\ \vdots & \vdots & \ddots & \vdots \\ a_M(\theta_0) & a_M(\theta_1) & \cdots & a_M(\theta_{L-1}) \end{pmatrix} \begin{pmatrix} s_0 \\ s_1 \\ \vdots \\ s_{L-1} \end{pmatrix} + \begin{pmatrix} n_1 \\ n_2 \\ \vdots \\ n_M \end{pmatrix} \quad (2.2)$$

where $a_k(\theta) = e^{-j2\pi(k-1)\frac{d}{\lambda}\sin(\theta)}$ represents the phase shift between the signals received by the reference and the k^{th} antenna, d the half-wavelength inter-antenna spacing and λ the wavelength. For practical reasons we express the model in a compact form:

$$\mathbf{x} = [\mathbf{a}(\theta_0) \quad \mathbf{a}(\theta_1) \quad \cdots \quad \mathbf{a}(\theta_{L-1})] \cdot \mathbf{s} + \mathbf{n} \quad (2.3)$$

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{s} + \mathbf{n} \quad (2.4)$$

where \mathbf{x} is the $M \times 1$ received signal vector, \mathbf{s} is the vector of the L impinging signals, \mathbf{A} is the $M \times L$ steering matrix whose columns represent the steering vectors $\mathbf{a}(\theta_l)$ of the l -th impinging signal, and \mathbf{n} is the Gaussian white noise vector with zero mean. Here, lowercase, bold lowercase and bold uppercase denote respectively scalars, vectors and matrices. The above model allows computing the value of θ , which represents the AoA.




FIGURE 2.3: MIMO system with N -element send and M -element receive coplanar ULA with omnidirectional antennas.

2.3 The ORION system

In this section, we present the system design of ORION. We first describe the system model and the techniques used to estimate the Angle of Departure; then we detail the different phase corrections applied to the Channel State Information to obtain accurate and precise estimates. Finally, we describe our technique used to estimate the terminal's orientation.

2.3.1 A look into the system model

In a MIMO system with N -element send and M -element receive coplanar uniform linear arrays (ULA) with omnidirectional antennas like the one depicted in Figure 2.3, joint estimation of azimuth AoD and azimuth AoA is possible. At the transmitter side, N signals are emitted with identical bandwidth and center frequency. The steering vector of an M -antenna array representing the relative phases at each antenna of a signal received at an angle θ , can be written as stated above: $\mathbf{a}(\theta) = [1 \ a_2(\theta) \ \cdots \ a_k(\theta) \ \cdots \ a_M(\theta)]$. In a multipath environment, the i -th path is represented by θ_i and ϕ_i , which are respectively the AoA and AoD of this path. Hence, the received signal vector is in the form:

$$\mathbf{x} = [\mathbf{a}_r(\theta_0) \otimes \mathbf{a}_t(\phi_0), \ \mathbf{a}_r(\theta_1) \otimes \mathbf{a}_t(\phi_1), \ \cdots, \ \mathbf{a}_r(\theta_{p-1}) \otimes \mathbf{a}_t(\phi_{p-1})] \cdot \mathbf{s} + \mathbf{n} \quad (2.5)$$

where \mathbf{x} is the $MN \times 1$ received signal vector, $\mathbf{a}_t(\phi)$ is the transmit steering vector, $\mathbf{a}_r(\theta)$ is the receive steering vector, $\mathbf{s} = [s_1 \ s_2 \ \cdots \ s_p]^T$ is the vector representing the complex gain of the p paths with $s_i = \alpha_i e^{j\omega_i}$, α_i being the reflection coefficient of the i -th path and $e^{j\omega_i}$ its phase component mainly due to the Doppler effect. As the Doppler shift has almost no effect on the orthogonality of the signals, we chose to ignore it in the remainder of the chapter. \mathbf{n} is the Gaussian white noise vector with zero mean. Here, \otimes represents the Kronecker product and $(\cdot)^T$ denotes the vector/matrix transpose.

Number of Impinging Signals A precise and accurate estimation of the parameters imposes the knowledge of the number of signals present in our measured sample. Before applying an estimation algorithm, we need to identify the signal subspace and the noise subspace. This is mainly conducted by computing the covariance matrix from the sample vector, and then by applying an Eigen decomposition. Thereafter we will subdivide our space according to the eigenvalues. Generally, the highest eigenvalues that contribute to most of the signal are the ones mainly related to the signal of interest. One can make sure of this observation by plotting the scree plot (eigenvalues) and look for the knee of the plot, which represents the point or the frontier between the noise space and signal space.

Estimation algorithm We used a maximum likelihood estimation (MLE) to estimate our AoA and AoD. This approach offers a higher resolution solution to estimate the parameters of a signal impinging on an antenna array. Generally, there is no closed-form solution for the MLE problem, but we can approach it using iterative methods such as expectation-maximization (EM) [98] [47]. This iterative procedure comprises two parts: the expectation step (E-step) and the maximization step (M-step). The first step is the approximation of the data using conditional expectation, while the second step maximizes the data likelihood. We repeat this procedure until the convergence. Knowing that the EM is expensive in terms of computation and suffers from slow convergence, the space alternating generalized expectation-maximization (SAGE) algorithm was proposed to speed up the estimations [54].

SAGE relies on updating sequentially the parameters of the minimization problem on small missing data spaces instead of using the complete data space as in the EM method. It takes advantage of the fact that using smaller missing data spaces will achieve faster convergence. On top of being faster than EM, SAGE keeps the main advantage of maximum likelihood methods, which is its robustness and accuracy even if multipath signals occur. This contrasts with subspace methods that need to apply spatial smoothing to reduce the effect of correlated signal sources [116].

In our case, we adapted the SAGE algorithm to estimate the AoA and AoD in the frequency domain jointly.

We are trying to estimate the AoA and AoD that characterize the line-of-sight signal impinging of our antenna array. In the maximum likelihood estimation method, we could interpret this problem as maximizing log-likelihood function $L(\psi_i; Y_i^{obs}(f))$ of $\psi_i = [\theta_i, \phi_i]$ given an observation $Y_i(f) = Y_i^{obs}(f)$:

$$L(\psi_i; Y_i^{obs}(f)) = 2 \int \Re\{P^H(f'; \psi_i) Y_i^{obs}(f')\} df' - \int \|P^H(f'; \psi_i)\|^2 df' \quad (2.6)$$

Where ψ_i comprises the parameters of the i^{th} path and $(\cdot)^H$ denotes the vector/-matrix conjugate transpose (Hermitian). We estimate the angle parameters by maximizing the above function. The maximum likelihood estimator is:

$$(\hat{\psi}_i)_{ML}(Y_i^{obs}) = \underset{\psi_i}{\operatorname{argmax}}\{L(\psi_i; Y_i^{obs})\} \quad (2.7)$$

$$(\widehat{\theta}_i, \widehat{\phi}_i)_{ML}(Y_i^{obs}) = \underset{[\theta, \phi]}{\operatorname{argmax}}\{|z(\theta, \phi; Y_i^{obs})|\} \quad (2.8)$$

with

$$z(\theta, \phi; Y^{obs}) = v^H(\theta, \phi) Y^{obs} \quad (2.9)$$

where $v^H(\theta, \phi) = \mathbf{a}_r(\theta) \otimes \mathbf{a}_t(\phi)$. Given that $\hat{\psi}_i = [\widehat{\theta}_i, \widehat{\phi}_i]$, we use the following cost function to estimate our AoA and AoD:

$$(\hat{\alpha}_i)_{ML}(Y_i^{obs}) = \frac{1}{MN} z((\theta_i, \phi_i)_{ML}(Y_i^{obs}); Y_i^{obs}) \quad (2.10)$$

Given we have access to a $M \times N$ virtual antenna array, we use the spatial diversity for the estimation procedure instead of time or frequency samples. By using this method, we can carry the estimate of our angles using signals from a single packet.

In our expectation step, we start by canceling interference. We iteratively cancel-out paths from the one with the highest contribution to the lowest. We carry this procedure by iteratively synthesizing a signal according to the estimated parameters and subtracting it from the received signal. The goal is to keep only the contribution of one path at a time before updating its parameters.

The procedure relies on estimating the path parameters to reduce the computation complexity. We get the updated parameters θ'' and ϕ'' of the i^{th} path by sequentially varying one parameter at a time:

$$(\hat{\theta}_i'') = \underset{[\theta]}{\operatorname{argmax}} \{ |z(\theta_i', \phi_i'; \hat{Y}_i^{obs}(f, \hat{\psi}_i'))| \} \quad (2.11)$$

$$(\hat{\phi}_i'') = \underset{[\phi]}{\operatorname{argmax}} \{ |z(\theta_i'', \phi_i'; \hat{Y}_i^{obs}(f, \hat{\psi}_i'))| \} \quad (2.12)$$

$$(\hat{\alpha}_i'') = \frac{1}{MN} z(\theta_i'', \phi_i''; \hat{Y}_i^{obs}(f, \hat{\psi}_i')) \quad (2.13)$$

where $\hat{Y}_i^{obs}(f, \hat{\psi}_i')$ represents the signal after the interference cancellation procedure, (\cdot') and (\cdot'') represent respectively the previous and the current update of the SAGE iterative procedure.

Proper initialization of our iterative method is essential to have convergence. We use spatial correlation to find the AoA and AoD of the Line-of-Sight signal. Then, we reconstruct the LoS signal and subtract it from the received signal. We carry the same procedure for all the other signals.

2.3.2 COTS Wi-Fi and AoD Estimation

We are mainly interested in estimating the orientation of the receiver. Therefore, the receiver plays the role of the target in our study in a line-of-sight (LoS) setting. The idea behind AoD measurement is to send a signal from each element of an antenna array and to measure the relative phase between the signals at the receiver antenna, see Figure 2.4.

The varying nature of the wireless medium imposes that measurements are performed with respect to a specific time delay constraint. In order to collect accurate measurements, we need to make sure that the frame transmission time is lower than the channel coherence time, to avoid phase changes due to channel variation. Otherwise, the phase shift measured between the streams would not be consistent with




FIGURE 2.4: Signal radiation from an N -element uniform linear antenna array: The k -th antenna transmits a signal with a phase shift of $e^{-j2\pi(k-1)\frac{d'}{\lambda}\sin(\phi)}$ relatively to the reference antenna.

the phase shift relative to the AoD. With measurements on the 5.32 GHz band using 802.11 OFDM with an environment speed (walking speed) of $0.33m/s$, we typically have 5.85 Hz of Doppler frequency according to Clarke's model, which equals a coherence time of $T_c = 72.3ms$ [14] and which is much bigger than the time symbol of $T_s = 3.2ms$ and the transmission time of $10ms$.

The AoD estimation is based on the measurement of the phase difference between signals coming to one antenna from the same origin. The idea is to transmit temporally delayed signals on different antennas and then to estimate the phase shift between the received signals. This measurement is difficult to perform when no fine grain synchronization between the signals is available. This criterion is critical knowing that the phase measurement is done based on the wrapped phase, and thus if one does not fine track the phase delays, it will be difficult to retrieve the phase difference between the signals. Thus, a strategy that consists in sending signals at different time delays with coarse time coordination is not suitable. In summary, we need a solution that respects a specific time coherence constraint, and that maintains a fine-grained synchronization between the transmitter's RF chains.

Recent wireless NICs are supporting more advanced MIMO techniques such as SM. SM is one of the most common features, it was mainly adopted for achieving higher data rates, and it consists of transmitting independent data streams over different RF chains in parallel. To estimate the channel between each receiving chain and spatial stream transmitter, the 802.11n amendment adopted the high-throughput long training field (HT-LTF) mentioned in Section 2.2. Thus, by enabling SM at the transmitter wireless network interface controller (NIC), we have access at the receiver to CSI values for each of the N streams, provided the system allows us to collect channel station information.

2.3.3 Phase correction for AoA/AoD estimation

Most of the commercial off-the-shelf (COTS) wireless cards are not intended to be used as measurement instruments, and thus, the measured metrics need to be corrected before launching our estimation procedures. In the following parts, we identify the different problems encountered regarding the phase information, and then we detail our correction design choices to achieve accurate estimation.

Phase correction for AoA estimation The wireless NICs used in this project provide a CSI value for each receiving antenna, each transmitted signal, and each sub-carrier. The CSI's phase and amplitude information capture the channel conditions necessary for computing the AoA. The CSI matrix for each transmitted signal is in the form:

$$\mathbf{CSI}_1 = \begin{pmatrix} h_{1,1} & h_{1,2} & \cdots & h_{1,M} \\ h_{2,1} & h_{2,2} & \cdots & h_{2,M} \\ \vdots & \vdots & \ddots & \vdots \\ h_{C,1} & h_{C,2} & \cdots & h_{C,M} \end{pmatrix} \quad (2.14)$$

This matrix representation with a numbering subscript ($\mathbf{CSI}_{i=1,2}$) will showcase the adopted phase corrections for the AoA estimation. Here, C is the number of subcarriers or frequencies, M is the number of receiving antennas, and $h_{c,m}$ is the channel complex gain for the c -th subcarrier and the m -th antenna. However, exploiting phase information provided by COTS wireless cards is rather challenging. Some design choices are at the origin of several distortions in the phase information. RF oscillator offset, carrier frequency offset (CFO), and sampling frequency offset (SFO) are among the most prominent ones.

RF oscillator phase offset happens when, upon starting up a wireless card, RF chains are locked at different instants. Therefore, each RF chain will have a different constant ψ value added to the measured phase. These phase values are constant for a wireless card till the next recalibration or reset, and thus, we only need to correct the phase offset once for each session. More precisely, this offset is due to the phase-locked loop (PLL) which locks in on a reference frequency. The design of the PLL, which is in our case a fractional-N synthesizer, is not capable of tuning in to the absolute value of the phase; it makes it possible to have a deterministic phase and thus a specific phase coherency between the reference signal and the synthesized one. So, knowing that all the RF chains have an oscillator in common, RF chains are all coherent with the frequency generated by the reference crystal and thus also between themselves. However, the RF chains will not have the same phase because each one of them has their own synthesizers, as shown in Figure 2.5.

Correcting the phase at the hardware level is neither practical nor feasible, so we choose software pre-processing. As the phase offset is constant during the entire session, tuning on a reference signal arriving at a known angle θ_{ref} helps in attenuating the undesired hardware-induced phase shift ψ . This implies applying a phase rotation on the measured target signal phase considering the reference signal phase. Thus, the new CSI matrix will be:

$$\mathbf{CSI}_2 = \begin{pmatrix} h_{1,1} & h_{1,2}\Delta_2 & \cdots & h_{1,M}\Delta_M \\ h_{2,1} & h_{2,2}\Delta_2 & \cdots & h_{2,M}\Delta_M \\ \vdots & \vdots & \ddots & \vdots \\ h_{L,1} & h_{L,2}\Delta_2 & \cdots & h_{L,M}\Delta_M \end{pmatrix} \quad (2.15)$$

where $\Delta_j = \psi_1 \cdot \overline{\psi_j}$ is the relative phase between the first and j -th RF chain for the reference signal coming at θ_{ref} . Here $\overline{(\cdot)}$ denotes the conjugate.

CFO occurs when the transmitting oscillator and the receiver are not synchronized: when the baseband down conversion is done, the signal will be rotated by a constant frequency. As the CFO is applied equally over all the RF chains, it does not affect the AoA estimation.




FIGURE 2.5: Partial view of the transceiver Block with a focus on the phase ambiguity generated by the signal synthesizers.

SFO is because the sender and the receiver sample the signal at different times with an offset δt . The same sampling delay offset is applied to all the subcarriers. In Wi-Fi OFDM, the subcarriers have equally spaced frequencies with $\Delta f = 312.5$ kHz and thus the phase rotation is different for all the subcarriers. The phase rotation applied on the phase data of the k -th subcarrier is $e^{-j2\pi k \Delta f \delta t}$. The peculiar feature of this phase offset is that it varies linearly across subcarriers. However, this linearity is not at first visible in the phase data as the phase wraps around every 2π , and phase unwrapping is the only way to observe phase tendencies across subcarriers and particularly the linearity. Thus, an SFO phase correction that allows consistent inter-subcarrier phase shift information is not needed, as we only focus on phase shifts sanitation between RF chains while estimating the AoA.

Phase correction for AoD estimation Using SM for estimating AoD on COTS Wi-Fi cards imposes to apply RF oscillator phase correction, as explained in Section 2.3.3. Implementing SM in Wi-Fi cards is usually paired with two other mechanisms that are spatial mapping and cyclic shift diversity (CSD), see Figure 2.6.

Spatial mapping is used for matching streams to RF chains. It consists of multiplying the stream matrix by the spatial mapping matrix $\mathbf{V}(k)$ whose columns are orthonormal. Hence, the transmitted signal corresponding to subcarrier k is represented as:

$$\tilde{\mathbf{s}}(k) = \mathbf{V}(k)\mathbf{s}(k) \quad (2.16)$$

where $\mathbf{s}(k)$ is the original signal before spatial spreading. Thus, the channel matrix as seen at the reception is:

$$\tilde{\mathbf{H}}(k) = \mathbf{H}(k)\mathbf{V}(k) \quad (2.17)$$




FIGURE 2.6: Partial view of the MIMO Transmission Block as described and specified in the IEEE 802.11n standard.

Then we need to retrieve the original CSI matrix that is related to the channel, by multiplying the received CSI matrix with the inverse of the spatial spread matrix. Different types of spatial spreading are available among them: direct mapping that consists of sending each stream to an independent RF chain and Walsh Hadamard that relies on using the Walsh matrix to mix the space-time streams. According to the standard, wireless NIC manufacturers may specify their own custom spatial mapping matrix. We have found that the adopted spatial mapping matrix in our scenarios involving two streams is Walsh Hadamard.

CSD is a mechanism adopted by the wireless NICs for delaying the streams (signals) in the time domain to avoid unintentional beamforming. However, it introduces for each stream a phase rotation of $\xi_k = e^{-j2\pi k\Delta f\delta t}$ in the phase data of the streams (signals), with δt , a constant delay applied to all subcarriers of the stream. The channel matrix as seen at the receiver for a subcarrier k and N streams is in the form:

$$\tilde{\mathbf{H}}(k) = \mathbf{H}(k)\mathbf{U}(k)\mathbf{V}(k) \quad (2.18)$$

where $\mathbf{U}(k) = \text{diag}(1, e^{-j2\pi k\Delta f\delta t_1}, \dots, e^{-j2\pi k\Delta f\delta t_{N-1}})$. We compensate the CSD applied by the wireless NIC like for the spatial mapping according to IEEE 802.11n standard, which specifies the applied time delay. We also noticed, while processing the collected phase data, that the delays specified in the standard are not strictly respected, which imposes a fine tuning of the CSD correction. Figure 2.7 shows the effect of phase correction enabling the use of the phase difference between the streams for all the subcarriers to estimate the AoD. Knowing that the CSD induced delay is in the order of hundreds of nanoseconds, it does not affect our condition regarding the coherence time.

2.3.4 Channel correlation and Spatial Diversity

The performance of a MIMO system depends highly on the availability of independent channels. It is well known that channel correlation will degrade the performance of a MIMO system, especially its capacity: this phenomenon is called spatial correlation fading [128]. When the channel correlation is high, using MIMO techniques (like spatial multiplexing that rely on the presence of spatial diversity), do not achieve higher capacity. This is mainly explained by the receiver is incapable of




FIGURE 2.7: Signal phase in radian of each spatial stream as seen by the first receiving antenna with and without phase correction for CSD.

detangling the signals of interest. In our experience, we faced packet loss in environments poor in multipath clusters. Such clusters generally serve as secondary antennas and therefore provide multiple statically independent paths between a transmitter and receiver. This provides the receiver with different copies of the signals of interest and as a result achieves a higher signal-to-noise-ratio (SNR) value. This issue was mainly observed when the experiments were conducted in an environment like an anechoic chamber as shown in Figure 2.9, where the modification of the antenna array positioning, inter-element distance and shape and the presence of reflectors or scatters are necessary conditions for decoding packets sent using a MIMO technique such as spatial multiplexing. Thus, as per our observation in the anechoic chamber about the efficiency of the SM technique, there is a need for an environment rich in multipath clusters such as an office.

Another phenomenon that could influence the channel capacity is mutual coupling between closely spaced antenna elements of an array. This proximity generally influences the performance of the antenna array negatively. However, using our antenna array that respects a half a wavelength distance between elements, we did not observe problems in our real indoor environment in terms of packet decoding.

2.3.5 Estimating orientation

The simultaneous estimation of the AoA and AoD opens the opportunity of estimating the orientation of a target terminal. Rather than relying on an RSSI-based fingerprinting, we propose an approach that exploits the phase information provided while the terminals are engaged in a MIMO communication process. There have been multiple approaches proposed [19, 22] especially in the radar signal processing for aircraft attitude estimation based on the phase difference estimation between the antennas of a particular antenna array geometry.

Our approach, illustrated in Figure 2.8, is similar in substance to the aircraft attitude estimation. However, it differs in that the terminal in question does not have to be equipped with a GPS chip. We exploit an antenna array at the sender and at the receiver, which respectively provides the AoD and the AoA. Intuitively, when the antenna arrays are parallel to each other and on the sample plane, the AoA would be equal to the AoD, and thus any rotation of the terminal would be accounted for

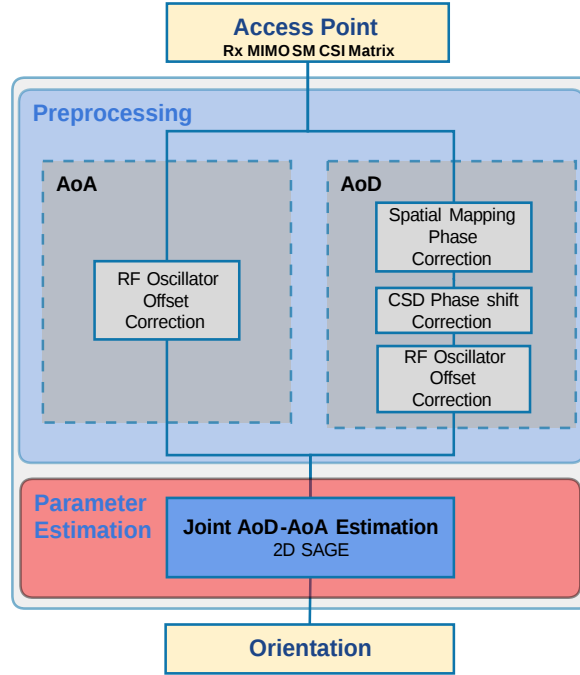


FIGURE 2.8: ORION System Design for orientation Estimation.

as the difference between the estimated AoA and AoD. In summary, the difference between the AoD and AoA gives back the rotation applied to the transmitting terminal. Orientation estimation, even though achievable through an inertial module unit (IMU) installed on the transmitting terminal, imposes more technical constraints to share these measurements with the access point (AP). In our method, all the signal processing is done at a central server. Only measurements are done at the AP level, and no requests are made for the terminal's IMU readings, which is much more convenient as the AP rarely provides these services.

2.4 Evaluation

In this section, we present the hardware, the software and the environment of calibration and test. Then we give the performance results of our design.

2.4.1 Implementation

Our implementation uses COTS Wi-Fi network interface cards (Intel Wireless Link 5300 AGN) with a modified version of the firmware to extract the CSI matrices using the Intel CSI tool for 802.11n HT packets [11]. In the case of OFDM systems, we can extract a CSI matrix for each subcarrier. In our case, the Wi-Fi cards offer up to 30 subcarriers. Knowing that the wireless card's CSI suffers from phase inconsistencies in the 2.4GHz band, we use instead the 5 GHz band with 20MHz of bandwidth that does not suffer from the same issue. Those wireless cards are suffering from cross-talk on the third RF chain port. Cross-talk happens when the wires in the integrated




FIGURE 2.9: Inria Sophia Antipolis' faradized anechoic chamber (R2lab) with 37 Wi-Fi nodes equipped with 3-elements ULA and Intel 5300 cards. The antennas are placed under the nodes (blue boxes) to have clear line-of-sight (LoS) communication between nodes.




FIGURE 2.10: Omni-directional antenna array rotated through a computer controlled module (red circle).

circuit are adjacent, due to magnetic field interference. So, we set up our wireless cards in injection mode using the Lorcon package^{*}. This avoids the need of association with an AP and allows raw Wi-Fi packets transmission. The injection mode relies on raw socket programming, and thus, we are not tied to the IEEE 802.11n media access control (MAC) layer mechanisms. As for the receptor, knowing there is no association, we use the monitor mode of the wireless card to receive packets. All the packets are processed at a central server where the phase corrections are conducted. The transmission is controlled by disabling the antenna selection algorithm and by specifying the desired number of antennas, streams and transmission techniques (SM). We have to make sure that the number of streams corresponds to the number of antennas chosen for the experiment. The number of streams corresponds to the modulation and coding scheme (MCS) index. We explain in further detail the method for computing the rate index in the ORION notebook[†]. The rotation is

^{*}<https://github.com/dhalperi/lorcon-old>

[†]ORION Notebook: <https://github.com/naoufal51/Orion>






FIGURE 2.11: Experimental cumulative density function of orientation error according to absolute reference.



(A) Box-plot of Orientation Error using SAGE. (B) Box-plot of Orientation Error using MUSIC.

FIGURE 2.12: Box-plot of Orientation Error using SAGE and MUSIC

realized using extension cables that connect the coplanar transmitter and receiver antenna arrays to the access points. To control the rotation of the antenna array in a fine-grained manner, we use a computer-controlled module for applying a specific rotation angle on the transmitting antenna array, see Figure 2.10. The calibration process of our system was conducted in the R2lab anechoic testbed before launching real life tests.



(A) Box-plot of Orientation Error using SAGE (B) Box-plot of Orientation Error using MUSIC w/ outliers.

FIGURE 2.13: Box-plot of Orientation Error using SAGE and MUSIC w/ outliers.



FIGURE 2.14: Office room where the experimentation is conducted.

2.4.2 Experimental Results

We conducted our experiments in an office room, which contains many scatters, reflectors, and other multipath sources. This environment has co-channel and near-channel interference from Wi-Fi hot spots with 48 functional independent basic service sets (IBSS) present in the building. The transmitter and receiver are placed in a coplanar setting with a maximum distance of 3m.

Orientation accuracy: In Figure 2.11, we compare the performance of multiple signal classification (MUSIC) and SAGE in terms of orientation estimation accuracy. We use only one packet and one sub carrier for each estimation and apply strictly the same phase corrections for both methods. In Figure 2.11, we can observe that the orientation accuracy and precision of MUSIC are low, with 9 degrees of error for 80% of the orientation estimations. As for SAGE, we observe a dramatic increase in accuracy and precision with 1.32 degrees of error for 80% of the estimations. The accuracy has jumped to 63.7% using the SAGE algorithm.




FIGURE 2.15: Experimental cumulative density function of orientation error for long distances according to absolute reference.

Impact of transmitted streams: We then vary the number of transmitted space-time streams (STS) and observe that using 3 STS instead of 2 allows for a 51.8% increase in accuracy when using SAGE, and a 10.2% increase when using MUSIC. The impact of using 3 STS is more perceptible when using SAGE than when using MUSIC. We can even observe that SAGE using only 2 STS still outperforms MUSIC with 3 STS by 23.9%.

Impact of wide angles on estimation accuracy: In Figures 2.12a and 2.12b, we examine the impact of the antenna array orientation angle on the estimation performance. The results show that when the rotation angle of the antenna array is superior to 50 degrees, added errors are observed, and the accuracy drops for both SAGE and MUSIC. This is mainly due to the fact that the elements of the antenna array are not completely visible to the receiver. Even though this error is common to both methods, the impact is different. The MUSIC algorithm cannot give any consistent estimation, and the precision drops tremendously with 178% of median error, in comparison SAGE still holds a fairly good precision with a 19.7% median error.

In Figures 2.13a and 2.13b, we can observe that MUSIC has a fair number of outliers far from the median value; e.g. for 5 degrees, the median error is -10% and the maximum outliers is 2280%. However, for SAGE, even if outliers still exist, they are located close to the median compared to the MUSIC case.

Impact of distance on estimation accuracy: The goal here is to show we can get decent orientation estimation performance even for longer distances. For that purpose, we deploy a testbed on a corridor where we test the estimation accuracy for distances up to 11 meters. We place the gateway at the same level as the terminal. The corridor is narrow and rich in multipath, which increases the number of multipath clusters. The results of the experiment, see figure 2.15, show that we have a median error of 5.6 ° with 10.8 ° of error for 90% of estimations for 11 meters and 3.9 ° with 8.6 ° of error for 90% of estimations for 6 meters.

2.5 Related Work

Wi-Fi based positioning is prolific, however, little focus was given to orientation estimation using Wi-Fi signals. Most of the contributions go towards making a robust estimation of a terminal's position using IMUs (gyroscope, accelerometer, and magnetometer) [9, 13, 25]. The main idea is to identify the physical orientation and compensate for the localization estimation error it generates. Authors of other related works [15, 115] considered the estimation of the position and the orientation by building a radio map using Wi-Fi received signal strength (RSS) and with orientation estimation limited to 4 directions. The ORION system allows for orientation estimation with an accuracy of the order of the degree as shown by the estimation of AoD and AoA usually used in bistatic radar systems [76] [36] to ensure the position detection of a passive target.

2.6 Conclusion

We described ORION, a system that can estimate the orientation of a Wi-Fi terminal using unmodified off-the-shelf Wi-Fi devices. We proposed phase correction techniques for both AoA and AoD estimations, and we applied a joint estimation of both parameters to enhance the performance accuracy. ORION exploits MIMO techniques widely adopted in Wi-Fi equipment such as spatial multiplexing, which makes it deployable in practice. We evaluated our system to be accurate for a joint AoA and orientation estimation in a non-controlled environment (an office room). We plan to extend ORION for estimating the pitch and the roll as well using two-dimensional antenna arrays. Detailed information on how to reproduce the experiments made on the R2lab anechoic chamber and the office room is available at <https://github.com/naoufal51/Orion-Augmented> and <https://github.com/naoufal51/Orion>.

Chapter 3

Joint localization and range extension for LPWAN

3.1 Introduction

LPWANs are sought to be the networks of choice for large-scale IoT communication in urban areas. Their main selling points are a low power consumption that can span a decade and broad geographical coverage of a few kilometers [110]. They also offer a cost-effective alternative to their cellular counterparts by operating on the sub-gigahertz industrial, scientific and medical (ISM) radio band, thus avoiding costly licenses. Unlike cellular networks which give great importance to dimensioning and have adopted industry standards concerning the quality control, privately owned LPWANs suffer from an inherent heterogeneity and poor dimensioning that hinder their performance in terms of rate, or poses serious reliability questions. This chapter presents Snipe [92] *, a diversity combining LoRa system with spatial filtering techniques providing: (1) an enhanced communication range for end-devices and gateways alike and (2) a precise and accurate position estimation system based on radar techniques. By exploiting the multipath effect omnipresent in urban and in-building settings, the Snipe gateway can increase the received SNR and hence enhance its decoding capabilities. The system uses spatial filtering techniques to achieve a higher SNR at the end-device. While most contributions are towards improving the uplink [52, 49], Snipe enhances both the uplink and downlink by using a MIMO design that respects the LoRa energy consumption restraints. MIMO support at the gateway provides a low-cost AoA and end-device position estimation. By offloading the position estimation processing to the gateway, the Snipe system alleviates the energy cost incurred when using GPS-enabled end-devices. By hosting the MIMO signal processing locally at the gateway, the proposed system allows combining the signals coherently from each antenna element indiscriminately. This proposal contrasts with cloud-based approaches that involve renting compute and storage nodes and providing links with large bandwidth [49].

We organize the rest of the chapter as follows: Section 2 reminds some basic LoRa principles. Section 3 presents the range extension problem and the approach to address it by leveraging coherent combining and spatial filtering. Section 4 presents in more detail the AoA estimation technique needed to perform spatial filtering.

*https://github.com/naoufal51/MIMO_LoRa.git




FIGURE 3.1: LoRa Gateways each equipped with a uniform linear array apply spatial filtering by adopting a more focused radiation pattern (beamforming) for Transmitting Down-link packets to clients instead of the usual omnidirectional radiation pattern.

Section 5 presents the power gain for both the uplink and the downlink. Section 6 describes the related work, and Section 7 concludes the chapter.

3.2 LoRa primer

Many IoT applications require low cost, low energy, and long-range communication. Spread spectrum techniques provide enhanced robustness with frequency diversity and low power density with signals transmission below the noise floor. However, classical direct-sequence spread spectrum (DSSS) systems for IoT applications need highly accurate and expensive reference clock source violating the low-cost requirement. LoRa uses Chirp Spread Spectrum, a modulation technique that uses chirps as carrier signals to encode information. An up (down) chirp is a tone in which frequency increases (decreases) linearly with time. An advantage of this specific spread spectrum technique is that timing, and frequency offsets between transmitter and receiver are equivalent, which reduces the complexity (and therefore the cost) of the receiver design. Different parameters can customize the modulation, such as the bandwidth (BW) and the spreading factor (SF). The SF is equal to the number of raw information bits per symbol. At the cost of a reduction of bit rate, adopting a higher SF results in higher time on air, which provides better coverage and a higher packet delivery ratio. LoRa can change the SF according to the link quality, and different orthogonal SFs can be supported to trade data rate for sensitivity within a fixed channel bandwidth. In the EU863-870 unlicensed radio spectrum in the ISM band, any LoRa end-device [89] has to support at least three channels of 125 kHz.

Figure 3.4 shows a simplified view of the LoRa demodulation process. The frame payload (part a of the figure) comprises a chirp sequence transmitted with an SF. Chirps are cyclically shifted, and frequency jump (Δf) at the beginning of each chirp encodes user information. To decode this information, the received signal is multiplied by a down-chirp sequence with the same length and SF (part b). Then, FFT is applied to the product to provide the encoded information (part c).




FIGURE 3.2: LoRA demodulation process.




FIGURE 3.3: Spectrum showing the different SF, where the higher the SF the longer the time-on-air.

3.2.1 Physical Frame Format

LoRa transmitters and receivers implement the following physical frame format. Each frame starts with a preamble made with a sequence of up chirps; the last two up chirps encode the sync word used to discriminate between devices that use the same frequency band. A device with a given sync word will discard packets with a different sync word. After the sync word, there are two and a quarter down chirps followed by a sequence of choppy up-chirps representing the payload as in Figure 3.4. An optional header containing information about the size of the payload (in bytes) could follow the preamble, the code rate used for the end of the transmission, and an optional 16-bit cyclic redundancy check (CRC) for the payload at the end of the frame.




FIGURE 3.4: LoRA Physical frame format composed of sequential Chirp Signals.

3.3 LP-WAN MIMO Gateway for Range extension

Obstruction and multi-path propagation can cause severe channel degradation in indoor environments. This may significantly reduce the communication range in the case of in-building LoRa deployments. In these environments, it is therefore important to try increasing the range in both the uplink (from the end-device to the gateway) and the downlink (from the gateway to the end-device) directions. This will help to widen the coverage and allow the cost-effective deployment of these LPWANs. Hopefully, leveraging spatial diversity with MIMO techniques can mitigate the range reduction and gain some dBs.

About the uplink, several state-of-the-art techniques leveraging temporal, frequency, or spatial diversity can be envisaged to enhance the SNR. Temporal diversity is not very appropriate in LPWAN networks because it would require a high duty cycle of end-devices (supposed to remain low) and increase their energy consumption. Frequency diversity is already well exploited in LoRa with the Chirp Spread Spectrum modulation that allows robust signal decoding. It can also improve the communication range by increasing SF. However, this will increase the Time on Air and then energy consumption, which is detrimental for the end-devices.

In this chapter, the proposed approach is to exploit spatial diversity by adding a second antenna at the gateway to intelligently combine the signals received on both antennas to increase the SNR at the gateway and therefore the uplink communication range. As for the downlink, a second antenna allows to beamform the signal to specific end-devices, thus improving the downlink range. However, this enhancement requires knowing the AoA of the signal sent by the end-device to the gateway to point the beam toward the end-device.

Using MIMO techniques makes it possible to provide the gateway with information on the arrival angle of the signal coming from end-devices. We may combine this additional information with various information from other sources such as the TDoA to improve the accuracy of the localization of LoRa devices. Traditional LoRa localization techniques based on TDoA lack in precision when multipath is present and rely on complex, fine-grained synchronization [53].

Regarding leveraging spatial diversity in LPWANs, previous work has proposed

combining signals from an end-device arriving at several gateways in the cloud to improve the SNR and thus increase the range of the uplink. The approach described in this chapter, supporting a second antenna at the gateway allows improving the range in the uplink direction while saving bandwidth between gateways and the cloud infrastructure. The cost of this approach in terms of CPU and energy consumption is reasonable because adding the processing functions takes place only at the gateway which, unlike end-devices, does not use a battery but is connected to a power source.

In the rest of the section, we present the range extension techniques used to enhance the transmission range from the gateway to the end-device and vice versa.

3.3.1 Coherent Combining for Uplink Communication

LPWAN's are getting very popular as a communication platform for the Internet-of-Things in large cities with big buildings and dense population. This situation mainly limits the signal range because of the attenuation from the buildings and other sources. This phenomenon will also result in battery drainage of the end-devices, forcing them to use of slow data rates to reach even the closest gateway. Hence, there is a need to improve the received signal's quality at the reception (gateways), allowing decoding the information of interest.

We take advantage of the MIMO capabilities of the proposed system to increase the SNR at the gateway and improve the decoding of the received signals. By leveraging multiple RF chains and antennas, we expect to benefit from the diversity factor that comes from the spatial separation of antennas, thus reducing the effects of multipath and co-channel interference (CCI) in wireless systems [ref12]. The rationale is that when the target signal suffers from multipath fading, different copies of the signals can be combined to provide a more robust representation of the signal. We chose in our system to adopt a weighted average approach, where the weights are selected to delay appropriately each copy of the signal to align the different copies (see Figure 3.6). Thus, the signals are added up constructively:

$$y_{comb} = \sum_{v=1}^N h_k^* y_k ,$$

where y_k is the received signal at the k^{th} antenna element and h_k is the channel estimate for the k^{th} antenna. The output y_{comb} is the summation of N matrices.

The combined LoRa signal shows an SNR almost equal to the highest SNR as seen from 3.6

3.3.2 Beamforming for Downlink Communication

As the name suggests, beamforming technology realizes directional communication. There are two types: digital beamforming and analog beamforming. We prefer to use analog beamforming as it works in the RF chain, and it is useful for generating a narrow directional beamform. Then, the higher the number of antennas, the narrower the pattern will be. In this work, we use adaptive beamforming to direct RF radiation patterns for the downlink messages according to the AoA from the uplink communication. For example, when an end-device attempts the transmission




FIGURE 3.5: Upon receiving a signal at the LoRa gateway's antenna array, the different signal copies are combined constructively according to their signal-to-noise ratio.

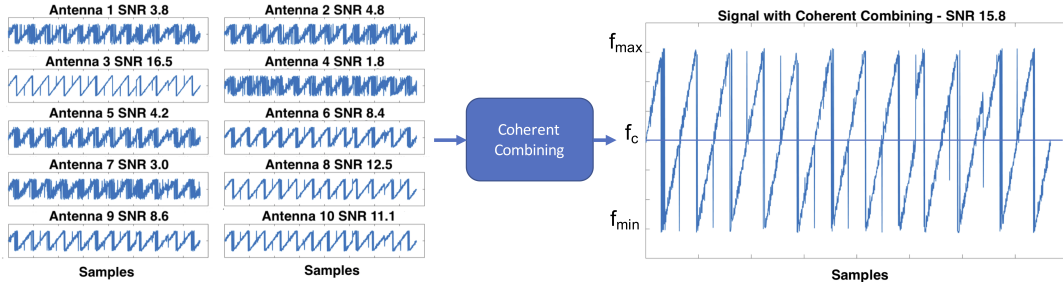


FIGURE 3.6: Distorted Chirps in the LoRa signals impinging on the antenna array are combined by aligning, scaling and summing the different copies producing a signal with a higher SNR.

of a 'confirmed' frame towards the network, it expects to receive an acknowledgment in one of the subsequent reception slots. In the acknowledgment's absence, it will try to retransmit the same data again. However, if we could make the gateway to relay the acknowledgment by targeting the end-device, we could avoid message retransmission of the same data, and, we will increase the battery life of the end-device. We achieve reliability in LoRaWAN through the message acknowledgment in downlink communication, which increases the capacity of the network. Therefore, it becomes mandatory to minimize the number of downlink frames to avoid the channel capacity drain. For all these restrictions, fair spectrum sharing is essential beyond the duty cycle limitation. With the beamforming technology, instead of minimizing the downlink messages, the acknowledgment/downlink messages can be sent efficiently, increasing the reliability of the transmission and also by abiding the duty cycle restrictions.

Spatial filtering As presented earlier, we propose to use beamforming or spatial filtering as a suitable candidate for not only increasing the decoding power of end-devices but also for reducing interference between devices. Beamforming allows a higher gain in the target direction while attenuating other directions with no directional antennas. More precisely, by sending simultaneous signals from the different




FIGURE 3.7: LoRa signal beamforming by applying complex weights to each RF chain. This allows combining the signals constructively towards the direction of the target.

elements of the Tx RF chain, we combine the omnidirectional radiated signals to form a higher gain beam towards the end-device. To exploit this feature, we need to control the direction of the beam by defining complex weights on each one of the RF chains, see Figure 3.7. By defining a weight that controls the phase of the radiated signal at each antenna, we can augment the response in the desired direction:

$$w_k = \alpha_k e^{-j2\pi(k-1)d \sin \theta} \quad (3.1)$$

$$z(t) = \sum_{k=0}^N w_k y_k(t) \quad (3.2)$$

Where θ is the beam steering angle, d is the distance between the radiating antennas, α_k the signal amplitude weight, N is the number of antenna elements in the array and y is the received signal.

Note that the weights we define are related to the position of the target. Thus, we need to accurately estimate the AoA of the signal radiated from the end-devices to steer the Tx beam to the direction of the target.

3.4 LP-WAN MIMO Gateway for AoA-based localization

To propose a system that allows continuous tracking of the position of end-devices while updating their steering angles, we propose to estimate the AoA of a LoRa end-device by using the gateways as reference points (localizers). The AoA estimation allows tracking the origin of a received LoRa chirp sent by the end-device through measuring the signal's phase from different positions (spatially separated antennas). The diversity provided by the chirp signal's distribution in time and frequency allows for a more robust estimation of the AoA.

The proposed AoA-based localization system will consume less power and provide a precise and accurate position estimation. On the one hand, adopting an opportunistic localization approach reduces the energy consumption burden on the end-device induced by GPS modules. But unlike TDoA localization techniques that require fine-grained time synchronization between the gateways and the end-device, AoA exploits phase data, thus providing an accurate localization system at a lower cost.

Measuring the AoA is based on the observation that the signal of interest impinging on the antennas is received at slightly different time delays, see Figure 3.8.




FIGURE 3.8: Signal impinging on an N -element uniform linear antenna array: Each antenna receives a copy of the signal with a phase shift $k.d.\sin(\theta)$ relative to the reference element (First antenna), with d , the antenna spacing, in the order of half a wavelength.

As the measured baseband signal is in the frequency domain, we compute the relative phase between the signals received at the antenna array. The distance between the antennas and the signal's direction of arrival define this phase shift. Usually, this distance is expressed in wavelengths and is in the order of half of the wavelength λ . Previous work [69] shows that respecting this distance constraint makes it possible to reduce the effect of mutual coupling, which usually occurs when the antennas element of the antenna array are too close (less than 0.2λ) and creating side lobes that can pollute the AoA estimations if they are too far (over 0.5λ).

The following matrix representation can express the received signal and the relative phase:

$$\mathbf{x} = \begin{pmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-1} \end{pmatrix} = \begin{pmatrix} a_0(\theta_0) & a_0(\theta_1) & \cdots & a_0(\theta_{r-1}) \\ a_1(\theta_0) & a_1(\theta_1) & \cdots & a_1(\theta_{r-1}) \\ \vdots & \vdots & \ddots & \vdots \\ a_{N-1}(\theta_0) & a_{N-1}(\theta_1) & \cdots & a_{N-1}(\theta_{r-1}) \end{pmatrix} \begin{pmatrix} s_0 \\ s_1 \\ \vdots \\ s_{r-1} \end{pmatrix} + \begin{pmatrix} n_0 \\ n_1 \\ \vdots \\ n_{N-1} \end{pmatrix} \quad (3.3)$$

where $a_k(\theta) = e^{-j2\pi k d \sin(\theta)}$ represents the phase shift between the signals received by the reference and the k^{th} antenna. For practical reasons we choose to express the signal in a compact form:

$$\mathbf{x} = [\mathbf{a}(\theta_0) \quad \mathbf{a}(\theta_1) \quad \cdots \quad \mathbf{a}(\theta_{r-1})] \cdot \mathbf{s} + \mathbf{n} \quad (3.4)$$

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{s} + \mathbf{n} \quad (3.5)$$

Where \mathbf{x} is the $N \times 1$ received signal vector, \mathbf{A} is the $N \times r$ steering matrix which columns are the steering vectors $\mathbf{a}(\theta_k)$, \mathbf{s} is the vector composed of r transmitted signal s_k and \mathbf{n} is the Gaussian white noise vector with zero mean. Here, lowercase, bold lowercase and bold uppercase denote respectively scalars, vectors and matrices.

To estimate the angle, we use the MUSIC algorithm [123]. The first step of the procedure consists in computing the sample covariance matrix. Usually, time distributed samples (packets) can obtain a robust estimation of the covariance. However, as very

few packets are exchanged in LoRa applications, we adopt an estimation procedure that involves time-frequency distributed samples.

Choosing such an approach allows taking advantage of the intrinsic diversity of chirp modulated signals. Thus, we compute the sample covariance matrix from the in-phase and quadrature (IQ) samples corresponding to L observations in time $\hat{\mathbf{R}}_{\mathbf{xx}} = \frac{1}{L} \sum_{i=1}^L \mathbf{xx}^H$. Then, we compute the eigenvectors and eigenvalues of the $\mathbf{R}_{\mathbf{xx}}$ by an eigen-decomposition, where $\mathbf{R}_{\mathbf{xx}} = \mathbf{Q}\mathbf{D}\mathbf{Q}^H$. Based on the assumption that the signal and noise vectors are uncorrelated, we partition the \mathbf{Q} matrix containing the eigenvectors into two subspaces; one spanned by the signal eigenvectors \mathbf{Q}_s corresponding to the largest eigenvalues and the other by the noise eigenvectors \mathbf{Q}_n . Then we use the MUSIC algorithm [123] to estimate the parameters by launching a search over all the possible values of the steering vectors to identify the ones related to the signal. This is possible because the signal and noise are uncorrelated. Thus, the signal steering vectors are orthogonal to the noise eigenvectors and ultimately to the noise subspace. Finally, we compute what we call the pseudospectrum from the quadratic function,

$$\mathbf{P}(\theta_i, \phi_i) = \frac{1}{\mathbf{a}(\theta_i)^H \mathbf{Q}_n \mathbf{Q}_n^H \mathbf{a}(\theta_i)}$$

where $(\cdot)^H$ denotes the Hermitian operation. The estimated AoA corresponds therefore to peaks in the computed spectrum.

By estimating the AoA for each one of the received signals, we can adapt the radiation pattern and the direction of the lobe. Then the gateway can now track a client whenever there is a change in either the environment or the client node's position.

To achieve a better estimation of the signal source position, a calibration procedure is mandatory in most cases where the different RF chains of a transceiver are not synchronized. During the frequency tuning phase, each one of the RF chains will lock at different phase offsets if they do not share the same phase reference. This is mainly mitigated by the calibration process where the phase delay between the reference RF chain and the others is estimated using a reference signal coming from a waveform generator with known properties, and applying a phase delay on the RF chains according to the reference RF chain [113].

3.5 SDR implementation

Most LoRa transceiver implementations [64] expect the Spreading Factor and Bandwidth and frequency information to be explicitly given as parameters for decoding. LoRa implements the frequency-hopping spread spectrum (FHSS) method to mitigate interference by allowing end-nodes to send packets on different frequencies 868.1 MHz, 868.3 MHz, and 868.5 MHz in a circular way. Therefore, to achieve our joint range extension and localization system, we need to infer these parameters for received LoRa signals. We implemented the following approach:

- **Spreading Factor** Algorithm 12 describes the procedure to detect the SF and the BW. LoRa modulated signals are received, and their IQ samples are de-noised with a threshold. Then the signal is dechirped by multiplying the received 'frequency-shifted LoRa chirp' with 'conjugated time-reversed chirp'

so we will get the constant frequency signal, with a specific frequency characteristic of the transmitted signal. This multiplication is done with the inverse chirps generated for all the SFs. Now if we take the fast Fourier transform (FFT) over the entire symbol, then the bin with the highest energy, which is the fundamental frequency, will represent the transmitted symbol. The plot with the highest peak represents the SF of the received signal. Here in 3.9 plot for SF 9 the highest peak than the other FFT plots.

- **Frequency Detection** We perform frequency detection by computing the power spectral density (PSD) of the received signal. The PSD describes how the power of the signal is distributed over frequency while the discrete Fourier transform (DFT) shows the spectral content of your signal, the amplitude and the phase of harmonics in your signal. Then we are applying thresholding on the signal to remove the noise according to an empirically determined limit, which corresponds to the half of the maximum value of this peak.

Algorithm 1 Spreading Factor Detection

```

1: procedure SF DETECTION
2:    $signal_{denoised} \leftarrow Threshold(signal)$  ▷ Denoising by thresholding
3:   for  $SF \leftarrow 7to12$  do
4:     for  $BW \leftarrow 125KHz, 250KHz$  do
5:        $downChirp \leftarrow complex(chirpI, chirpQ)$ 
6:        $convSignal \leftarrow signal_{denoised} * conj(downChirp)$ 
7:        $FFTresult \leftarrow abs(fft(convSignal))$ 
8:        $Peak(SF, BW) \leftarrow \sqrt{max(FFTresult)}$ 
9:     end for
10:  end for
11:   $SF, BW \leftarrow argmax(Peak)$  ▷ The estimated SF and BW
12: end procedure

```

3.6 System Evaluation

To test the system performance, we assess the localization precision and the power gain in a 700m² office and outdoor environments. We build both the end-device and the proposed MIMO gateway as custom software-defined radio systems using GnuRadio and Ettus Research's universal software radio peripheral (USRP) B210 software-defined radio on laptops with 32 GB RAM and a 2.70GHz i7-4800MQ processor. The MIMO gateway has a 2-element ULA (Vertical antenna VERT 900MHz 3dBi) for beamforming and a similar ULA for AoA estimation and coherent combining, while the end-device has only one antenna for transmission and another one for reception. We established two channels between the gateway and the end-device: the uplink on the 900 MHz frequency and the downlink on the 868MHz frequency. We use a band-pass filter with a bandwidth of 4MHz for the 869MHz frequency that covers the band of interest to filter out-of-band interference. In the following evaluation, we use a LoRa signal with an SF of 10 and 125KHz of bandwidth.

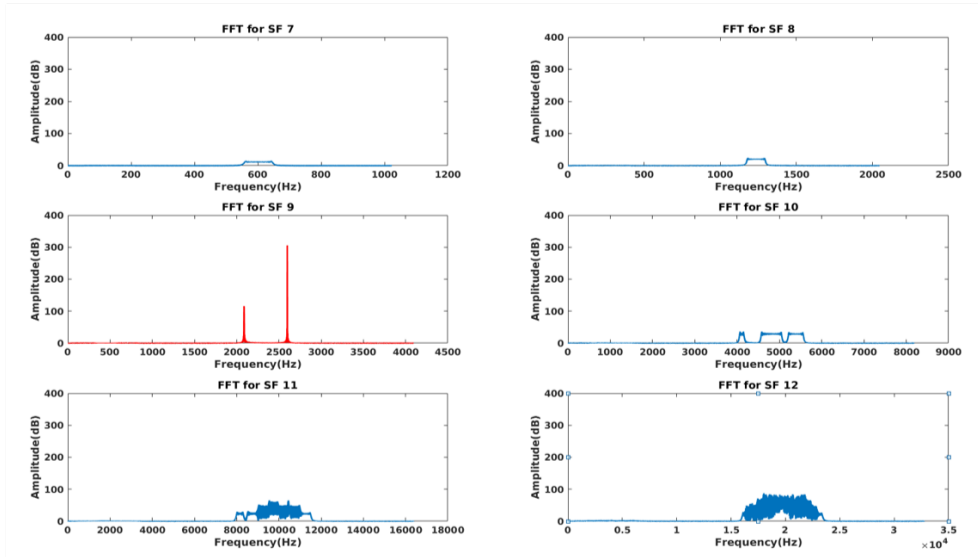


FIGURE 3.9: FFT for all SFs




FIGURE 3.10: USRP b210 equipped with a 2-element Uniform Linear Array for estimating the AoA and spatial filtering.

3.6.1 Position estimation precision

To evaluate the accuracy of our AoA localization system for LoRa, we conduct a series of experiments in various environments (indoor and outdoor) and conditions (LoS and NLoS).

Evaluation of the positioning system in Indoor Environments.

Distance, multipath, and environment geometry have a significant impact on localization accuracy and precision. Thus, we deployed our system in various typical indoor environments with abundant multipath clusters, namely a corridor, an office, and a meeting room with a total area of 700 square meters. The testbed used to evaluate our system is shown in Figure 3.11. As shown in Figure 3.12, the AoA estimation in the office environment has a mean error of 4.84° . This is explained by the relatively short distance between the gateway and the end-device, as the maximum distance between the transceivers is 12 meters. Regarding the corridor and

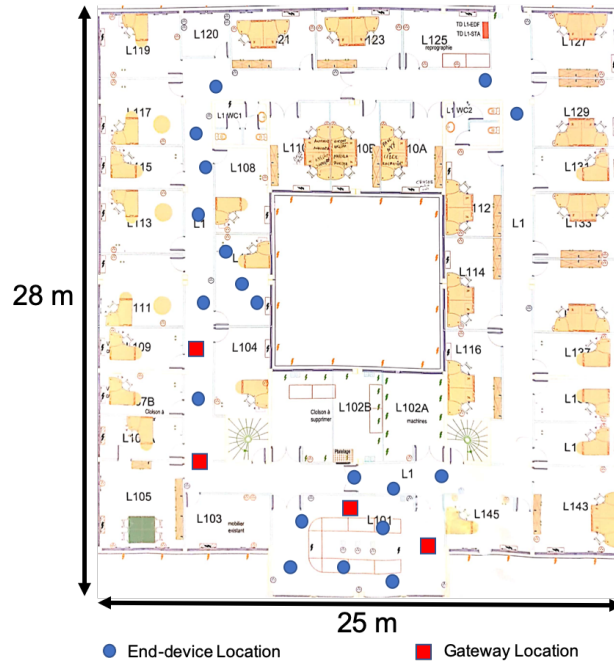


FIGURE 3.11: Experiment testbed showing the end-device locations (blue circles) and the gateway locations (red squares). The testbed is designed to evaluate Snipe in a wide variety of deployment scenarios indoors. Snipe has been tested when both the gateway and the end-device are along a corridor, office and meeting room.

meeting room environments, the mean error is 12.10° and 11.16° , respectively. We observe a drop in AoA accuracy due mainly to the estimation algorithm (MUSIC) that suffers from the correlation between the multipath signals and the Line-of-Sight signal. The longer the distance between the transceivers, the bigger the number of reflected signals, the higher the estimation error. However, the error is still relatively small, thanks to the estimation process that samples the signal in time and frequency. Finally, with a through-the-wall position estimation, the mean AoA error is 4.66° , which is explained by the limited distance between the transmitter and receiver. We observe that the effect of reflected signals on the wall is negligible, which is mainly because of signal power at the gateway where the LoS signal is attenuated but not occluded.

Evaluation of the positioning system in Outdoor Environments.

Outdoor environments are challenging because of vast distances between the end-device and the gateway and heterogeneous multipath sources, which entails extremely attenuated signals with difficulty to recover the phase information on top of having highly correlated signals. The task becomes even more challenging when the communication is near-line-of-sight and NLoS between the transceivers where the LoS signal is weakened and becomes indistinguishable from other multipath signals.




FIGURE 3.12: ECDF of the AoA estimation error in different environments.



FIGURE 3.13: Experiment testbed showing the end-device locations (blue circles) and the gateway locations (red squares). The testbed is designed to evaluate Snipe in outdoor LoS condition with large building, cars and vegetation.

Position	Distance (m)	Relative altitude (m)	GT AoA [°]	Est AoA [°]	Abs Est Error [°]
North Bench	5	0	-26.6	-21.01	5.59
East Bench	20	0	-37	-10.85	26.1
LEAT Bench	50	0	-13.7	-11.16	2.54
Down	70	-8	-12.9	-10.36	2.60
Up Bat1	90	0	-4.4	-11.22	6.82
Templiers1	130	14	12.92	-8.09	21

TABLE 3.1: The estimated Angle of arrivals and the incurred errors according to the distance from the gateway and the elevations for LoS Experiments.

Position	Distance (m)	GT AoA [°]	Est AoA [°]	Abs Est Error [°]
NLoS 1	20	32,4	8,75	23,7
NLoS 2	42	12,4	-3,28	15,7
NLoS 3	25	3,57	-22,3	25,9
NLoS 4	32	19,5	-21,9	41,4

TABLE 3.2: The estimated Angles of arrival and the errors according to the distance from the gateway for NLoS Experiments.

Line-of-Sight We evaluate the performance of our system for such a scenario in the outdoor testbed shown in Fig 3.13. We deploy this testbed on the Sophia Tech campus, where the line-of-sight conditions are preserved, with buildings, trees, and cars serving as multipath clusters. To estimate our AoA, we use the same setup and procedure described earlier with end-devices placed at different distances and elevations. As seen from the Table 3.1, the mean AoA error for Snipe is around 10.7°. This performance is comparable to the corridor’s results; known for being plagued by multipath due to its narrowness. Given that our array is one-dimensional, the difference in level influences negatively the AoA precision as it renders the broadside angle (a combination of azimuth and elevation) instead of the azimuth.

Near-line-of-sight and NLoS The goal of this evaluation is to assess the resilience of the system to partial or total blockage. We define NLoS conditions for our gateway where all direct paths to the end-device are obstructed, and then LoS becomes weaker than reflected signals. For that, we deploy a testbed with a gateway placed on an elevated vantage point and end-devices in positions where the direct path is blocked. The elevation in this scenario is negligible compared to the LoS outdoor testbed. As seen from the Table 3.2, the mean AoA error for Snipe is around 26.7°. This performance is worse than those obtained indoors, which can be explained by the EM characteristics of the blockage.

3.6.2 Coherent Combining

To evaluate the accuracy of the coherent combining system for LoRa, we conduct a series of experiments in various environments (indoor and outdoor).



FIGURE 3.14: Outdoor Experiment testbed deployed on Inria campus showing the end-device locations (blue circles) and the gateway locations (red squares). The testbed is designed to evaluate Snipe in an outdoor environment rich with multipath and in NLoS propagation conditions.

Position	Distance (m)	Relative altitude (m)	Coherent Combining Gain [dB]
North Bench	5	0	5.42
East Bench	20	0	5.41
LEAT Bench	50	0	4.93
Down	70	-8	5.95
Up Bat1	90	0	6.05
Templiers1	130	14	5.16

TABLE 3.3: The power gain using coherent combining according to the distance from the gateway and the elevations for LoS Experiments.

Coherent Combining in Outdoor Environments

To test the power gain using coherent combining, we have deployed experiments in an outdoor environment. The goal here is to show we can achieve a better decoding power by combining signals constructively in conditions where there are rich multipath conditions. For that purpose, we use the same outdoor line-of-sight testbed described earlier, see Fig 3.13, where the multipath comprises cars, buildings, and trees. In Table 3.3, we observe that mean power gain is 5.4 dB, thus making the link between the end-device and the gateway more robust.

Coherent Combining in Indoor Environments

To test the power gain using coherent combining, we have deployed experiments in both a corridor and a meeting room environment. We show in Figure 3.15 the




FIGURE 3.15: Empirical CDF power gain using coherent combining in a meeting room and corridor.

result of the coherent combining of the signals impinging of the 2-element antenna array. We observe an average power gain in the meeting room and corridor of 5.81 dB and 5.99 dB respectively, which is a 400% gain compared to the baseline. By applying the coherent combining, we can enhance the link between the Tx and Rx. We also observe that we retain roughly the same power gain across environments and conditions (LoS and NLoS).

3.6.3 Beamforming

The experiments to evaluate the beamforming based on AoA for LoRa were done in the following manner. We have used two USRP B210s. One as a gateway (beamformer) and other as end-device (beamformed). Since commercial LoRa devices do not offer power gain readings, we have been utilizing the USRP B210 as an end-device to estimate and compare the gain of the beamformed signal with a non-beamformed one.

Beamforming in Indoor Environments

To test the power gain of our system, we have deployed experiments in both a corridor and a meeting room environment. Figure 3.16 shows that when using beamforming, we have an average power gain of 2.7 dB in the corridor environment and 2 dB in the meeting room, which is roughly a 200% gain compared to the baseline. The performance difference comes from the partial or total occlusion between the gateway and the end nodes when experimenting on through-the-wall transmission.

Beamforming in Outdoor Environments

To further characterize the performance of our system in terms of beamforming in a real campus environment, we conducted a measurement campaign in which, taking different positions, we computed the SNR value at the receiving side of the end-device. The experiment was designed to assess the power gain at the end-device




FIGURE 3.16: Empirical CDF power gain when using beamforming in a meeting room and corridor.

Position	Distance (m)	Relative altitude (m)	Beamforming Gain [dB]
North Bench	5	0	3.48
East Bench	20	0	2.49
LEAT Bench	50	0	8.87
Down	70	-8	-0.87
Up Bat1	90	0	3.68
Templiers1	130	14	-2.36

TABLE 3.4: The power gain using beamforming according to the distance from the gateway and the elevations for LoS Experiments.

when applying beamforming at the gateway. We tune the Tx gain at the gateway and the Rx gains at End-Device to avoid signal clipping.

LoS To estimate the system efficiency in applying spatial filtering, we use the same outdoor line-of-sight testbed described earlier, see Figure 3.13. As seen in Table 3.4, the mean power gain for Snipe is around 2.54 dB. This performance is comparable to the corridor’s results. Given that our array is one-dimensional, the difference in level influences negatively the AoA accuracy, and thus, the beamforming gain drops and becomes negative. This limitation can be mitigated by using a 2D antenna array for both estimating the AoA and performing spatial filtering.

Near-line-of-sight and NLoS The goal here is to show we can get decent beamforming performance even under blockage and lack of line-of-sight conditions. For that purpose, we use the same testbed used earlier for localization, see Figure 3.14, where the gateway is placed at a vantage point, and either cars, buildings, or foliage partially or entirely obstruct the end-devices. We have shown earlier that our localization system sustains an additional 10° error on average. However, the results of the experiment, see Table 3.5, shows that even when the AoA is slightly off, we have large power gains using beamforming with an average of around 6.87 dB.

Position	Distance (m)	Estimation AoA Error [°]	Beamforming Gain [dB]
NLoS 1	20	23,7	14.6
NLoS 2	42	15,7	4.64
NLoS 3	25	25,9	0.18
NLoS 4	32	41,4	7.97

TABLE 3.5: The power gain using beamforming according to the distance from the gateway for NLoS Experiments.

3.7 Related Works

Range extension LoRa range extension has been the focus of multiple research projects trying to increase the decoding capabilities at the gateway level for the uplink while achieving a low power consumption by end-devices [52, 49]. A system with a single antenna gateway that enables collaboration between end-devices was proposed to overcome the challenges rising in high-density deployment [52]. However, this system requires a dense deployment of end-devices and a synchronization mechanism between them. The authors have extended the latter system in Charm [49], aiming to enhance the battery life of client devices and the uplink coverage of LPWANs in a sizeable urban deployment. The system is built in such a way that if the signal is decodable at any individual gateway, then these weak signals coming from multiple gateways are coherently combined in the cloud. However, this approach is only suitable for private operators, as it involves costs to maintain a cloud infrastructure and deploy geographically distributed gateways. In comparison, the Snipe system proposed in this chapter uses multiple antennas on a single gateway system to extend the LPWAN range for both the uplink and downlink transmissions. We combine signals coming from multiple antennas in the uplink communication, which increases the SNR of the signal and improves the downlink communication through beamforming[138, 60] based on the target signal’s direction of arrival.

Localization LoRa geolocation using the signal of opportunity methods has been studied in the context of TDoA to infer the position of the end-device [53] and [106]. The authors in the first paper [53] propose a LoRaWAN geo-positioning system capable of exploiting transmitted packets to compute the current position without GPS or global system for mobile communications (GSM), allowing low-power consumption. The authors use multilateration, which consists in computing the different TDoAs from the coordinated universal time (UTC) timestamp of the packets received by the end-device. This technique requires a fine-grained synchronization of gateways with a reported mean accuracy around 100 m. The authors in the second paper[106] quantify the TDoA geolocation performance in a publicly deployed LoRa network for different scenarios (walking, cycling, and driving). The authors propose using an additional tracking improvement algorithm to enhance tracking accuracy, which considers the road map and movement speed with an overall median error of the TDoA location output was around 200 m, and in 90 percent of the cases, the error is less than 500 m overall SFs and trajectories. Other authors [67] propose the design

of an energy-efficient LoRa GPS wristband tracker for dementia patients by leveraging the low energy consumption of the LoRa communication module and the GPS duty cycling strategy. The trackers use a GPS module to retrieve geolocation from satellites and a LoRaWAN module to transfer real-time locations back to a central server. Our work does not rely on synchronization between gateways reducing the complexity and provides an accurate estimate of the position [125]. Our system is energy efficient as all we offload the processing for the position estimation to the gateway, thus avoiding draining the end-device batteries. Our technique requires only one station, allowing it to be easily deployed in different environments; a flexibility lost when using TDoA receivers, which needs at least two stations to estimate the position, three for 2D localization. Our localization technique based on AoA uses low data rates to send to the network server, while TDoA systems require a high data rate for sending the collected waveforms to a remote server. TDoA systems are prone to signal de-correlation between receivers imposing costly frequency and time coherence tracking, which is not the case of AoA based localization. Given that TDoA is based on signal cross-correlation, a signal with high periodicity could induce erroneous estimations. The time to fix is orders of magnitude more significant than AoA. Some AoA systems support concurrent geolocation of many frequency separated signals, which is difficult for TDoA requiring higher data transmission requirements.

3.8 Conclusions

This chapter presents a novel MIMO-LPWAN system that not only extends the range between a gateway and an end-device but also enables a precise LoRa end-device localization system. We have shown that leveraging spatial filtering techniques and coherent combining across the signal impinging on a linear antenna array it is possible to increase the decoding capabilities at both the gateway and the end-device. We show that it is possible to achieve an accurate angle of arrival estimation by adopting model-based MIMO radar techniques for a signal distributed across time and frequency (Chirp). Therefore, allowing the estimation of the AoA information using a MIMO gateway would complement other already deployed ranging techniques for LoRa and thus strengthen the localization. Given that we are dealing with LoRa, which is a low-cost oriented IoT solution, providing a localization solution that builds upon this limitation is crucial. Therefore, we privilege using the estimated AoA in two distinct scenarios for LoRa localization:

- **Single homing:** Trilateration is not possible since a single gateway is used to serve the deployed end-nodes. By adding ranging capabilities (Time of arrival) to the AoA information, a single gateway would be able to localize the emitter, which would require synchronization between the Tx/Rx. Or in a coarser-grained manner, we could utilize a ranging system based path loss between an end-node and the gateway. Therefore, the level of the received signal power at the gateway would be used to approximate the distance; however, this requires prior knowledge of the Tx power, the antenna gain, AGC (automatic gain control), etc.

- Multiple-station homing: Multilateration based techniques (TDoA) are adopted by the LoRa community. AoA would be complementary to TDoA, which would create a hybrid AoA/TDOA system. This hybrid approach requires fewer stations than separate AOA or TDOA systems while preserving optimal area coverage, thus reducing budget requirements. This would mean computing the hyperbolic line of time difference from TDoA and the azimuth line from the AoA system, and therefore the position of the target would be the intersection of these two lines. For this case, we would need at least one gateway with AoA/TDOA, and all the other ones with TDoA capabilities.

Future work seeks to enhance AoA accuracy by augmenting spatial resolution using synthetic aperture radar and virtual antenna arrays. With the higher number of antennas, this approach would allow mitigating the multipath effect by decorrelating the signal of interest from its reflected copies using spatial smoothing techniques on a broader virtual array. Detailed information on how to reproduce the experiments is available at https://github.com/naoufal51/MIMO_LoRa.git.

Chapter 4

Robust Passive Localization Using Wi-Fi CSI: A Data-Driven Paradigm

4.1 Introduction

Passive localization technology can track the position of objects without requiring a specialized instrument on them. It represents a fundamental primitive for many ubiquitous computing applications. Some use cases include: (i) *Mobile healthcare*. The intensity and quantity of motor activities can indicate the healthiness of a person and assess the risk or the recovery from chronic diseases such as stroke and dementia [70, 121]. Therefore, by enabling continuous, non-intrusive passive motion tracking, the healthcare agents can remotely assess elder people’s health without hospitalization. (ii) *Retail 2.0*. By mining the location and sojourn time distribution of customers in a retail store, the retailer can predict the demands and realize a more precise inventory control [58]. (iii) *Public safety*. The real-time estimation of crowd distribution can enable law enforcement personnel to plan their dispatch. The ability to locate lives inside buildings can also enable more effective disaster recovery.

In most of these applications, the well-explored camera and computer vision technologies are not applicable because of occlusion or privacy restrictions, whereas RF-based passive localization becomes a niche solution. Over the past two decades, substantial research has been devoted to active RF localization using radio devices [151]. Passive localization leverages the RF signals reflection off target objects, much like radar tracking. It has been well established that the spatial resolution of RF signals is proportional to $\frac{\lambda d}{B\mathcal{A}}$ [101], *i.e.*, finer resolution can be achieved with shorter signal wavelength λ , a wider sampling bandwidth B , shorter distance d from transmitter/receiver to the target, and larger antenna aperture \mathcal{A} . Practical localization algorithms manipulate B and \mathcal{A} to achieve high precision.

For example, WiVi [5] used an inverse synthetic aperture radar (ISAR) algorithm to create a virtual antenna array, thus increasing \mathcal{A} and realizing sub-meter-level through-wall human tracking. WiTrack [6, 4] used a multi-GHz frequency-modulated continuous-wave radar (FMCW) radar (with large B) and multiple antennas (to increase \mathcal{A}) to resolve the locations of multiple human subjects.

This work has been carried out under the guidance of Prof. Xinyu Zhang at UC San Diego and with the help of Yifan Huang and Changhan Ge.




FIGURE 4.1: (a) LoS and (b) NLoS in passive localization.

In this chapter, we propose a challenging task for wireless sensing: Can we achieve similar spatial resolution as the wideband MIMO radar, by using low-profile COTS Wi-Fi radios? Unlike radar, COTS Wi-Fi radios operate at a narrower sampling bandwidth (20 MHz) with few antennas. They do not maintain phase coherence across channel samples (*i.e.*, packet transmissions and channel estimations), which precludes the use of coherent processing techniques like ISAR to increase \mathcal{A} . Essentially, meeting this grand challenge would entail the design of *super-resolution algorithms* that circumvent the RF resolution law.

A few systems [149, 107, 108] have started to address the challenge. These systems use Wi-Fi CSI to extrapolate the geometries of the propagation paths, including time-of-flight (ToF), AoA, AoD, *etc.*. They adopt super-resolution parameter estimation algorithms (*e.g.*, MUSIC[124] SAGE[55]) to jointly obtain the path metrics. However, these closed-form models have to make a fundamental assumption[108]: A first-order reflection path exists between the transmitter, target, and receiver (Fig. 4.1(a)), although the path may be attenuated by weak obstacles. Many obstacles (*e.g.*, brick/concrete walls, metal furniture) lead to non-line-of-sight (NLoS) conditions (Fig. 4.1(b)). Geometries of the resulting high-order reflection paths become highly dependent on the environment, which apparently cannot be resolved by closed-form models.

In lieu of the fundamental limitations of geometrical models used since the radar era, we propose to adopt a data-driven paradigm for passive localization, based on *RFLoc*[114]—an open-source data infrastructure we create. We use a COTS Wi-Fi 802.11ac radio to collect CSI samples in 4 distinct environments, with a variety of LoS/NLoS conditions and up to 4 human subjects walking simultaneously. Meanwhile, we use a depth camera sensor to label the target location. The *RFLoc* dataset comprises *1.1 million location labels* each associated with one CSI matrix.

Despite the stunning growth of deep learning (DL) tools, it is non-trivial to apply them to passive localization. We propose the following specific challenges.

Challenge 1: Designing DL models for wireless signals. DL architectures are

essentially designed to efficiently represent the distribution of high-dimensional features [62]. Whereas the majority of existing DL applications focus on images comprised of a dense array of pixels, wireless signals are intrinsically sparse due to a few "sensors" (antennas). In addition, rather than modeling the spatial distribution of pixels, DL based wireless sensing needs to represent the propagation environment, *i.e.*, distribution of the complex channel gains which can be affected by reflector locations, activities, *etc.* Such distribution may indirectly imprint the ToF, AoA, AoD metrics used in closed-form geometrical models. Therefore, we have to tailor classical DL mechanisms for wireless sensing. Examples include new regularization schemes (*e.g.*, task-specific loss function), training optimization (*e.g.*, adversarial training), CSI sequence modeling (*e.g.*, long short-term memory (LSTM) to model channel correlation), and structured probabilistic models to incorporate domain knowledge and more interpretable structures.

Challenge 2: Super-resolution passive localization in true NLoS conditions. To overcome the LoS limitations of closed-form models, data-driven solutions need to implicitly capture the geometrical relation between the target's position and the sophisticated high-order reflections (Fig. 4.1(b)). The corresponding DL models have to work under the Wi-Fi hardware constraints (\mathcal{A} and \mathcal{B}), while achieving even higher spatial resolution than first-order geometries [149, 107, 108, 77]. This is challenging, but feasible, because DL models can incorporate prior knowledge of the feature distribution (*e.g.*, continuity of reflector surface, diffraction/refraction effects) by training over massive data.

Challenge 3: Weakly supervised or unsupervised learning for environment-independent passive localization. Ambient multipath reflections can strongly affect CSI features associated with a target location. Even in LoS conditions, closed-form solutions' performance can deviate by 20%-40% depending on the reflection environment [108]. Under NLoS, even for the same target location, the CSI can vary drastically depending on the ambient reflectors' geometry, orientation, and reflectivity. Therefore, the passive localization model must be robust across different environments. Unfortunately, one cannot afford to collect location-labeled CSI samples and run supervised training for each environment, and even the same environment may change rapidly because of furniture placement, human activities, *etc.* So weakly supervised training (which requires a minor recalibration), or even unsupervised adaptation to a new environment, is more preferable. Recent work in Wi-Fi-based activity sensing [155] adapted transfer learning [62] to ease the labeling overhead. But they still need non-trivial ground-truth labels (> 500 samples) in each new environment.

In this work, we address challenges 2 and 3. We propose a DL framework, called Deep Passive Localization (DPLoc), that uses a deep convolutional neural network (CNN) to infer locations given CSI inputs, and an unsupervised generative model to automatically reconcile unlabeled CSI samples in a new environment with those labeled ones in a benchmark environment. In this way, DPLoc only needs a one-time CSI collection and location labeling in the benchmark environment, and can *automatically adapt to the new environment with zero additional labeling efforts*. We have implemented DPLoc and conducted extensive experiments based on the RFLoc dataset.

The results demonstrate that, compared with a baseline without environment adaptation, DPLoc reduces localization error by around 25%, consistently maintaining *sub-meter precision* when locating multiple subjects across a variety of LoS/NLoS settings. Under the same NLoS setup, the location error escalates to *around 6.7 m* when using state-of-the-art closed-form models [108]. We can summarize the main contributions of this work: (i) We create the RFLoc dataset as an open source campaign for passive RF localization. This effort will probably help establish standard practices, and encourage more reproducible research in wireless sensing, just like the ImageNet challenges [48] did in computer vision. (ii) We address the challenges via a novel DL framework, and validate its effectiveness through comprehensive experiments.

4.2 Related Work

Passive localization technology roots in the vast literature of radar tracking, but mostly targets indoor and civilian use cases. Early work, such as WiVi [5], focused on through-wall tracking. The main challenges involve removing wall-induced cluttering effects through static signal cancellation and improving tracking precision using ISAR. WiTrack further enables multi-people tracking [6, 4] using a wideband (multi-GHz) multi-antenna FMCW radar with high spatial and temporal resolution. WiDeo [77] uses a full-duplex radio to cancel the direct LoS signals, rendering the target reflections more prominent. These systems all require specialized wireless transceivers, such as radar or software-radio, to evade the limitations of commodity communication-oriented packet radio devices such as Wi-Fi.

COTS Wi-Fi devices have limited sampling bandwidth (and hence ToF or range resolution). However, recent work has shown that, by combining angular estimation from multiple antennas, and Doppler frequency estimation across multiple subcarriers, it is possible to achieve "super-resolution", in both spatial and temporal domains. For example, mDTrack and Widar [149, 107, 108] cast the problem as a multi-dimension joint parameter estimation and used the SAGE algorithm to estimate ToF and AoA simultaneously. The target location can then be computed as a point in the 2D space that satisfies both parameter constraints. Unlike radar [4], these systems can only track a *single subject*.

All the above approaches are contingent on the LoS setup. Although some experiments involve blockage effects [5], the obstacle is typically a drywall that attenuates (by a few dB at most) but never eliminates the LoS path. Alternative models, such as radio tomography [143, 147], suffer from the same limitation and often require a much denser deployment of radio devices. Fingerprinting based approaches [148, 127] do not rely on such geometries, and can implicitly account for the impacts of high-order reflections that detour NLoS.

But such approaches face a severe difficulty in multi-subject tracking since the fingerprinting load increases exponentially with the number of location spots. Fingerprinting has to be re-initiated upon minor environment changes.

Machine learning models have been adopted recently for RF localization. For example, deep learning algorithms can automatically extract robust features from

CSI fingerprints [145, 144] for locating active radios. RFPose [160] localizes prominent limb joints of the human body by using vision labels to train RF signals. Although RFPose achieves remarkable precision, it still relies on a specialized wide-band MIMO radar and assumes a LoS signal path (possibly attenuated by obstacles) exists.

Parallel to RF localization, image-based target tracking has been well studied in the past decades. Approaches include 2D-to-3D matching [119], CNN-based object localization [119] based on a massive ImageNet dataset [48], *etc.* In contrast, RF signals are intrinsically sparse and ill-shaped because of sophisticated multipath reflections. Some deep learning mechanisms we adopt in this chapter (*e.g.*, CNN) are well known in image, speech, and natural language processing. However, to our knowledge, we are the first to design a generative model to reconcile the heterogeneous NLoS reflection effects, making a single passive tracking model work across different environment without additional labeling.

4.3 The RFLoc Dataset

It is critical to have a large CSI dataset with location labels to enable deep learning-based passive tracking models. However, collecting fine-grained position labels comes with scaling and accuracy challenges. Using manual labeling would be too tedious or nearly impossible to achieve with limited guarantees on the label’s quality.

To overcome these constraints, we implement a system to automatically collect labeled CSI data.

Automatic dataset generation. Our CSI data are collected using a COTS Wi-Fi device, equipped with the 4×4 802.11ac Wave 2 wireless chipset QT3840BC [40] from Quantenna. Quantenna provided us a development version of the device, along with APIs to access the physical (PHY) layer information. Unlike the Intel 5300 or Atheros ath9k devices widely used for localization research [149, 107, 108], the Quantenna radio provides CSI for all the 48 data subcarriers on each 20 MHz Wi-Fi band, and across all 4 antennas (hence 192 CSI values per sample).

To autogenerate fine-grained location labels, we use the depth sensor of Kinect v2 [159], with Brekel Pro Body v2 [39]—a motion capture software driver that logs an accurate skeleton representation of each human subject. Typically, the skeleton consists of 62 body joints; each mapped to a 3-D Cartesian coordinate system. In RFLoc, we only extract the hip coordinates as location labels (which would correspond to the center of mass of the human body). To synchronize the data collection, we send a triggering signal in parallel from a local server to start the CSI and position measurement.

Although the optical-based system can generate highly accurate location labels, we had to take a few additional steps to preserve the quality of our ground-truth position labels: (i) We collect data in the evenings as the depth sensor performs poorly under sunlight interference. (ii) The subjects need to wear light color clothes; Otherwise, the motion capture system cannot create the skeleton frames.

(iii) In the multi-subject scenario, we have to avoid inter-subject crossing; Otherwise, the occlusion effect will fail the depth perception. However, note that one




FIGURE 4.2: The DPLoc dataset is built using Quantenna Wi-Fi radio for CSI collection and the Kinect v2 for location labeling which provides both depth maps and hip coordinate as location label for each CSI sample.

subject may still block the other from the perspective of the radio devices placed at different locations and possess a different field-of-view than the Kinect sensor. (iv) Brekel performs random ID allocation and has difficulties in maintaining the skeleton data integrity across sessions. During the pre-processing of skeleton data, we identify the initial position of each skeleton and apply an id reassignment.

Data format and diversity. Our dataset spans four propagation environments (yoga room, corridors, meeting room, and outdoor space) with a mix of LoS and NLoS conditions and up to four human subjects moving simultaneously. Fig. 4.3 illustrates the corresponding floor plans and setup for data collection. To create the NLoS, we use a large metal board to block the LoS path between the target and the Tx/Rx (Fig. 4.6). Both the Tx and Rx radios are co-planar, with the Rx located close to the Kinect sensor. We fix the Tx-Rx distance in all environments. We maintain the same Cartesian coordinate system with the xy -plane parallel to the ground, the receiver position is our origin (0,0), and the transmitter position is (0,5.4).

For each CSI sample, we log the target position, and a depth map of the experimental scene, which may be instrumental for vision-assisted deep learning models for RF sensing [160]. For each environment, we also log a set of baseline samples in the absence of human subjects. Our RFLoc dataset comprises approximately *1.1 million location labels and CSI samples*. Table 4.1 summarizes the structure of our dataset.




FIGURE 4.3: Floor plans of experiment environments.

4.4 Data-Driven Passive Localization in NLoS and Heterogeneous Environments

Our DPLoc framework consists of two distinct modules. The *position inference module* produces a position estimation given a phase matrix extrapolated from a CSI sample. It comprises a deep neural network, trained on a single benchmark environment with location labels. The *environment adaptation module* runs an unsupervised generalization network. It takes CSI samples in a new (unseen) environment and processes the samples to remove environment heterogeneity. Its output is fed into the inference module to obtain a location estimation.

4.4.1 Position Inference module

CSI phase sanitization. As in geometrical model-based approaches[107, 108, 149], DPLoc exploits the phase components from CSI as input. In general, the CSI in a multipath environment can be expressed by a channel transfer function:

$$H(f) = \sum_{n=0}^{N-1} \rho_n e^{j\phi_n} e^{-j2\pi f\tau_n} \quad (4.1)$$

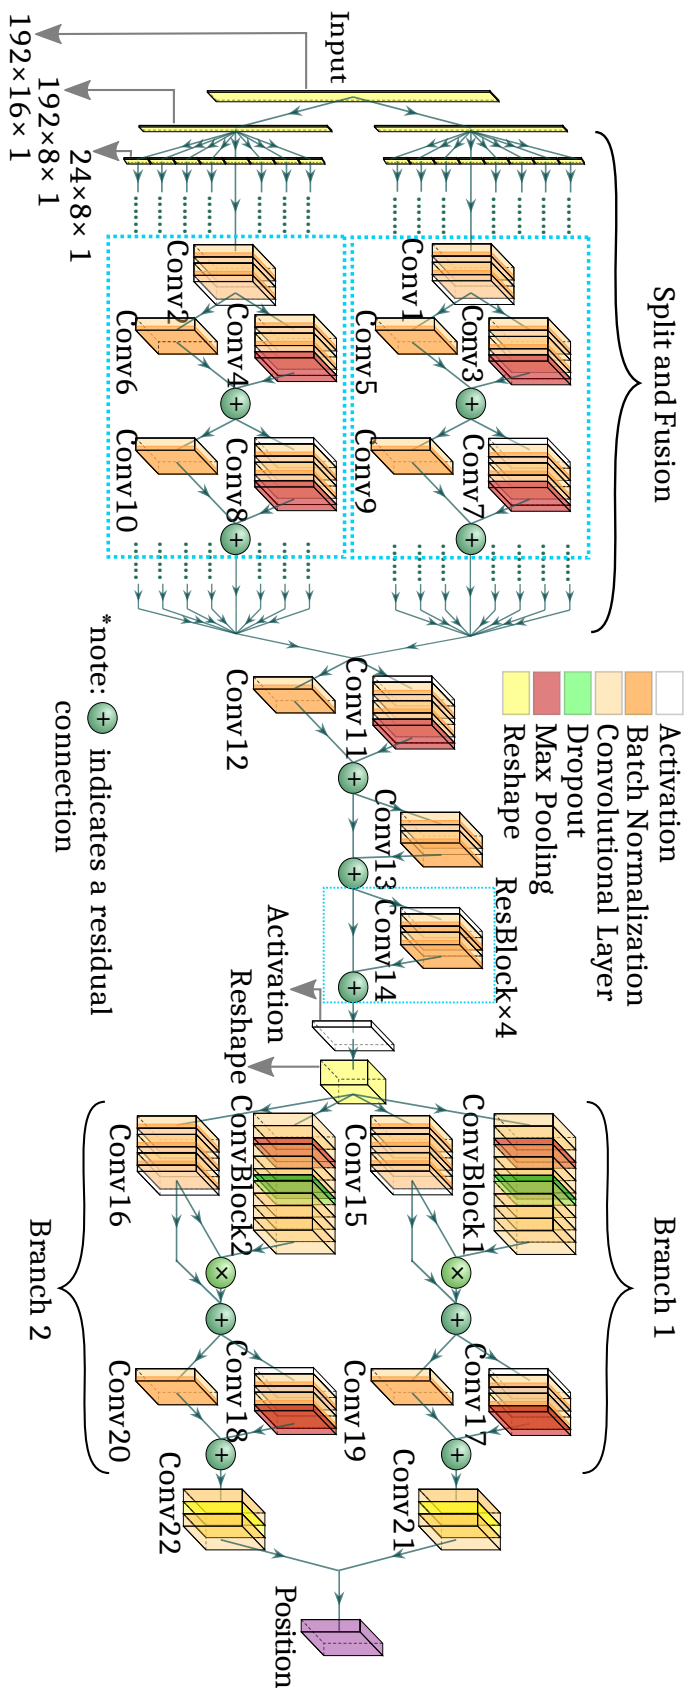


FIGURE 4.4: Network Design of the Inference Module

Environment	Obstacles	LoS		NLoS	
		# of People	# of Sample	# of People	# of Sample
Yoga Room	No	1	72720	1	—
		2	90540	2	54300
		3	—	3	72600
		4	—	4	36460
	Yes-1	2	54600	2	144800
	Yes-2	1	18130	1	18130
		2	36260	2	72340
3		18130	3	18130	
Corridor	No	1	118170	1	—
		2	36240	2	36240
		3	—	3	36220
Meeting Room	Yes	2	36220	2	36360
		3	—	3	36260
Outdoor	No	2	36180	2	72440
Total number of sample		1,115,210			
Total Video Time		9.38h			

TABLE 4.1: Structure of the RFLoc dataset.

	Conv1-2	Conv3-4	Conv5-6	Conv7-8	Conv9-10	Conv11	Conv12
Num Filters	64	128	128	256	256	4096	1024
Filter Size	3×3	3×3	3×3	3×3	3×3	3×3	3×3
Context	12×4	12×4	6×2	6×2	3×1	3×1	3×1
	Conv13	Conv14	ConvBlock1-2	Conv15-16	Conv17-18	Conv19-20	Conv21-22
Num Filters	1024	1024	1024	1024	1024	1024	3
Filter Size	3×3	3×3	3×3	3×3	3×3	3×3	1×1
Context	3×1	3×1	8×6	8×6	8×6	4×3	1×1

TABLE 4.2: Convolutional layers used in our location inference model.

where τ_n denotes the delay of the n -th path, f the frequency, and $\rho_n e^{j\phi_n}$ the complex gain of each propagation path. Thus, the target can be characterized as a multi-path cluster with a complex gain. In MIMO-OFDM WiFi, a separate $H(f)$ exists for each Tx-Rx antenna pair and each subcarrier.

As the phase data changes in time, our training data will suffer from offsets. To mitigate this problem, we use the relative phase data between antenna elements for our model as the inter-element spacing is fixed.

Specifically, for subcarrier k :

$$\omega_{ij}(k) = \angle(H_i(k)H_j^*(k)) \quad (4.2)$$

In addition, each RF chain on the MIMO radio may have a different constant ψ value added to the measured phase. So, the measured CSI on RF chain i is:

$$H_i(f) = \left(\sum_{n=0}^{N-1} \rho_n e^{j\phi_n} e^{-j2\pi f \tau_n} \right) e^{j\psi_i} \quad (4.3)$$

Fortunately, the phase offset remains stable [107] for a wireless card till the next re-calibration or reset. Therefore, we use a calibration signal to estimate this offset and subtract it from our relative phase data, in a similar way as [150].

Network architecture design choices and optimization. Fig. 4.4 illustrates the architecture of the position inference module. The input is a 192×16 matrix, representing the sanitized phase values of the 4×4 MIMO channels for each of the 192 subcarriers as we use CSI on 80MHz bandwidth. The input is further split into 16 parts in both the frequency (subcarrier index) and space (antenna index) dimensions, where each split is a 24×8 matrix (Fig. 4.4). Using smaller dimension inputs reduces the number of trainable filters in the deep neural network and hence the training time (Sec. 4.5). We feed each 24×8 matrix through convolutional neural blocks with residual connections. Each convolutional layer comprises learnable filters, so the result is a matrix of filter responses to the sanitized relative phase information.

Closed-form models [149, 107, 108] often extrapolate the relation between target location and relative phase. Here we expect the deep network to do so automatically through supervised training. Intuitively, using a large number of convolutional layers to learn more complex features would yield a better estimation precision. However, this comes with the vanishing gradient problem, where the gradient value becomes too small to update the filter weights. Therefore, we adopt the residual connections [71], which reduces the information loss through the convolutional layers. Our convolutional residual blocks apply multiple convolution operations on their input while preserving their original dimension. Then the blocks' inputs are added to the outputs of the subsequent convolutions (the \oplus in Fig. 4.4).

As the network goes deep, the number of convolution operations snowballs, which challenges trainability. To mitigate this problem, we adopt depth-wise separable convolution layers [43], which reduces the convolution operations and training time, while maintaining accuracy. To stabilize the training process of the inference network, we use batch normalization [74] in the residual blocks to provide a more stable gradient decent during the back-propagation operation and allowing a better convergence of the network. Our residual design uses 3×3 stride-2 depth-wise separable convolution layers with ReLU activation [100] and a batch normalization layer (Fig. 4.4).

After processing each matrix, we recover a single stream provided to the subsequent layers using a concatenation layer along the matrices' third dimension. The output of this fusion operation is first processed through 4 residual blocks to learn more complex features. The result is then propagated in two branches (Branch1 and Branch 2 in Fig. 4.4). Each branch is composed of two convolutional residual blocks. The first block applies the residual attention method [142] to help identify the critical features in the input. This method relies on skip connection with multiple convolutional layers. As for the second block, it has the same architecture as the described residual blocks. To shape our output to the dimension of the location label, we use two stride-2 convolutional layers with ReLU activations. Finally, as our position estimation system deals with a regression problem on unbounded continuous values, the last layer uses a linear activation. We summarize the output dimension of each convolution layer in Table 4.2.

Loss function for training the position inference network. To train the position




FIGURE 4.5: Network design for environment adaptation.

inference module, we adopt a loss function that aims to minimize the Euclidean distance between the estimated positions and the ground-truth labels:

$$\mathcal{L}_{\text{position}} = \sum_{i=1}^n \mathbf{M}(n) \sqrt{\sum_{i=1}^n (p_i - \hat{p}_i)^2} \quad (4.4)$$

where p_i and \hat{p}_i are the ground-truth and the estimated location coordinates, respectively; while n represents the number of dimensions. To accommodate the multi-subject scenarios and to minimize the impact of missing/corrupted labels (due to occasional Kinect sensor failure), we apply a binary mask function $\mathbf{M}(n)$ that rejects imperfect labels. The label matrix used by our loss function is sparse with a dimension corresponding to the maximum number of subjects supported by the model. Thus, the mask's role is to exclude missing position labels in the model's calculation. Otherwise, the training process would diverge.

4.4.2 Environment Adaptation Module

To ensure DPLoc's model applies to a new environment with zero additional position label collection, we design an unsupervised generative model.

Each environment (also called a domain) can be characterized by the many-dimension distribution of CSI, conditioned on the distribution of target locations. Our generative model allows the joint correspondence between distributions of the original (source) and new (target) environments.

Let $csi^A \sim p^*(csi^A)$ and $csi^B \sim q^*(csi^B)$, where $p^*(csi^A)$ and $q^*(csi^B)$ are the empirical distributions of source domain A and target domain B , respectively. Our goal is essentially to convert a sample from the domain A into a sample that matches

the distribution of domain B , and vice versa:

$$csi^B = \mathbf{Gen}_\Theta(csi^A), csi^A \sim p^*(csi^A) \quad (4.5)$$

$$csi^A = \mathbf{Gen}_\Psi(csi^B), csi^B \sim q^*(csi^B) \quad (4.6)$$

In this way, whenever a CSI sample is collected in the new environment, we can use the generative model to map it to the original environment and use the pre-trained inference model (Sec. 4.4.1) to predict a subject's location.

As illustrated in Fig. 4.5, our generative model extends the GAN [63][161]. The traditional GAN model contains a generator and a discriminator. The generator network relies on the feedback of the discriminator network, which has access to both the ground-truth samples and the "fake" samples created by the generator. We can consider the discriminator as a learned loss function, that outputs $D(real) = 1$ when the sample is from the real dataset and $D(G(fake)) = 0$ otherwise. We train the generator to produce new samples that are good enough to be classified as genuine by the discriminator, whereas the discriminator aims to maximize its discrimination accuracy. Our design comprises two GANs, for the source and target domain, respectively. The generator in GAN uses its CSI input to produce CSI samples realistic enough to be mapped to the target environment, while the discriminator is in charge of detecting fake CSI samples from real ones. The overall process for producing the training samples follows:

$$csi^B = G^B(E^A(csi^A)) \quad (4.7)$$

where $E^A(\cdot)$ is a neural network encoder that produces feature maps from CSI data for a generator G . The superscript A indicates the domain it operates on. While GANs are often used for image-to-image translation, they are hard to optimize and usually suffer from instability.

Empirically the discriminator is usually easier to converge, thus often overpowering the generator, which leads to unsatisfactory samples.

To mitigate this problem, [29] proposed to use the Wasserstein distance. Given this distance is hard to optimize, the authors use the Kantorovich-Rubinstein duality to approximate this metric [140]. Thus, we adopt it for training the GANs:

$$\mathcal{L}_{\text{GAN}} = \mathbb{E}_{csi \sim \mathbb{P}_{\text{real}}} [D(csi)] - \mathbb{E}_{\tilde{csi} \sim \mathbb{P}_{\text{Gen}}} [D(\tilde{csi})] \quad (4.8)$$

Intuitively, this metric represents the distance between the distribution of our CSI data \mathbb{P}_{real} and the distribution of the generated samples \mathbb{P}_{Gen} . Instead of generating a Boolean decision, the discriminator in Wasserstein GAN gives a score ranging between -1 and 1, which is seen as how real or fake the input samples are. By using this loss function, there is a 1-Lipshitz condition on the discriminator D :

$$|D(x_1) - D(x_2)| \leq |x_1 - x_2| \text{ for all } x_1, x_2 \in \mathbb{R}. \quad (4.9)$$

The derivative of the discriminator function is bounded. This 1-Lipshitz continuity condition was at first fulfilled by clipping weights to a fixed range.

$$w_{\text{clipped}} = \text{clip}(w, -c, c) \quad (4.10)$$

However, this method suffers from vanishing gradient when the clipping value is small, and it increases the training time when the value is large [140]. Because of this problem, we adopt the Wasserstein GAN with a regularization term called gradient penalty [65], used for constraining the gradient to respect the 1-Lipshitz condition. Thus the critic loss function is as follows:

$$\mathcal{L}_{\text{GAN}_c} = \mathbb{E}_{\tilde{csi} \sim \mathbb{P}_{\text{Gen}}} [D(\tilde{csi})] - \mathbb{E}_{csi \sim \mathbb{P}_{\text{real}}} [D(csi)] + \lambda \mathbb{E}_{\hat{csi} \sim \mathbb{P}_{\hat{csi}}} [(\|\nabla_{\hat{csi}} D(\hat{csi})\|_2 - 1)^2] \quad (4.11)$$

where the first two terms correspond to the Wasserstein loss; the last term is related to the gradient penalty where $\hat{csi} = \rho csi + (1 - \rho)\tilde{csi}$ with ρ is uniformly distributed on $[0,1]$. We apply the gradient penalty to samples produced by a combination between the distribution of the real data and the one produced by the generator.

Reconstruction consistency between different environments. Ideally, we want the generative process to produce a one-to-one mapping between the two domains, *i.e.*, the original input can also be recovered from its converted version. However, with the adversarial loss function alone, the generator may produce a CSI in the target domain even given a random CSI input. Inspired by the image-to-image translation scheme [161], we tackle this problem using a loss term that controls the cycle consistency:

$$\mathcal{L}_{\text{cc}} = \left\| G^A(E^B(\tilde{csi}^B)) - csi^A \right\|_1 \quad (4.12)$$

The previous loss function conserves the one to one mapping between a CSI sample and its reconstructed version after domain translation. But we also need to preserve the identity mapping between the input sample and the output of its corresponding generator without domain translation. Otherwise, the translated samples would not be completely consistent with the target domain [134, 161]. Therefore, we add the following regularizer into the loss function to penalize wrong identity mapping:

$$\mathcal{L}_{\text{ae}} = \left\| G^A(E^A(csi^A)) - csi^A \right\|_1 \quad (4.13)$$

Retraining the position inference model.

After mapping the samples from the original environment to the target environment, we retrain the position inference model against these generated samples to learn the characteristics of the target domain. This process could be likened to a data augmentation process that improves the generalization capability of the network and reduces the estimation error when making predictions on new samples from unseen environments. This requires that the generated samples used for retraining should conserve the position information of the original CSI samples. As we have access to the labels from the original domain, we feed the generated CSI samples to our pre-trained model to produce position estimations further used in our loss function:

$$\mathcal{L}_{\text{pc}} = \lambda_3 \mathcal{L}_{\text{inference}}(\tilde{csi}^B) \quad (4.14)$$

Algorithm 2 Training algorithm for environment adaptation.

```

1: for  $k$  training steps do
2:   for  $t = 1, \dots, n_{\text{Disc}}$  do
3:     Process data  $csi^A$  and  $csi^B$ 
4:     Update  $D^A$  with  $\nabla_{D^A}[\mathcal{L}_{D^A}(D^A, (E^B, G^A))]$ 
5:     Update  $D^B$  with  $\nabla_{D^B}[\mathcal{L}_{D^B}(D^B, (E^A, G^B))]$ 
6:   end for
7:   Process data  $csi^A$  and  $csi^B$ 
8:   Update  $G$  with  $\nabla_G[\mathcal{L}_{\text{global}}(E^A, E^B, G^A, G^B)]$ 
9:   Update  $Pos_{\text{est}}$  with  $\nabla_{Pos_{\text{est}}}[\mathcal{L}_{\text{position}}(E^A, G^B)]$ 
10: end for

```

where λ_3 is a weighting parameter that sets the importance of the regularization and $\mathcal{L}_{\text{inference}}$ is the Euclidean loss function described in Sec. 4.4.1.

In sum, after the training process, our complete model will concatenate a generator used to produce new CSI samples in the original domain and the retrained inference model (Fig. 4.5). The global objective function of the complete model is:

$$\begin{aligned}
& \max_{D^A, D^B} \min_{E^A, E^B, G^A, G^B} \mathcal{L}_{\text{pc}}(E^A, G^B) + \lambda_4 \mathcal{L}_{\text{GAN}}(E^A, G^B, D^B) \\
& + \lambda_4 \mathcal{L}_{\text{GAN}}(E^B, G^A, D^A) + \lambda_5 \mathcal{L}_{\text{ae}}(E^A, G^A) + \lambda_5 \mathcal{L}_{\text{ae}}(E^B, G^B) \\
& + \lambda_6 \mathcal{L}_{\text{cc}}(E^A, G^B, E^B, G^A) + \lambda_6 \mathcal{L}_{\text{cc}}(E^B, G^A, E^A, G^B)
\end{aligned}$$

where the weighting parameters represent the importance of each regularizer (Sec. 4.4.3).

Network architecture design and optimization. Within the environment adaptation module, we use convolutional layers with no fully connected layers, which has proven to stabilize the training procedure [109]. As illustrated in Fig. 4.5, the encoder part uses convolutional blocks to compress the CSI input into a latent space, and the generator uses transposed convolutional blocks to recover the input. The transposed convolutional layers can upsample input to the desired dimension, [50]. We adopt five 5×5 2D convolution layers at the encoder level, and five of their 5×5 transposed counterparts at the generator level, with a ReLU activation and a batch normalization layer interleaved between every two convolutional layers.

The discriminator architecture is based on five 4×4 stride-2 convolutional blocks with an increased number of filters, ReLU activations for the first and intermediate layers and a *tanh* activation for the final layer. The final activation, which outputs values within $[-1, 1]$, is used with the Wasserstein loss function. As we are using gradient penalty, we avoid batch normalization layers in our discriminator design as it was shown to create a correlation between the samples of a batch and thus hindering the effectiveness of the regularization process [65].

4.4.3 Implementation and End-to-End Training

We implement our deep neural networks using Keras 2.2.4 with a TensorFlow 1.12.0 back-end. For training and testing, we use the Google Cloud Platform instances with



FIGURE 4.6: NLoS setup between a Tx and a subject.



FIGURE 4.7: Yoga room environment with obstacles.

NVIDIA Tesla K80 GPUs (12GB of video memory each) and Intel Broadwell processors. We first pre-train the inference network using a batch size of 64 samples and train with Adam optimizer [80] with a learning rate of $1e-4$ which is reduced after every epoch, and momentums of 0.90 and 0.99. As for the environment adaptation network, we reduce the learning rate for the pre-trained model to $5e-5$, and set the learning rate to $1e-4$ for the generators and discriminators with a batch size of 30. The training process of the discriminator is slow [65]. Thus, during each training iteration on a group of CSI data, the discriminator is trained for $n_{\text{Disc}} = 5$ iterations followed by a single generator and inference model training (see Alg 2). To avoid over-fitting; We chose a weight decay strategy in the form of an L2 regularization with a rate of 0.01 in the inference module.

The weighting hyper-parameters are empirically set to: $\lambda = 10$, $\lambda_3 = 0.1$, $\lambda_4 = 1$, $\lambda_5 = 1$ and $\lambda_6 = 10$.

4.5 DPLoc Evaluation

In this section, we validate the efficacy of DPLoc. By default, we use the 2-subject scenarios in the RFLoc dataset.

4.5.1 Effectiveness of the Inference Module

Input splitting strategy. We used two criteria to test the advantages of our input design: the estimation accuracy and the training time. We have tested multiple splitting configurations, which spans from using the raw 192×16 relative phase matrix to the 24×8 matrix produced by splitting the input space dimension by 2 and the frequency dimension by 8. From the results (Table 4.3), we observe our input splitting strategy provides a relative gain in position accuracy compared to other configurations. The advantage of this approach is clearer when it comes to training time optimization, which is reduced by almost 36%. Note that the batch time means the training time for each batch of samples.

Position estimation accuracy with the inference module alone. To test DPLoc’s position estimation capabilities, we have trained its inference network on LoS CSI data inside the empty yoga room. Table 4.4 summarizes the mean error, and Fig 4.8




FIGURE 4.8: Location error CDF when changing the layout in the same environment (yoga room).

Splitting Strategy	No split	Freq/2	Freq/4	Freq/8	Spatial/2	Freq/8 & Spatial/2
Position (cm)	20.01	20.42	20.33	19.99	19.74	19.29

TABLE 4.3: Split and estimation error on same validation set.

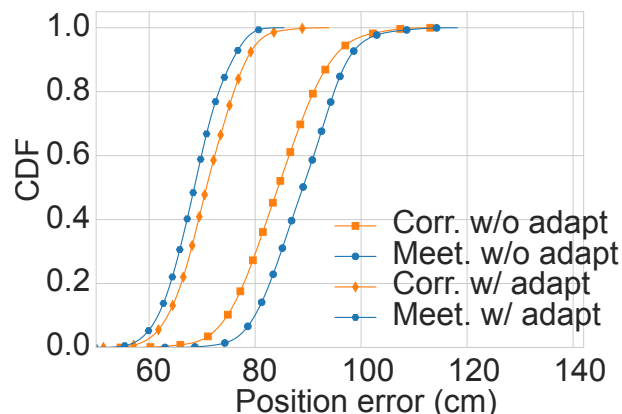


FIGURE 4.9: Location error CDF when using adaptation in a different LoS environments.

plots the CDF across all test samples. The average location error is only 19.05 cm, and 90th percentile 21.93 cm, when the testing data are generated in the same environment as the training data.

We then change the testing environment, by placing metallic cabinets and blocks of stacked chairs with metallic frames around the radios, while maintaining the LoS path towards the target. The resulting average error increases to 49.62 cm. Finally, we change the propagation environment by blocking the LoS path between the target and both the Tx and Rx radios. As expected, we see a sharp decrease in the estimation accuracy, with an average error around 96.40 cm. These results verify that data-driven passive localization is highly sensitive to the environment.

Domain	w/o Env. Mitigation		w/ Env. Mitigation	
	Mean (cm)	90 th pctl (cm)	Mean (cm)	90 th pctl (cm)
Outdoor	83.77	92.31	68.79	75.43
Yoga room	96.56	103.34	77.42	85.65
Corridor	87.42	95.61	69.57	77.39
Meeting room	83.25	89.73	71.50	78.10

TABLE 4.4: Estimation error in different environments in NLoS.

Network	Batch Time (ms)	Total (h)
No split	21	8.11
Freq/8 and Spatial/2	13	5.23

TABLE 4.5: Training time delay with different splitting methods.




FIGURE 4.10: Estimation error with and without retraining on generated corridor data.

4.5.2 Effectiveness of Environment Adaptation

We repeat the yoga room experiments by enabling the environment adaptation module. The results (Fig. 4.12) demonstrate that our generative model can reduce the mean localization error by 19% (from 96 cm to 78 cm), when the training and testing data are produced in the same environment, but with LoS and NLoS setup, respectively. We further run experiments across different environments, starting with an LoS setup. All the environments reuse the same LoS training data in the yoga room. Fig. 4.9 plots the CDF of localization errors. We observe that the environment adaptation module can reduce the LoS mean localization error from 84 cm to 70 cm, and

Domain	Mean (cm)	90 th pctl (cm)
LoS Empty room	19.05	21.93
LoS w/ Obst	49.62	67.15
NLoS	96.40	103.34

TABLE 4.6: Estimation error when changing the layout in the same general environment (yoga room).




FIGURE 4.11: Location error CDF using the retrained inference w/ and w/o adaptation generator.




FIGURE 4.12: Location error CDF when enabling the environment adaptation module.

from 89 cm to 69 cm, for the corridor and meeting room, respectively. We then repeat the above experiment but with LoS blocked, and include the outdoor setup. Fig. 4.12 shows that, when the environment adaptation is added atop the position inference module, it reduces the mean error by 28%, 20.4%, 23.5%, for the outdoor, corridor, and meeting room, respectively. These experiments verify the significance of our generative model in dealing with the heterogeneous multipath environment.

We have also conducted experiments on the Widar 2.0, which builds on geometrical models and has an open-source implementation [108]. Fig. 4.13 plots the CDF of location errors within the yoga room environment (Fig. 4.8). We see that the average error in LoS is 1.8 m, slightly larger than [108], likely because of more abundant multipath reflections in our setup. Remarkably, under NLoS, the mean error escalates to 6.7 m, which verifies the failure of the closed-form models in NLoS (Sec. 4.1). In

contrast, DPLoc maintains a small error of 0.77 m under the same setup (Table. 4.4).

Effect of inference model retraining on localization error.

Furthermore, we were able to reduce even more the localization error by roughly 4 cm compared to a strategy without retraining our inference model and using the generative model, see Fig. 4.10. For instance, the mean error without the retraining strategy is 73.07 cm in the corridor environment, while retraining reduces it to 69.55 cm. We tested our inference system after retraining without using the domain adaptation generator (Fig. 4.11). We observe that our inference model can achieve better performance after retraining, even without using the adaptation module, which is explained by the data augmentation procedure based on the generated samples using our domain adaptation method.




FIGURE 4.13: Location error of Widar 2.0.

4.6 Conclusion

Through this work, we have established RFLoc, a large labeled dataset for Wi-Fi CSI-based passive localization. This dataset was designed with accurate labels collected through a computer-vision-based motion capture system. We further identify the advantages and fundamental challenges in deep learning-based passive localization. In addition, we design a generative model to address the NLoS and environment-dependency issues. The Dploc model exploits state-of-the-art generative adversarial networks to ensure an unsupervised migration of a localization model to new unseen propagation environments. By adopting such approach, we were able to reduce the localization error by 20 % in average. We plan on expanding this work to achieve super-resolution on new environments with limited measured data. We expect that the RFLoc campaign can inspire more open-source, data-driven research in wireless sensing.

Chapter 5

Reproducibility for Wireless Networking Experiments

5.1 Introduction

Reproducibility is a crucial aspect of establishing a trustful scientific contribution. Its importance does not lie in proving the correctness of published results but the ability to retrace all the steps carried during a research project. When Faced with sophisticated methodologies and complex data analysis, reproducibility ensures a coherent and transparent recount of information for explaining the reported claims. Given its importance, one could think reproducibility is one of the main aspects of a researcher's work. Most scientific communities still consider reproducibility as a marginal activity optionally carried and upon explicit demand of conferences or journal editors. This situation has led to a so-called reproducibility crisis. According to a survey carried by Nature [1], 70% of researchers could not reproduce experiments from other peers, and 50% failed at repeating their experiments. This survey establishes the severity of the current state and the need for more efforts towards reproducibility. ACM has exerted more effort to spread awareness about the importance of reproducibility as an integral part of a researcher's occupations. While ACM issued a policy defining a specific terminology to ensure the homogeneity of the review process, it proposed no additional guidelines.

Along with technical challenges, we consider sociological issues as prominent circumstances hindering reproducibility. Most factors related to scientific research irreproducibility (selective reporting, the pressure to publish) stem from these issues. In this work, we will try to address aspects related to each challenge regarding networking and wireless networking communities. The rest of this chapter is: In Section 5.2, we identify the main problems related to reproducibility and identify the reasons and factors behind the current state. Later on, in Section 5.3, we present the results of a survey on the state of reproducibility in the community to draw a more accurate diagnostic. In Section 5.4, we offer two approaches to tackling the reproducibility problem from our experience in reproducing wireless networking and sensing projects. In Section 5.5, we discuss procedures and guidelines for improving reproducibility in the community from our experience and through our interactions with members of the ACM community in workshops and other talks. We conclude by drawing our perspective on reproducibility in the wireless experimentation domain.

5.2 Why is Reproducibility is so hard to achieve?

Introducing reproducibility in the community as an indispensable part of an investigator's workflow requires considering the peculiarities of the scientific community and its inherent technical challenges. Members of the ACM community have organized a reproducibility workshop during SIGCOMM 2017 to identify the pitfalls and propose recommendations to enhance reproducibility under the prevailing circumstance [120].

Author incentives In our community, reproducibility is still frowned upon and has little to no influence on scientific recognition. Given this state of mind, it is a costly time-wasting activity that slows down original scientific contributions. Even in the absence of intellectual property restrictions and privacy issues, providing a reproducible contribution solicits an added effort including more resources (human operator, material) to repeat the experiments and rigorous maintenance of the documentation consistency. However, these added organizational efforts are beneficial. By ensuring the experiment reproducibility, authors can reuse the developed tools and artifacts and maintain a convenient archive that outlives the project's life-cycle. It could quicken project development by avoiding the time-consuming re-implementation of pre-existing tools.

Reviewer incentives While the ACM reproducibility initiative provides authors with reproducible projects a seal of approval and thus more credibility, reviewers have little to gain. Reviewing reproducibility is a difficult task; It involves a deep understanding of the internal functioning of the provided artifacts and checking their sanity and consistency. It involves time-consuming investigative work such as setting up custom environments, launching code, or accessing datasets. To conduct a fair and trustworthy review, the reviewer has to be keen on diving in projects with aspects involving diverse expertise. A reviewer with an adequate skill set is rare. So other than the sense of duty, there is rarely any incentive for her to delve in such endeavors.

Technical limitations and Artifacts Availability Artifact dissemination and its technical ramifications impose adopting a new review process. Depending on the reviewed project, artifacts could be subject to restrictions, which would complicate the reviewing process. Sometimes, the reviewer must sign a non-disclosure agreement to review confidential artifacts. Also, the current review process cannot guarantee anonymity, and reviewers may reuse artifacts to their discretion.

Ensuring the perennality of artifacts is crucial to maintain reproducibility. The heterogeneity of the artifacts (hardware, software, hybrid) proposed by the community surfaces new types of challenges. Given these circumstances, the usefulness of an imposed unified platform such as IEEE Code Ocean [32] to run the author submitted code is doubtful. Usually, artifact description and documentation are vague, too broad, or superficial. This situation could be fatal when the used material or resources are not available anymore, and therefore finding alternatives becomes unlikely.

5.3 Survey of Reproducibility in the Community

Investigating the factors behind the lack of reproducibility calls for surveying the current state of the produced research artifacts [56] *. This survey has a critical aim to sound out the likeliness of reproducibility, given the current practices and the nature of the provided artifacts.

5.3.1 Method

We conducted the study among participants from four leading ACM SIGCOMM conferences, including venues with broad topics (SIGCOMM, CoNEXT), and more domain-specific venues (IMC and ICN). CoNEXT is a conference on novel and emerging networking technologies; ICN is a conference on Information-Centric Networking; IMC is a conference on Internet measurement and analysis, and SIGCOMM is a major generalist conference in communications and computer networks.

Questionnaire We designed an informal questionnaire (Google form) forwarded by the conference chairs to the authors of the accepted papers. We sectioned the questions into three groups. The first group aims to collect the metadata (paper title, conference name, and author email). Then, we provide an open-ended question to allow the authors to explain their artifacts in their own words, where we asked the authors to provide a brief but precise and complete description of their tools and data. We also asked the participants to provide a URL to access these artifacts. Responders provided a link to their artifacts with a more comprehensive description than the one in the paper. Among all responses, three papers were not listing any link to their artifacts in the paper itself. The last section is a self-assessment provided by the authors in a scaled closed-ended question about the reproducibility ease of their publication with a scale from 1 ("easy – an undergraduate can do that") to 10 ("hard – only I can do that").

Recruiting Participants To solicit participation in the survey, we asked the chairs of the technical program committees to contact all authors of accepted papers after the conferences. Some authors did not receive the email if they disabled the notifications in the conference submission system. However, at least one author per paper was reachable, and we removed duplicate submissions. We increased the number of survey participants of ACM ICN by sending a reminder. After we collected the provided data, we analyzed the artifact descriptions and the actual artifacts in more detail. Those authors who replied provided at least the same information in their published paper.

The form and all the data we collected are available online [129].

5.3.2 Results

Out of the 137 potential respondents, 49 researchers (35.8%) took part in the survey. The response rate was diverse among the conferences. Most of the ICN authors

*Artifact as defined by ACM is a digital object either created by the authors to be used as part of the study or generated by the experiment itself".

	Conference Potential Re- sponses [# papers]	Actual Responses [# papers]	Response Rate [%]
CoNEXT	40	8	20
ICN	19	12	63
IMC	42	17	40
SIGCOMM	36	12	33
Total	137	49	35.8

TABLE 5.1: Summary of artifact survey, compared to the overall number of published papers per conference.

(63%) were responsive, followed by IMC (40%) and SIGCOMM (33%). Only 20% of the CoNEXT authors participated in the survey; see Table 5.1.

Given the surveyed conferences’ diversity, we identified three main categories —*Architectural, Measurements, Miscellaneous*—from the 49 collected responses. Among the surveyed papers, 22 were aiming at providing a new network algorithm, protocol, or architecture [45, 72, 35, 78, 41, 42, 85, 87, 84, 118, 153, 68, 31, 82, 79, 97, 51, 30, 88, 28, 66, 162]. The scientific work described in 19 papers is focused on measuring an already installed system as measurements [139, 94, 90, 75, 154, 157, 96, 46, 99, 86, 131, 61, 146, 141, 57, 117, 111, 130, 81]. 8 papers are classified as miscellaneous, which do not fit in the other topics, optical networks or security [133, 95, 152, 158, 59, 26, 37, 27].

Then from the open-ended questions intended for artifact description, we propose an artifact classification —*Tools, Hardware, Simulation, Dataset, Testbed*—according to their nature. The surveyed artifacts are based on custom-made software (*Tools*), depend on specialized hardware (*Hardware*), relies on numerical evaluation, simulation, or emulation (*Simulation*), based on an external dataset (*Dataset*) or uses a testbed or a specific infrastructure (*Testbed*). For each type of artifact, we identified three options. Either the artifact is *new* (i.e., researchers had to build the artifact on their own) or is built upon existing material. When at least two papers have used existing material, we highlight the artifact by naming it in the table. If only one paper uses a specific existing material, we summarize those artifacts by *Other*. We count an artifact in the table only if it is provided (by some sort) to the community. The only exceptions are private testbeds, which have been used by users but that cannot be shared.

Table 5.2 shows the number of papers for each research topic and the applied method. Significant differences are visible among the fields. For architectural papers, researchers use existing tools or change the operating system. However, researchers in the measurement domain created their tools (automation scripts). A trade-off is followed by the ICN community, which extends libraries and well-established tools or creates new tools from scratch.

The conference’s scope and themes influence the artifacts provided by the authors. From the survey, we have identified these tendencies:

	Arch.	Measurements	Misc.
Tools			
New	8	11	6
NDN [156]	3	–	–
CCN-Lite [122]	2	–	–
Linux/RIOT [33]	6	–	–
Other	8	4	2
Hardware			
New	–	–	1
Smartphones	3	–	1
Specific	3	1	1
Simulation			
New	1	–	–
MATLAB [137]	–	–	2
ndnSim [8]	2	–	–
Other	4	–	1
Dataset			
New	–	12	3
CAIDA [136]	2	4	–
Other	4	4	1
Testbed			
Private	2	5	2
IoT-Lab [7]	2	–	–
RIPE [132]	–	4	–
Other	1	6	1
Average rank	4.2	3.5	2.0

TABLE 5.2: Summary of artifact nature. Please note: Some artifacts are counted in multiple rows (if applicable). But only once per column.

- **Architectural Communities**, e.g., ICN, depend on experiments and hardware-specific solutions. They require a non-negligible start-up cost. For instance, several proposed projects rely on specialized hardware or modified platforms (Smartphone). However, we see a tendency of using publicly available testbeds with dedicated operating systems, which promotes more reusability of the project artifacts and a smoother review process.
- **Measurement Communities**, e.g., IMC, rely on publicly available datasets, testbeds, or measurement platforms. Given the importance of such resources, and the effort required for such endeavors, IMC has established *the community contribution award* as an incentive to authors who provide high-quality reusable artifacts [3]. However, some papers published in the measurement community use private infrastructures or testbeds.

A common trait between these communities is the importance of testbeds and datasets, highlighting the need for publicly available reliable testbeds, not only to allow external users to conduct their experiments but also to allow comparison of solutions by using the same infrastructure in multiple studies. As a general remark, researchers in the measurement and simulation domains are much more confident in the ability to reproduce their work, compared to other researchers. The least confident researchers are those who worked with sophisticated platforms or testbeds.

Data Storage The responding authors provide information for reproducing their results. Overall, papers point to webpages containing artifacts (tools, data). Less than 20% of researchers store artifacts on their personal or project websites; Instead, they prefer popular code platforms such as GitHub. Four papers had broken links further establishing the need for well-maintained platforms such as the ACM Digital Library [2]. Using such a platform would improve the durability data access, with a snapshot of the status of the artifacts at publication time.

5.3.3 Discussion

As presented earlier, and from our analysis of the survey results, we can identify two classes of scientific work:

- **Network architecture contributions that propose new standalone systems or changing existing systems.** In IoT, wireless, and optical networking, researchers often have to rely upon specific hardware and testbeds. However, they comprehensively describe the used environment. A tendency that is much less common in network architecture contributions that share most challenges as the mentioned research areas but do not offer the same level of descriptive effort.
- **Measurement contributions that rely on datasets.** The measurement community uses well-known public datasets and public measurement platforms and makes their collected data available to everyone. However, there is also a broader set of measurement papers that use confidential data, and thus, instead of publishing the raw data, they provide aggregated data instead. In

large measurement projects, companies often provide data. These data are confidential because of business reasons or cannot be shared because of size.

As seen by categories, for developing new network architectures, researchers emphasize on implementing real systems and make sure their code is available, putting aside from the actual data used in their evaluation or the precise description of their evaluation environment. Measurement papers insist on the data they used more than on the tools themselves. With a system relying on specific hardware, researchers describe the hardware and software they use.

Many papers in our survey used existing public datasets, testbeds, infrastructures, or source code. For instance, many papers are reusing artifacts from other research projects. Given the importance of such tools and artifacts, we should give special care when disseminating and maintaining them.

5.4 Reproducibility for Wireless Experiments:

Testing and comparing the performance of wireless systems, like for any other scientific area, requires the ability to reproduce experimental results. In this section, we describe the specific issues we encountered when focusing on reproducing the experiments described in a paper related to wireless systems.

5.4.1 Reproducing a wireless experiment: A case study

We selected the OpenRF [83] paper published in SIGCOMM 2013, an exciting research work proposing a beamforming system on commodity Wi-Fi devices. We illustrate how reproducibility depends strongly on the hardware used and why extensive knowledge of the hardware used and its design is necessary. Based on this experience, we propose recommendations and lessons for the design of reproducible wireless experiments.

Our initial goals when starting our work around OpenRF were to study and design cross-layer performance enhancements in commodity MIMO devices. This effort required the ability to easily modify the network, MAC, and physical layer parameters of the wireless system. For example, accessing the channel state information and using it to pre-code the transmission matrix is essential for beamforming and interference nulling. Upon receiving an explicit beamforming request from the transmitter, the receiver computes and sends back a compressed form of the measured CSI matrix respecting the IEEE 802.11n recommendations. In this way, the transmitter can pre-code its transmission matrix by allocating complex weights on each antenna element, which influences the phase and the power of the radiated signal. This method results in focusing the radiated signal to the receiver's region, which increases the signal-to-noise ratio at the receiver and reduces interference for others. Even though the IEEE 802.11n draft [73] mentions this feature, very few chip manufacturers enabled it natively. It is usually impossible to access the CSI or to pre-code the transmission matrix in Consumer-Off-The-Shelf devices because of firmware lock.

OpenRF

In the absence of close collaboration with chip manufacturers, we used the OpenRF extensions* of the 802.11n CSI Tool, initially developed by the University of Washington [11] for the Intel Wireless Link 5300 802.11n MIMO cards. It is interesting to use OpenRF as it enables commodity Wi-Fi access points to perform several MIMO management techniques: interference nulling, implicit beamforming, and interference alignment. Contrary to the explicit beamforming technique described in the IEEE 802.11n draft, OpenRF uses an implicit beamforming mechanism that requires no feedback from the receiver, assuming that the channel conditions are the same in both directions. The OpenRF code includes a MATLAB function that pre-codes the channel for either nulling interference or beamforming according to the MAC address of the receiver. This function reads the CSI matrix on the transmitter side, computes the steering or beamforming matrix in user-space, and sends it back to the wireless chip driver. We identify the OpenRF system components in Figure 5.1.

To sum up, the perspective of using beamforming within a cross-layer system was the main reason for us to reproduce the OpenRF experiments.

Reproducibility challenges

The OpenRF code release web page¹ contains installation and testing instructions with an old Ubuntu 10.04.4 long term support (LTS) ISO image to download. It was late 2015 when we were trying to reproduce the OpenRF experiment, and we faced the problem that this deprecated Ubuntu 10.04.4 LTS release would not support our recent hardware; so we first attempted to install the OpenRF tool on a more recent release (15.04), keeping the same kernel as the one used for the original implementation. Within a 15.04 image, compiling OpenRF raised issues related to deprecated packages. We could not conduct and maintain our experiments, even after applying the changes and improvements. Suspecting that the problem would come from the OS distribution, we tried to identify and use the previous Ubuntu long term releases, more precisely 14.04 LTS and 12.04 LTS. Installing OpenRF was much simpler within these distributions, as it did not raise compilation issues as before.

However, this setup did not yield satisfactory results as we could not reproduce the experiment, which comprises pre-coding our transmission with a beamforming matrix. When reached through email, the original authors of the OpenRF paper suggested that using 12.04 LTS was a potential source of the problem. So, we took measures to install 10.04 LTS regardless of the many incompatibilities - video card, SSD hard drive, and network card were not supported - and we finally installed it after much effort. Therefore, we adopt the following system configuration to reproduce the experiment:

OS	Kernel	Wireless Cards
Ubuntu 10.04.4 LTS	3.5.4-csitool+	533AN MMW Full / Half

At that point, we had precisely matched the original conditions of the OpenRF paper but still could not reproduce its results. Getting the CSI matrix was straightforward

*<https://www.andrew.cmu.edu/user/swarunk/openrf.html>

with the Intel CSI tool. However, for beamforming using OpenRF, we observed no variation in the signal-to-noise ratio at the receiver. Hence, we adopted a verification process that eventually allowed pinpointing the problem. To test the beamforming matrix injection at the transmitter level, we check the CSI matrix computed at the receiver end by plotting the SNR for each one of the sub-carriers before and after injecting the transmission matrix. Usually, the spatial mapper should use the information from the injected matrix to compute the beamforming vectors. However, we observed that the beamforming matrix was never applied, and the spatial mapper was still using Intel’s indirect spatial mapping every time we send a packet. This situation is peculiar since we were getting a positive acknowledgment from the driver about the beamforming matrix injection. We double-checked every step down the path for sending the CSI report to the wireless card firmware. We resorted to checking the driver messages which report the state of the card and problems but could spot no message related to a firmware crash or a rejection of the matrix injection by the firmware. In addition, we adopted several transmission scenarios while injecting the beamforming matrix. At first, we kept the rate adaptation mechanism functioning, which yielded no changes other than those related to the rate variation. Then, we specified the number of space-time streams (between 1 and 3), and the modulation and coding scheme (MCS) code, and we noticed no variations on the SNR plots. We adopted this verification process on different Ubuntu releases starting from a recent 15.04 Ubuntu, including 14.04 and 12.04, down to a deprecated 10.04 LTS. We made this effort to match the software setup of the original implementation of both the Intel CSI tool and OpenRF. In our work, we used almost 45 wireless cards that vary in sizes (Half or Full) and origins, and none supported the injection of the beamforming matrix. We describe further details on our verification process in 5.A, stating the technical path we followed. Our goal is to promote such practices that lead to better replicability and to ease the learning process for future experimenters.

From this experience, even though the authors provided their experiment source code, there was not enough information about the hardware settings. Fully describing the experimental setup is critical to reproduce the experiment, especially when the equipment presents peculiarities, e.g., related to the manufacturers. We discovered several types of Intel Wi-Fi Link 5300 cards on the market. Some were engineering samples, which are a source of problems and complications. We contacted the authors to seek their support. However, this interaction was not fruitful as none of the proposed solutions succeeded. The OpenRF experiment reproducibility process took three weeks in four months.

Observation and Discussion

Although *reproducibility is the ultimate goal* as stated by the ACM Result and Artifact Review and Badging policy*, the replicability of experimental results by independent researchers using author-supplied artifacts represents an important intermediate step. The simplest case is when experiments can be done using hardware and software available in open testbeds such as ORBIT[†] and R2lab[‡]. Here, authors only

*<https://www.acm.org/publications/policies/artifact-review-badging>

†ORBIT testbed at WINLAB, URL: <http://www.orbit-lab.org/>

‡FIT Reproducible Research Lab (R2lab) at Inria, URL: <http://fit-r2lab.inria.fr>

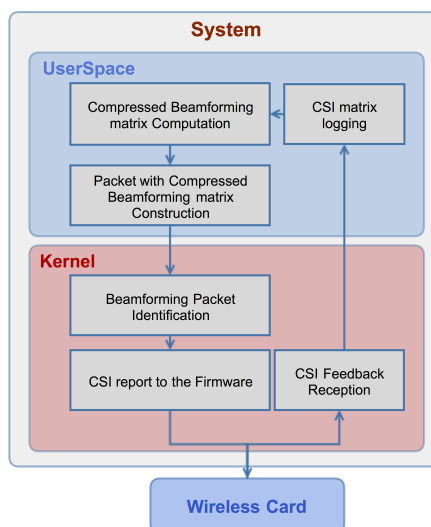


FIGURE 5.1: Beamforming System Design

need to ensure they explain the scenarios and to provide all the code and scripts used to replicate the experiments.

However, these testbeds are not convenient for many experiments, and for instance, the experiment repeatability could be tied to a specific hardware setting not available in those testbeds. Therefore, it is essential to explain the configuration. The description should focus on the specifications and requirements necessary for reproducing the experiments. A simple description of the OS reference is not sufficient; There should be a mention of what makes this version of the OS necessary for the success of the experiments (packages, drivers, ...). The full description of the defining characteristics has to be provided and not just a simple description of the setup used for the experiment. This could help future experimenters to find a suitable equivalent when the hardware or software used is not available anymore. This allows for a future-proof reproducible experiment. It is helpful that authors verify that their solution works with hardware from different manufacturers or OS and, as importantly, that they mention the setups that did not work. Again, mentioning the motivation behind the choice of particular hardware or software is desirable. If the verification work cannot be done, an explicit warning should be present, and all the details about the hardware should be provided with the corresponding references (serial number, manufacturer, vendor, etc.). As for the software used for the system design, authors should avoid when possible licensed software. For example, using MATLAB code also hindered our setup because of license management difficulties on experimental machines.

It is worth considering managing a research project more like a software development project. In many aspects, the challenges for reproducibility have strong similarities with the ones of software development, and tools like source code management tools, maybe even test suite frameworks, and interactive computing concepts like notebooks, can be beneficial in building more reproducible research.

5.4.2 Towards Reproducible Wireless Experiments Using R2lab

We have launched a project called ORION [93] that focuses on the design of an orientation measurement system, based on the joint estimation of the angle of departure and the angle of arrival of a Wi-Fi signal. Our system exploits antenna arrays at both the transmitter and receiver ends. When a signal is radiated or received by an antenna array, a phase shift is created between the signals of adjacent antenna elements. By computing this phase difference from the CSI measurements, an estimation of the angles mentioned above is possible. As estimation techniques for the angle of arrival were already mentioned in several works before, we proposed a method for estimating the angle of departure, exploiting a well-known MIMO mechanism called spatial multiplexing, and supported in commercial off-the-shelf (COTS) wireless cards. The reproducibility of our experiments was a driving force, and a necessity in the system design since an orientation measurement system requires consistency and resilience against recalibration or reset.

In our previous work [93], we explain in greater length the scientific challenges addressed in this work, which focuses on properly dealing with various sources of phase shifts and inconsistencies that need to be accounted for to achieve decent accuracy.

In the present part, we wish to describe into more detail the actual experimental method and work carried out to reach the ORION results, with a focus on what is easily reproducible and what is not, and to discuss improvements.

Experimental Methodology

Throughout our experiment design, we have adopted R2lab as an initial testbed for designing and testing features. By using a controlled wireless environment such as an anechoic chamber simplifies the design process of a wireless system. It allows identifying design flaws before launching costly realistic wireless environments [112]. It has been crucially important for us to use the R2lab testbed* in the early stages of the project, for an initial calibration phase.

R2lab, testbed for reproducible wireless experiments

Room characteristics


The R2lab platform sits in a 90m² insulated anechoic chamber in a basement of a building at Inria, Sophia Antipolis, France. Figure 5.2a shows a snapshot from inside the room and Figure 5.2b the topography. It hosts thirty-seven PC nodes on the ceiling scattered on a fixed grid; more than half of the nodes feature an software defined radio (SDR) board, of various kinds.

It is insulated from the outside electromagnetic conditions by a Faraday cage. RF absorbers are needed to prevent a high level of reflections on the copper foils.

*FIT Reproducible Research Lab (R2lab) at Inria, URL: <http://fit-r2lab.inria.fr>



(A) R2lab room



(B) R2lab topology

Wi-Fi nodes

The 37 wireless nodes are Icarus* computers provided by NITlab[†] with the following features:

- CPU Intel® Core™ i7-2600, 8M Cache at 3.40 GHz
- 8GB DDR3
- 240 GB Solid State Drive
- 3 Gigabit Ethernet interfaces: one for remote node power and reset management, one for control used by the testbed management framework for providing access, and one for data, dedicated to experimentation, e.g., to create wired link or to connect to an SDR device such as USRP2 or N210.
- 2 Wi-Fi MIMO NICs dedicated to experimentation: one Atheros 802.11 93xx a/b/g/n and one Intel 5300. Each card is connected to 3 dual-band 5dBi antennas, operating on both 2.4GHz and 5GHz. Antennas are spaced of 2.8cm, which corresponds to half the wavelength at 5GHz, see photo in Figure 5.3.

To control and monitor each Icarus node, we use the NITlab's Chassis Manager Card (called CM card); this device embeds a tiny web server that can serve HTTP requests to power on/off and reset the motherboard, or one attached USB device.

SDR devices

About half of Icarus nodes are attached to an SDR device. R2lab currently supports USRP1, USRP2, N210, B210 and X310 Ettus[‡] devices and also LimeSDR from Lime Microsystems[§], see Table 5.3.

Each N210 device includes an SBX-40 USRP daughterboard (400 MHz - 4.4 GHz, 40 MHz BW), whereas the X310 board includes two SBX-120 USRP daughterboards (400 MHz - 4.4 GHz, 120 MHz BW).

*Icarus node: <https://nitlab.inf.uth.gr/NITlab/>.

[†]NITlab:

<https://nitlab.inf.uth.gr/NITlab/hardware/wireless-nodes/icarus-nodes>.

[‡]Ettus Research: <https://www.ettus.com/>.

[§]Lime Microsystems: <http://www.limemicro.com/>.



FIGURE 5.3: Wi-Fi dual-band antennas for a node

Device type	USRP1	USRP2 & N210	B210	X310	LimeSDR
# of	4	8	6	1	2

TABLE 5.3: Available SDR devices

Using R2lab for calibration

The benefit of using R2lab [91] [103] is that its basic hardware components, like nodes and wireless devices and antennas, remain in a relatively constant configuration, which is a good thing in terms of reproducibility. Running experiments in a controlled environment, allowed us to draw observations between runs that differed in only a small set of parameters, which is something hardly possible in an open environment, if at all. Launching our first experiments in R2lab helped to provide us the means to classify the various phase shifts and inconsistencies at work, and namely:

- **Phase inconsistencies due to hardware defect:** Intel Wi-Fi Link 5300 wireless cards suffer from a defect on the third RF chain on the 2.4 GHz band, which further alters phase data. Instead, we used the 5 GHz band that does not exhibit this issue.
- **Phase inconsistencies not affecting the estimations:** we could experimentally verify that within our approach - that does not use time-of-flight - even though SFO and CFO do bring phase shifts, they are equally applied to all RF-chains and thus can be safely ignored.
- **Phase shift due to the adopted transmission mode:** spatial multiplexing in the wireless card involves two techniques, spatial mapping SM and cyclic shift delay, that create phase shifts that need to be compensated for, using a mapping matrix from either the manufacturer's datasheets, or from the IEEE 802.11n draft [73].




FIGURE 5.4: Pseudospectrum for AoA (left) and AoD (right) before (blue) and after (red) phase correction (from [93]).

- Phase inconsistencies impairing reproducibility of results:** RF oscillator phase offset is a constant phase shift added to each one of the RF chains of a wireless card. This phenomenon occurs because the RF chains are locked at different instants when starting up the wireless cards. Therefore, each RF chain will have a different constant offset added to the measured phase. This offset remains constant across one session*.

Using an anechoic chamber gave us a unique way to address each potential issue. As an illustration of the outcome of dealing with the latest mentioned source of phase shifts, we showcase in Figure 5.4 the estimation results of AoA and AoD before and after applying the phase correction method introduced in ORION[93]. This technique relies on measuring the RF oscillator offset from a known reference point. By doing so, we can reproduce comparable results throughout different measurement sessions.

Details on the experimental setup

In terms of software, the ORION paper uses the same technical substrate as OpenRF [83], namely Intel Wi-Fi Link 5300 wireless cards, with the Intel CSI tool [11], which allows interactions with the firmware, such as reading or writing the CSI matrix for 802.11n HT packets.

In terms of radio, after determining that the 2.4GHz band could be an issue, we used the 5GHz band with 20MHz of bandwidth. We also set up the cards in the injection mode, which avoids the need for an association with an AP and allows raw Wi-Fi packet transmission. The antenna arrays installed in R2lab are fixed and thus inoperable for rotation estimation. Hence, we used two external uniform linear antenna arrays (ULAs), which are connected to the nodes through extension cables to have more liberty of movement. These cables are 3 meters long and compatible with both the Wi-Fi bands with a 2dB signal attenuation for the 2.4GHz band. To respect the coplanar aspect of the experiments and to accommodate the cable rigidity, we placed the ULAs on ladders at the same height as the R2lab nodes. The antennas used for our experiments are 5dBi omnidirectional, compatible with 2.4GHz and the 5GHz band. Since we are operating at a 5.32GHz frequency, we created a 2.8cm inter-antenna spacing, which corresponds to a half-wavelength.

*A session is a period during which hardware configuration is fixed. Typically a node reset, a node reboot, or simply a change of frequency in the card, yields a new session.

In our experiments, we faced issues while attempting to decode a packet at the receiver end when using three antennas with three spatial streams. Therefore, we operated with two spatial streams and a 2-antenna system at the transmitter end. As for the receiver, we had no problems while using the 3-antenna setting.

Using an orchestration engine

Most wireless experiments face a need for fine-grained synchronization of tasks otherwise intrinsically parallel. So as experimental scenarios grow in complexity, easy and effective management of synchronizations in an experimental scenario is critical. We chose `nepi-ng` [104][38] as an orchestration engine for running experiments, as it enhances the experiment reproducibility:

- **Lean software dependency.** The only requirement for managing a computing resource is for it to be reachable via `ssh`. OMF requires a dedicated messaging infrastructure to be in place for controlling resources.
- **Convenience.** `nepi-ng` is usable from an experimenter's laptop in terms of both experiment control and data analysis.
- **Efficiency.** `nepi-ng` transparently and seamlessly share and re-use `ssh` connections expensive to establish.
- **Modularity.** The modules defined in `nepi-ng` can be re-used, shared, or so that the testbed can expose convenience helpers such as decorating the core of one experiment with utility features like loading images, turning off unused nodes, checking for a valid reservation, or similar.
- **Clarity.** `nepi-ng` adopts a job-oriented programming model to enable parallelism and synchronization where dependencies are explicit between jobs, allowing for an explicit and visual representation of dependencies, as illustrated below.

Thus `nepi-ng` focuses on providing an efficient paradigm for orchestrating experiments and a clean way to write these gory details using the most appropriate tool - often a plain shell script.

Given all the complexity described so far, many runs of the same scenario, or small variants of it, need to be carried out to fine-tune and validate the overall method. Being able to launch such runs efficiently is very desirable, and to achieve this we have used `nepi-ng`^{*}, a tool designed for R2lab that addresses this usage[†].

Based on this tool, we have been able to script[‡] the actual data collection process, which can be achieved in only few minutes even with a reasonably large set of nodes, while letting the experimenter focus on meaningful issues, instead of having to focus on the tedious task of adequately coordinating the various stages.

^{*}Read the Docs documentation for `nepi-ng` at <https://nepi-ng.inria.fr/>

[†]Tutorials for `nepi-ng` at <https://r2lab.inria.fr/tutorial.md>

[‡]ORION data gathering script at <https://github.com/parmentelat/r2lab/tree/public/demos/jobs-angle-measure>

Furthermore, given the efficiency of this orchestration tool, we have also been using it for launching and controlling 4G networks in R2lab using OpenAirInterface [91]. We have deployed a standalone 4G network in less than 5 minutes. All the network components (base station, subscriber management, serving and packet gateways, network traffic analyzers) were run automatically using the nepi-ng experiment orchestration tool.

Experimental material in a *git* repository

As an additional step towards a more reproducible experiment, we have gathered in a single public *git* repository* a detailed description of the system setup and the hardware, with illustrations, that hopefully provides sufficient information on the measurement steps. To avoid any licensing problems when using MATLAB the repository contains a python version of our code data post-processing tools, and a *jupyter notebook*[†] that allows running all the steps of our angle estimation technique.

Running in a different environment

To make our system deployable, we needed to test our system in an open, realistic, and non-controlled wireless environment. We reproduced the same experiment in an office room. This setup involved several multipath clusters; the main difference with the original system setup in R2lab is the elevation of the antenna arrays, as we were careful to preserve their co-planarity. We have been using the same hardware as in R2lab, including antennas cables and antennas spacing.

With all these pieces in place, it has been rather straightforward to re-run the same experiments in such an open environment. As a side effect, the same tools can be used by other experimenters to reproduce our results with reasonable effort. Here is a tentative list of the topics that may require extra work to do, given the feedback we have gathered so far:

- **Hardware setup** when reproducing outside of our premises, setting up the right wireless cards and antennas should be rather straightforward, but will require some initial effort though.
- **Position in space** of the antennas: whether the experiment is run in R2lab, in our open environment or other premises altogether, antennas crucially should be properly spaced, and in a common plane; This is hardly automatable, and so accounts for most of the time spent in variously tedious and possibly time-consuming activities.
- **Software image** we do provide a ready-to-load image for the R2lab nodes for running the experiment, but in the current state, this is not usable as such on other types of hardware, due to our imaging technique. It would make sense to use more standard techniques like e.g., *docker* to manage images, although it is not yet clear if running in a container-based system can provide a level of hardware interaction typically needed in wireless experiments.

*ORION git repository <https://github.com/naoufal51/Orion>

[†]ORION notebook at <https://www-sop.inria.fr/teams/diana/orion/>

Discussion

In the research work's early stages, and along with the steps of our system design, the main concern was control and calibration of the hardware. In this context, having access to an anechoic chamber has been a tremendous asset. Extensive knowledge of the experiment hardware with a thorough understanding of its capabilities and limitations is essential to decide on a measurement system based on COTS hardware. For instance, in our case, the card was functioning abnormally in the 2.4GHz band. In the same spirit, we have published a study on some pitfalls to avoid when using COTS hardware in experimental testbeds [135]. As the project became more mature, our concerns shifted to making the experiment more reproducible and have led us to manage the project more like a software development project. In many aspects, the challenges for reproducibility have strong similarities with the ones of software development, and tools like source code management tools, maybe even test suite frameworks, and interactive computing concepts like notebooks, can be beneficial in building more reproducible research.

In terms of operating the R2lab testbed, this study has brought very fruitful insights into what user expectations can be, and even if it is impossible for remote users to control for example the position of antennas, it is crucial to describe such details accurately. R2lab's website formally describes a reference configuration for all such elements in the room so that users can return the chamber in a known and well-documented configuration.

5.4.3 A step further towards runnable papers using R2lab:

For illustrating a methodological approach to publish experiment results, we have defined an experiment, that comprises measuring, for all couples of nodes (a, b) in the testbed, the power received by b when a is sending. We designed the experiment to characterize the radio channel in the room according to the measured power distribution and infer the impact of the disposition of the room and its peculiarities, e.g., pillars and blind corners, on the wireless channel. This experiment is particularly interesting as it helped us characterizing the propagation environment before experimenting on Wi-Fi mesh networks [38] [103]; it allowed identifying the right nodes to chose to enable multi-hop scenarios.

Jupyter notebooks are hybrid documents that mix formatted text and executable code. In this section, we are exploring the capabilities offered by using Jupyter notebooks, with publicly available infrastructures like `github.com` and `mybinder.org`, as a step towards runnable papers. Again, the experiment is straightforward, as our focus here is purely methodological. We want to exhibit one way of publishing this experiment to maximize reproducibility.

Scenario Description

In general, we define two main successive phases for describing the experiment workflow.




FIGURE 5.5: A glimpse at the Jupyter notebook for R2lab radiomap
click the images to run on mybinder.org

Data collection consists of generating traffic between a Tx and Rx node and measuring the RSSI at the level of the receivers. The procedure used for collecting the measurements is as follows:

1. Create an ad-hoc network at a given frequency.
2. Generate traffic using `ping` with different parameters (see Table. 5.4) from a Tx node while all the other nodes are listening. The nodes take turn in sending the traffic.
3. Capture the incoming packets with `tcpdump` on the monitor interface of each receiving node.
4. Process the dump files in the nodes and collect the generated RSSI to the experimenter's local machine.

Data processing can then occur to provide a higher level of quantitative feedback as a result of the experiment. The processing involves the visualization of some representative outcome of the experiment. With this simplified experiment, we will provide a small visualization tool, that allows a user to select a specific sender node, and that displays the perceived power from all other nodes.

Data collection requires actual physical access to the experimental resources, while data processing can be carried out in the void. Also, several environmental parameters can affect the measurements (see below). Collecting data for many combinations of these environmental parameters can require a vast amount of time, and so it is reasonable to perform data collection incrementally in several complementary measurement campaigns.

The notebook

We provide a Jupyter notebook that can be used to perform either of the stages described above.

Data collection: the environmental parameters of the experiment protocol that can be controlled in the notebook are listed in Table 5.4.

NIC Driver	ath9k	iwlwifi
Transmit Power (in dBm)	5 to 14	0 to 15
Number of antennas	1, 2, 3	3
Physical Tx Rate	1 to 54 Mbps	
Wi-Fi Channel	1-11, 36, 40, 44, 48, 52, 56	

TABLE 5.4: Controllable parameters from the notebook

The data collection code leverages both `nepi-ng` and `rhubarbe`. Data gathered along the various runs accumulates in a simple directory structure, so that data collection is trivially incremental.

Data processing: once performed in python code, data can be plotted interactively right inside the notebook using various visual tools, in the sense that the reader can select either the environmental parameters he is interested in - when available or select a sender node, and in the case of 3D diagrams, can navigate inside the figure space. These possibilities naturally reach way beyond what is usually workable from a static printable paper.

The reader can download the notebook as part of a `git` repository*, and run from anyone's computer. Clickable Figure 5.5 will lead directly to a pre-deployed instance of that notebook, hosted in the `mybinder.org` public infrastructure, as a best-effort attempt to make it usable as seamlessly as possible.

The experimenter can study the impact of parameters on the transmission range and the channel response. Researchers can use the radio maps as a prospection tool for researchers to select the most convenient nodes to be used for their experimentation scenarios by taking advantage of the R2lab's layout, for example, when various

*The notebook source from its GitHub repository

<https://github.com/parmentelat/r2lab-demos/radiomap/radiomap.ipynb>

types of communication scenarios are required, e.g., with line-of-sight, non-line-of-sight, near-line-of-sight. Radio maps can be useful as a diagnostic tool to identify issues with NICs and antennas.

The notebook is part of a git repository that contains other required artifacts. This repository comes with some measurements; in a real experience, this would correspond to the data gained by the author. Thus, in a first step, a reader could quickly reproduce the visualizations of the paper from that data, before she can go further and perform her own data acquisition.

Discussion

Using the runnable paper methodology for publishing results is attractive; However, it is not perfect.

On the pro side, notebooks are a very effective way to convey ideas and implementation. The format provides executable documents allowing authors to have a significant level of flexibility to show the essential pieces of the code and to hide less crucial details in code stored separately.

Still, on the bright side, the current public offering allows combining `github.com` and `mybinder.org` and to offer single-click access to a runnable notebook, that removes the burden of software installation, at the price of more limited features though.

On the downside, however, running the notebook from that public environment is limited to post-processing, as gaining physical access to R2lab requires SSH credentials. Whether the notebook is run from that public spot or local machine, a reservation needs to be obtained before data collection can be carried out.

Sustainability is a significant challenge here. We cannot assume that `mybinder.org` will be up forever; therefore we rely on this platform only as best-effort mode, essentially to smoothen the adoption process of Jupyter, which besides is gaining increasing popularity in several scientific communities. The notebook, in our case, with its depending code and notes, can be considered a regular artifact where its real purpose is to be run locally.

5.5 Thoughts On Improving Reproducibility

ACM Members organized a SIGCOMM Reproducibility workshop to discuss and come up with some guidelines and general recommendations to improve reproducibility in the community. Guided by the participants' diverse experiences, we organized a brainstorming session to tackle some of the most sensitive and challenging aspects of the scientific work that hinder the widespread adoption of reproducibility.

5.5.1 Author incentives

As suggested earlier in the chapter and through the conclusion of the survey conducted in the community, researchers are keen to provide and encourage reproducible research projects. However, given the restraints and the exerted pressure

by funding agencies, they had to prioritize novelty at the expense of reproducibility, which is undervalued and with a low return on investment. Thus, encouraging reproducibility should consider the sensitivities of researchers and avoid antagonizing the community by imposing the reproducibility of each accepted paper. A more positive alternative is to provide an incentive in added recognition to the presented research work. The badging policy proposed by ACM should be applied as an opt-in and focus on helping the volunteer author achieve the highest badge. By doing so, the author gets more recognition from the achieved work and provides the community with reusable tools and artifacts. Producing reproducible and reusable artifacts would help reduce the burden and starting costs for other research projects and foster more collaboration. Promoting more reproducibility in the community should avoid shallow initiatives and adopt a more hands-on approach that comprises guiding the volunteer authors to gain the reproducibility reflexes and introduce reproducibility as an organic part of the research activity.

5.5.2 Reviewer incentives

Unlike classical reviews, reproducibility is challenging and time-consuming; therefore, the process should be straightforward, well-structured, and efficient. For that purpose, the review process should be deployed progressively and only focus on authors that provide artifacts. This approach would allow us to identify the shortcomings and the pitfalls and develop recommendations and an efficient process. For this purpose, during the workshop, participants proposed to keep a volunteering-based approach where authors are encouraged to apply for the badging imposing no other constraints. By doing so, the reviewer will deploy a more focused effort and plan her workload more efficiently. The participants proposed to keep the review committee public and independent dedicated only to reviewing reproducibility and only get involved after the acceptance of the paper; this is motivated by the need for a specific set of skills to fairly assess the reproducibility and avoid problems related to intellectual property infringement and privacy concerns. So, to make this process more immune to possible derives, the review process and the comments are made public in the same spirit as openreviews.

5.5.3 Review Form Design

Reviewing reproducibility is impossible using the current review tool. We have put in place a task force to produce a draft of the review form*. This form would allow a structured review that considers pre-defined criteria and would at the give the appropriate badge to the candidate papers. This initiative would be focused on two main goals, produce a fair review and guide the reviewer in this process in the simplest way possible. We have designed the form to be included in the ACM Artefact Review HOTCRP for assisting reviewers in their task. Even though this form has seen multiple versions since it was introduced in the ACM SIGCOMM Artefact Review, it keeps four main aspects. The form comprises closed-ended and open-ended questions to the reviewer axed on presenting the key aspects of the provided artifact,

*Artifact review form <https://github.com/naoufal51/Review-Template.git>

to assess the completeness of the provided information, the artifact consistency with the paper content and a recommendation (Artefacts Available, Artefacts Evaluated - Functional, Artefacts Evaluated - Reusable.)

5.5.4 Sharing Artifacts

Verifying reproducibility is conditioned by the quality of the artifacts' presentation. As the review process is limited in time and is already strenuous enough, the community should adopt a standardized method to present artifacts and should privilege using benchmarking datasets or proven tools when it is possible. The conducted survey shows that different communities have different challenges about artifacts sharing. However, most suffer from underserved documentation or lack of structure, which hinders first the reviewing process and in-fine the reproducibility process. Given that most research projects are continuously updated, a versioning effort in artifacts is critical for conserving coherence, especially when the modifications change the system's behavior. The documentation should be accurate, organized, and intelligible to simplify the evaluation process. The authors should provide the level of detail necessary to reproduce the results, which could be identified by asking lab mates to redo the experiments by only using the description provided. Through this exercise, the authors can enhance the quality of the description and provide troubleshooting procedures if problems occur. To ensure a longer reproducibility cycle, authors should be able to motivate the choice of tools according to the needed functionalities; Hence, avoiding problems related to packages deprecation, end-of-life operating systems, non-commercially available equipment, etc.

After badging, the artifacts should be hosted in long-term storage and adopt a versioning procedure that would tag the version used for reproducing the results. With new features, authors should use other dedicated branches, and the documentation should stay consistent with the published source code.

5.6 Conclusion

Through this work, we address some reproducibility challenges in our community by first providing a comprehensive understanding of the current state and also by proposing methodologies to improve the reusability of the artifact, especially in the wireless community. Through this study, we have shown that the reproducibility problem is not only technical but also societal, where a system of incentives for all the involved parties (authors and reviewers) is crucial. Moreover, we have shown that there is a dire need for well-designed and well-dimensioned testbeds capable of providing a comprehensive benchmarking base. And finally, we have shown that there also a need for a new approach to publishing scientific contribution which could enhance the understanding of the contribution by the reviewers and readers alike, e.g, runnable papers, notebooks. Even though these few steps represent a good start, achieving reproducibility is still laborious to achieve, especially for communities relying on specialized hardware, cutting-edge equipment, or fine-grained control over experimentation parameters.

Appendix

5.A OpenRF reproducibility: Detailed verification process

When trying to reproduce the OpenRF results, our symptom was that no matter how hard we try to change, in our Intel WiFi Link 5300 card, the CSI matrix used for emission, we could measure no noticeable change in the emitted waves. In this appendix, we describe the technical steps we have taken to get our Intel cards to take our modified CSI matrix into account - but to no avail. We are thus left with tracking down data interchange between kernel space - namely the modified `iwlwifi` driver - and userland, and back.

Upon the reception of a packet, a CSI report is sent from kernel to userspace using a Netlink socket. Then the CSI report is decoded, the CSI matrix is extracted, and a singular value decomposition of the matrix is performed. This decomposition results in the computation of the V matrix used for either beamforming, interference nulling or alignment. This matrix is compressed using specific rotations respecting the IEEE 802.11n recommendation for explicit beamforming. These operations are conducted through the MATLAB function `precod_channel` in userspace. After computing the V matrix, we need to send it to the wireless card. Two structures are used when sending this matrix back to the kernel: the MIMO Control subfields and the Compressed Beamforming Report field (`iwl-command.h`), both defined in the standard [73]. More precisely, the MIMO Control field is implemented to handle beamforming feedback information, and the Compressed Beamforming Report field is used to carry explicit feedback in angles to be used by a transmitter when computing the corresponding steering matrix.

After storing the matrix in the structures, a netlink socket is used to send this CSI report from userland to kernel space with a specific ID (`REPLY_BFER_VCOMP_CONFIG = 0xbc`). This operation is implemented in the `test_send_weight_matrix.c` function. When the CSI report structure is received in kernel space, it is used in function (`iwlagn_send_bfer_config`) defined in file `iwl-agn-lib.c`. In this function we are using the command `iwl_dvm_send_cmd_pdu` for sending this CSI report to firmware. When checking the driver messages after the beamforming operation is completed, we obtain these messages:

```
$ dmesg
...
iwlwifi 0000:04:00.0: Setting beamforming matrix
iwlwifi 0000:04:00.0: Set bf: Returned (0)
iwlwifi 0000:04:00.0: In iwlagn_send_rxon_assoc_wsdn
```

These driver messages indicate that the beamforming was performed and that the wireless card is already using the compressed beamforming matrix instead of the indirect mapping matrix. However, despite all that the SNR plots showed no changes in the signal levels. From our correspondence with the authors, we applied a transmitter reset to verify if it could fix the problem, but it has no effect either.

```
$ sudo bash -c "echo 0 >
/sys/kernel/debug/ieee80211/phy0/iwlwifi/debug/bf_flag;
echo 1 >
/sys/kernel/debug/ieee80211/phy0/iwlwifi/debug/bf_flag"
```

5.B Artifact Evaluation Review Form

1. Reviewer confidence

- No familiarity
- Some familiarity
- Knowledgeable
- Agree
- Expert

2. Type of Artifact

- Dataset
- Script/Tools
- New Artefact
- Extension of Existing Artefact

3. What is the degree of efforts to obtain the artefact?

Rate from 0% to 100%, where 0% means NOT AVAILABLE; 20%, VERY HARD TO OBTAIN up to 100% meaning VERY EASY TO OBTAIN.

- 0%
- 20%
- 40%
- 60%
- 80%
- 100%

4. Comment about the degree of efforts to obtain the artefact

E.g., did the paper have explicit pointer to the artefact? The artefact is simply not available. I need to register to obtain the artefact. The artefact is a testbed and I cannot use it. The artefact is actually hardware, I cannot obtain it. The artefact is licensed and I should pay...

5. Do you consider that the nature of the artefact (or any component of it) may raise (non-technical) issues? Of which Nature?

E.g., ethical issues; non FRAND - Fair, Reasonable and Non Discriminatory - licenses; environmental issues; safety issues;...

6. Has the artefact been deposited in what can be reasonably considered the most appropriate available repository, which ensures perennially availability?

7. In relation to the previous point, how long do you expect or estimate the artefact will be available:

- Less than 1 year
- Between 2 and 5 years
- Between 5 and 10 years
- More than 10 years

8. Are the instructions on how to install and make the artefact functional provided?

Rate from 0% to 100%, where 0% means NOT INSTRUCTION PROVIDED; 20% VERY HARD TO UNDERSTAND up to 100% meaning VERY EASY TO INSTALL AND RUN.

In case of a data set the question has to be answered considering 0% means NOT DATA FORMAT DESCRIPTION; 20% VERY HARD TO UNDERSTAND THE DATA FORMAT up to 100% meaning VERY EASY TO UNDERSTAND AND USE THE DATA.

- 0%
- 20%
- 40%
- 60%
- 80%
- 100%

9. Comment about the efforts to obtain, install, and use the artefact

E.g., very limited instructions were given; I abandoned at some point because it was not clear what to do; it was difficult to figure out what exactly are the steps to follow; easy to follow but some degree of familiarity is necessary; very simple step-by-step instructions have been provided; piece of cake, the authors provide a guide for dummies; I had nothing to do I just downloaded a VM image.

10. If the artefact reached the "functional" level, meaning that it seems to function correctly, did the authors provide all of the information and meta-data needed for others to reuse the artefact?

Rate from 0% to 100%, where 0% means NO DOCUMENTATION PROVIDED; 20% DOCUMENTATION HARD TO UNDERSTAND up to 100% meaning DOCUMENTATION VERY CLEAR AND EASY TO UNDERSTAND.

- 0%
- 20%
- 40%
- 60%
- 80%
- 100%

11. Comment about the quality of the documentation
E.g., documentation is complete and clear; documentation is lacking fundamental explanations; authors even explain in details how to reproduce the paper; documentation is not provided; documentation is lacking the following basic information;...
12. If the artefact reached the "functional" level, meaning that it seems to function correctly, can the artefact be easily reused and/or repurposed?
Rate from 0% to 100%, where 0% means IMPOSSIBLE TO REUSE; 20% VERY LIMITED REUSABILITY up to 100% meaning ARTEFACT VERY EASY TO REUSE/REPURPOSE.
 - 0%
 - 20%
 - 40%
 - 60%
 - 80%
 - 100%
13. Comment about the degree of reusability/repurposing
E.g., I already have some ideas how to reuse it; the artefact does not sufficiently feature-rich to be reused; I can see opportunities to use it for; the artefact is functional but reusing it is kind of a hassle...
14. Artefact Weight in the Research Paper
Considering the portion of the paper supported by artefacts (included other artefacts not proposed by the authors), what share the key results therein are clearly directly linked to the provided artefact? Please choose between 0% and 100%. Where 0% means no key results are directly related to the artefact and 100% means that all of the key results therein are clearly directly obtained with the artefact.
 - 0%
 - 20%
 - 40%
 - 60%
 - 80%
 - 100%
15. Badge Recommendation based on the following ACM Badge definitions
 - Artefacts Available: Author-created artefacts relevant to this paper is publicly available. A DOI or link to this repository along with a unique identifier for the object is provided.
 - Artefacts Evaluated ? Functional: The artefacts associated with the research are found to be documented, consistent, complete, exercisable, and include appropriate evidence of verification and validation.
 - Artefacts Evaluated ? Reusable: The artefacts associated with the paper are of a quality that significantly exceeds minimal functionality. That is, they have all the qualities of the Artefacts Evaluated ? Functional level, but, in addition, they are very carefully documented and well-structured to the extent that reuse and repurposing is facilitated. In particular, norms and standards of the research community for artefacts of this type are strictly adhered to.

- No Badge
- Artefact Evaluated - Available
- Artefact Evaluated - Functional
- Artefact Evaluated - Reusable

16. Comments to the authors

17. Confidential comments to the AEC

18. Rebuttal

In case of a claim made by the authors to dispute the badge assignment, would you like to participate in the rebuttal phase?

- Yes
- No

Chapter 6

Conclusion

In this thesis, we have taken on the challenge of providing robust and accurate localization and orientation estimation systems by exploiting antenna array signal processing and data-driven approaches. This work specifically contributes to RF-based sensing:

- **Estimating orientation with COTS Wi-Fi devices:** In Chapter 2, we describe an orientation estimation method called ORION; a system that can estimate the orientation of a Wi-Fi terminal using unmodified off-the-shelf Wi-Fi devices. We designed our approach by enabling the angle of arrival and angle of departure estimation, which are primitives essential to compute the orientation of a target terminal by the infrastructure. As phase inconsistencies plague COTS hardware, we propose a specific procedure to mitigate the phase shifts influencing the orientation estimate. We preferred using a joint estimation framework that provides a more robust and accurate estimation of both AoA and AoD.
- **Joint Localization and Range Extension for LP-WAN:** In Chapter 3, we have presented a novel MIMO-LPWAN system that enables a precise AoA based localization system on top of extending the range between an LPWAN gateway and an end node. We propose a system design that makes it possible to achieve accurate position estimation by adopting model-based MIMO radar techniques. Our study also shows that using beamforming and coherent combining across the signals impinging on the antenna array element increases the gain at both the gateway and the end device.
- **Robust Passive Localization Using Wi-Fi CSI: A Data-Driven Paradigm:** In Chapter 4, we have presented both a new large-scale dataset called RFLoc for Wi-Fi CSI-based passive localization and a data-driven passive localization system called DPLoc. Our design takes advantage of the generative model paradigm to address the Non-Line-of-Sight and environment-dependency issues. By using this approach, we are trying to address the problem of domain dependency, which translates into a generalization problem of trained models. Our design shows that we can mitigate the environmental effect on position estimation by adopting transfer learning techniques. We predict that the RFLoc campaign can inspire more open-source, data-driven research in wireless sensing.

- **Reproducibility for Wireless Networking Experiments:** In Chapter 5, we opened a discussion about reproducibility challenges in networking and wireless networking communities. Addressing this challenge is far exceeding this thesis. However, we attempt to understand the situation by assessing the blocking factors while providing our take on reproducibility according to our requirements and limitations. We tried to present the point-of-view of the different actors of the community (Authors, Reviewer, Conference organizers) and highlighting their struggles to either provide or review or organize reproducibility and the axes of enhancements.

6.1 Future Work

In this thesis, we have presented contributions for passive and active localization using COTS RF devices. Wireless remote sensing has advantages compared to optical methods in terms of occlusion, range, and privacy. However, the level of accuracy using the existing models is in the centimeter level, which is not enough to be a viable alternative for application demanding a high level of accuracy, for instance, human pose estimation for virtual reality application and augmented reality.

To address these challenges, data-driven approaches and particularly deep learning could play a key role in providing enhanced accuracy. However, deep learning tools are tailored for computer vision applications, which makes it challenging to adapt to our wireless sensing problem. In this sense, we project on working on these aspects:

- **Novel deep learning model custom-tailored for wireless signals** To cater to the sparse nature of wireless signals because of the limited number of "sensors" (antennas), which contrasts with the models built on exploiting the dense nature of images. This approach will focus on **modeling the propagation environment** rather than using models designed for exploiting the spatial distribution of pixels.
- **Novel deep learning model for Super-resolution passive localization in true Non-Line-of-sight conditions.** To achieve a spatial resolution higher than first-order geometry radar methods, even in NLoS conditions and under Wi-Fi hardware constraints (few antennas and limited bandwidth). This solution will need to capture the geometrical relation between the target's position and the sophisticated reflection patterns.
- **Weakly supervised or unsupervised learning for environment-independent passive localization.** To mitigate the effect of the environments' multipath on the CSI features associated with the target's location. The CSI can vary depending on the ambient reflectors' geometry, orientation, and reflectivity. Therefore, the passive localization model must be robust across different environments while avoiding the burden of collecting location-labeled CSI samples, carried by human operators.

Reproducibility In this thesis, we have also presented our reproducibility project. To promote more reproducibility in the community while reducing the strain on the conference organizers and reviewers, we are working on changing the format of the

ACM reproducibility committee to be more representative of the topics discussed in the community. We are planning on simplifying the review process and make it more inclusive, which means more interaction between authors and reviewers with a common goal of providing a quality artifact reusable by the community.

Bibliography

- [1] *1,500 scientists lift the lid on reproducibility Reproducibility*. <https://www.nature.com/news/1-500-scientists-lift-the-lid-on-reproducibility-1.19970>. Accessed: 2019-05-12.
- [2] *ACM Digital Library*. <https://dl.acm.org/dl.cfm>.
- [3] *ACM IMC 2017 – Call for Papers*. <https://conferences.sigcomm.org/imc/2017/cfp/>.
- [4] Fadel Adib, Zachary Kabelac, and Dina Katabi. “Multi-Person Localization via RF Body Reflections”. In: *Proceedings of USENIX NSDI*. 2015.
- [5] Fadel Adib and Dina Katabi. “See Through Walls with WiFi!” In: *Proc. of ACM SIGCOMM*. 2013.
- [6] Fadel Adib et al. “3D Tracking via Body Radio Reflections”. In: *Proceedings of USENIX NSDI*. 2014.
- [7] Cedric Adjih et al. “FIT IoT-LAB: A large scale open experimental IoT testbed”. In: *Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on*. IEEE. 2015, pp. 459–464.
- [8] Alexander Afanasyev, Ilya Moiseenko, Lixia Zhang, et al. “ndnSIM: NDN simulator for NS-3”. In: *University of California, Los Angeles, Tech. Rep 4* (2012).
- [9] A. Mariakakis et al. “SAIL: Single Access Point-based Indoor Localization”. In: *ACM MobiSys*. Bretton Woods, New Hampshire, USA, 2014, pp. 315–328. ISBN: 978-1-4503-2793-0. DOI: 10.1145/2594368.2594393. URL: <http://doi.acm.org/10.1145/2594368.2594393>.
- [10] AFZAL et al. “Multi-Magnetometer Based Perturbation Mitigation for Indoor Orientation Estimation”. In: *Navigation* 58.4 (2011), pp. 279–292. ISSN: 2161-4296. DOI: 10.1002/j.2161-4296.2011.tb02586.x. URL: <http://dx.doi.org/10.1002/j.2161-4296.2011.tb02586.x>.
- [11] D. Halperin et al. “Tool Release: Gathering 802.11N Traces with Channel State Information”. In: *ACM CCR* 41.1 (2011), pp. 53–53. ISSN: 0146-4833. DOI: 10.1145/1925861.1925870. URL: <http://doi.acm.org/10.1145/1925861.1925870>.
- [12] D. Vasisht et al. “Decimeter-Level Localization with a Single WiFi Access Point”. In: *ACM NSDI*. Santa Clara, CA, 2016, pp. 165–178. ISBN: 978-1-931971-29-4.
- [13] E. C. L. Chan et al. “Orientation-based Wi-Fi positioning on the google Nexus One”. In: *IEEE WiMob*, 2010, pp. 392–397. DOI: 10.1109/WIMOB.2010.5645038.

- [14] E. Perahia et al. "Investigation into the Doppler Component of the IEEE 802.11n Channel Model". In: 2010, pp. 1–5. DOI: [10.1109/GLOCOM.2010.5684207](https://doi.org/10.1109/GLOCOM.2010.5684207).
- [15] H. Mohd et al. "Indoor Human Localization with Orientation Using WiFi Fingerprinting". In: ACM ICUIMC. Siem Reap, Cambodia, 2014. ISBN: 978-1-4503-2644-5. DOI: [10.1145/2557977.2557980](https://doi.org/10.1145/2557977.2557980). URL: <http://doi.acm.org/10.1145/2557977.2557980>.
- [16] J. Gjengset et al. *ArrayPhaser: Enabling Signal Processing on WiFi Access Points*. RN, 04. 2014.
- [17] J. Xiong et al. "ArrayTrack: A Fine-Grained Indoor Location System". In: ACM NSDI. Lombard, IL, USA, 2013.
- [18] J.J. Leonard et al. "Mobile robot localization by tracking geometric beacons". In: *IEEE T-RO* 7.3 (1991), pp. 376–382. ISSN: 1042-296X. DOI: [10.1109/70.88147](https://doi.org/10.1109/70.88147).
- [19] K. Liu et al. "Attitude determination for unmanned aerial vehicles via an antenna array". In: *IEEE WSA*. 2012.
- [20] M. Kotaru et al. "SpotFi: Decimeter Level Localization Using WiFi". In: ACM SIGCOMM. London, UK, 2015.
- [21] S. Souvik et al. "You Are Facing the Mona Lisa: Spot Localization Using PHY Layer Information". In: *ACM MobiSys*. Low Wood Bay, Lake District, UK, 2012, pp. 183–196. ISBN: 978-1-4503-1301-8. DOI: [10.1145/2307636.2307654](https://doi.org/10.1145/2307636.2307654). URL: <http://doi.acm.org/10.1145/2307636.2307654>.
- [22] T.F.K. Cordeiro et al. "Improved Kalman-based attitude estimation framework for UAVs via an antenna array". In: *DSP* 59 (2016), pp. 49 –65. ISSN: 1051-2004. DOI: <http://dx.doi.org/10.1016/j.dsp.2016.07.006>. URL: <http://www.sciencedirect.com/science/article/pii/S1051200416300914>.
- [23] Y. Xie et al. "Precise Power Delay Profiling with Commodity WiFi". In: *ACM Mobicom*. Paris, France, 2015, pp. 53–64. ISBN: 978-1-4503-3619-2. DOI: [10.1145/2789168.2790124](https://doi.org/10.1145/2789168.2790124).
- [24] Y. Zheng et al. "From RSSI to CSI: Indoor Localization via Channel Response". In: *ACM Computer Survey* 46.2 (2013). ISSN: 0360-0300. DOI: [10.1145/2543581.2543592](https://doi.org/10.1145/2543581.2543592). URL: <http://doi.acm.org/10.1145/2543581.2543592>.
- [25] Z. Zhang et al. "I Am the Antenna: Accurate Outdoor AP Location Using Smartphones". In: ACM MobiCom. Las Vegas, Nevada, USA, 2011. ISBN: 978-1-4503-0492-4. DOI: [10.1145/2030613.2030626](https://doi.org/10.1145/2030613.2030626). URL: <http://doi.acm.org/10.1145/2030613.2030626>.
- [26] J. Amann et al. "Mission Accomplished? HTTPS Security after DigiNotar". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [27] D. Anderson et al. "A High-Performance Algorithm for Identifying Frequent Items in Data Streams". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.

- [28] O. Aponte and P. Mendes. “Now@ - Content Sharing Application over NDN”. In: *The 4th ACM Conference on Information-Centric Networking – ICN’17*. Berlin, Germany, 2017.
- [29] Martin Arjovsky, Soumith Chintala, and Léon Bottou. *Wasserstein GAN*. 2017. eprint: [arXiv:1701.07875](https://arxiv.org/abs/1701.07875).
- [30] O. Ascigil et al. “A Native Content Discovery for the Information-centric Networks”. In: *The 4th ACM Conference on Information-Centric Networking – ICN’17*. Berlin, Germany, 2017.
- [31] H. Augé et al. “Virtualized ICN (vICN): Towards a Unified Network Virtualization Framework for ICN Experimentation”. In: *The 4th ACM Conference on Information-Centric Networking – ICN’17*. Berlin, Germany, 2017.
- [32] J. L. Ayala. “Code Ocean Is Live: Upload Your Algorithms”. In: *IEEE Design Test* 34.3 (2017), pp. 108–109. ISSN: 2168-2356. DOI: [10.1109/MDAT.2017.2692207](https://doi.org/10.1109/MDAT.2017.2692207).
- [33] Emmanuel Baccelli et al. “RIOT OS: Towards an OS for the Internet of Things”. In: *Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on*. IEEE. 2013, pp. 79–80.
- [34] Constantine A Balanis. *Antenna theory: analysis and design*. John Wiley & Sons, 2016.
- [35] R. Ben Basat et al. “Constant Time Updates in Hierarchical Heavy Hitters”. In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [36] M. L. Bencheikh and Y. Wang. “Joint DOD-DOA estimation using combined ESPRIT-MUSIC approach in MIMO radar”. In: *Electronics Letters* 46.15 (2010), pp. 1081–1083. ISSN: 0013-5194. DOI: [10.1049/el.2010.1195](https://doi.org/10.1049/el.2010.1195).
- [37] J. De Benedetto, B. Wissingh, and X. Fu. “ICN personalized global-scale testbed using GTS”. In: *The 4th ACM Conference on Information-Centric Networking – ICN’17*. Berlin, Germany, 2017.
- [38] Yonathan Bleyfuesz et al. *Demo: Using nepi-ng for Mesh Networks Experiments*. WINTECH 2018 - 12th ACM International Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization. Poster. Nov. 2018. URL: <https://hal.inria.fr/hal-01869979>.
- [39] Brekelman. *Brekel Pro Body V2*. 2019. URL: <https://brekel.com/brekel-pro-body-v2/>.
- [40] Quantenna QT3840BC Chipset. *Quantenna*. 2014. URL: <http://www.quantenna.com/wp-content/uploads/2016/12/QV840V3.3.pdf> (visited on 03/13/2019).
- [41] I. Cho, K. Jang, and D. Han. “Credit-Scheduled Delay-Bounded Congestion Control for Datacenters”. In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [42] K. Cho. “Recursive Lattice Search: Hierarchical Heavy Hitters Revisited”. In: *ACM Internet Measurement Conference 2017 – IMC’17*. London, UK, 2017.
- [43] Francois Chollet. “Xception: Deep Learning With Depthwise Separable Convolutions”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.

- [44] S. Colieri et al. "A study of channel estimation in OFDM systems". In: *Proceedings IEEE 56th Vehicular Technology Conference*. Vol. 2. 2002, 894–898 vol.2. DOI: [10.1109/VETECONF.2002.1040729](https://doi.org/10.1109/VETECONF.2002.1040729).
- [45] Q. De Coninck and O. Bonaventure. "Multipath QUIC: Design and Evaluation". In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT'17*. Seoul/Icheon, South Korea, 2017.
- [46] J. DeBlasio et al. "Tripwire: Inferring Internet Site Compromise". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [47] A.P. Dempster, N.M. Laird, and D.B. Rubin. "Maximum likelihood from incomplete data via the EM algorithm". In: *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B* 39.1 (1977), pp. 1–38.
- [48] J. Deng et al. "ImageNet: A Large-Scale Hierarchical Image Database". In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2009.
- [49] Adwait Dongare et al. "Charm: Exploiting Geographical Diversity Through Coherent Combining in Low-power Wide-area Networks". In: *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks*. IPSN '18. Porto, Portugal: IEEE Press, 2018, pp. 60–71. ISBN: 978-1-5386-5298-5. DOI: [10.1109/IPSNS.2018.00013](https://doi.org/10.1109/IPSNS.2018.00013). URL: <https://doi.org/10.1109/IPSNS.2018.00013>.
- [50] Vincent Dumoulin and Francesco Visin. "A guide to convolution arithmetic for deep learning". In: *CoRR* abs/1603.07285 (2016).
- [51] S. Dwyer and P. Mendes. "Named-Data Networking in Opportunistic Networks". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.
- [52] Rashad Eletreby et al. "Empowering Low-Power Wide Area Networks in Urban Settings". In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. SIGCOMM '17. Los Angeles, CA, USA: ACM, 2017, pp. 309–321. ISBN: 978-1-4503-4653-5. DOI: [10.1145/3098822.3098845](https://doi.org/10.1145/3098822.3098845). URL: <http://doi.acm.org/10.1145/3098822.3098845>.
- [53] Bernat Carbones Fargas and Martin Nordal Petersen. "GPS-free Geolocation using LoRa in Low-Power WANs". In: *Proceedings of 2017 Global Internet of Things Summit (GIoTS)*. IEEE, 2017. DOI: [10.1109/GIOTS.2017.8016251](https://doi.org/10.1109/GIOTS.2017.8016251).
- [54] J. A. Fessler and A. O. Hero. "Space-alternating generalized expectation-maximization algorithm". In: *IEEE Transactions on Signal Processing* 42.10 (1994), pp. 2664–2677. ISSN: 1053-587X. DOI: [10.1109/78.324732](https://doi.org/10.1109/78.324732).
- [55] J. A. Fessler and A. O. Hero. "Space-alternating generalized expectation-maximization algorithm". In: *IEEE Transactions on Signal Processing* 42.10 (1994).
- [56] Matthias Flittner et al. "A Survey on Artifacts from CoNEXT, ICN, IMC, and SIGCOMM Conferences in 2017". In: *SIGCOMM Comput. Commun. Rev.* 48.1 (Apr. 2018), pp. 75–80. ISSN: 0146-4833. DOI: [10.1145/3211852.3211864](https://doi.org/10.1145/3211852.3211864). URL: <http://doi.acm.org/10.1145/3211852.3211864>.

- [57] R. Fontugne, E. Aben, and C. Pelsser. "Pinpointing Delay and Forwarding Anomalies Using Large-Scale Traceroute Measurements". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [58] Ghose, Anindya and Li, Beibei and Liu, Siyuan. *Mobile Targeting Using Customer Trajectory Patterns*. 2016. URL: <https://ssrn.com/abstract=2962044>.
- [59] Y. Gilad, O. Sagga, and S. Goldberg. "MaxLength Considered Harmful to the RPKI". In: *The 13th International Conference on emerging Networking Experiments and Technologies – CoNEXT'17*. Seoul/Icheon, South Korea, 2017.
- [60] M.H Golbon-Haghighi. "Beamforming in Wireless Networks." In: *IntechOpen*, Dec. 2016. Chap. 8, pp. 163–192. ISBN: 978-953-51-2833-5. DOI: [10.5772/66399](https://doi.org/10.5772/66399).
- [61] B. Goodchild et al. "The Record Route Option Is An Option!" In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [62] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016.
- [63] Ian J. Goodfellow et al. *Generative Adversarial Networks*. 2014. eprint: [arXiv:1406.2661](https://arxiv.org/abs/1406.2661).
- [64] *gr-lora gr-lora*. <https://github.com/rpp0/gr-lora>. Accessed: 2019-05-12.
- [65] Ishaan Gulrajani et al. "Improved Training of Wasserstein GANs". In: *CoRR abs/1704.00028* (2017). arXiv: [1704.00028](https://arxiv.org/abs/1704.00028). URL: <http://arxiv.org/abs/1704.00028>.
- [66] C. Gündogan et al. "Information-Centric Networking for the Industrial IoT". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.
- [67] Tim Hadwen et al. "Energy efficient LoRa GPS tracker for dementia patients". In: *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2017. DOI: [10.1109/EMBC.2017.8036938](https://doi.org/10.1109/EMBC.2017.8036938).
- [68] O. Hahm et al. "Low-power Internet of Things with NDN and Cooperative Caching". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.
- [69] P. Hannan. "The element-gain paradox for a phased-array antenna". In: *IEEE Transactions on Antennas and Propagation* 12.4 (1964), pp. 423–433. ISSN: 0018-926X. DOI: [10.1109/TAP.1964.1138237](https://doi.org/10.1109/TAP.1964.1138237).
- [70] S. M. Hatem et al. "Rehabilitation of Motor Function after Stroke: A Multiple Systematic Review Focused on Techniques to Stimulate Upper Extremity Recovery". In: *Front Hum Neurosci*. (2016).
- [71] Kaiming He et al. "Deep Residual Learning for Image Recognition". In: *CoRR abs/1512.03385* (2015). arXiv: [1512.03385](https://arxiv.org/abs/1512.03385). URL: <http://arxiv.org/abs/1512.03385>.

- [72] T. Holterbach et al. "SWIFT: Predictive Fast Reroute". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [73] "IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput". In: *IEEE Std 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009)* (2009), pp. 1–565. DOI: [10.1109/IEEESTD.2009.5307322](https://doi.org/10.1109/IEEESTD.2009.5307322).
- [74] Sergey Ioffe and Christian Szegedy. "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift". In: *CoRR abs/1502.03167* (2015). arXiv: [1502.03167](https://arxiv.org/abs/1502.03167). URL: <http://arxiv.org/abs/1502.03167>.
- [75] C. Iordanou et al. "Who is Fiddling with Prices? Building and Deploying a Watchdog Service for E-commerce". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [76] Ming Jin, Guisheng Liao, and Jun Li. "Joint {DOD} and {DOA} estimation for bistatic {MIMO} radar". In: *Signal Processing* 89.2 (2009), pp. 244–251. ISSN: 0165-1684. DOI: <http://dx.doi.org/10.1016/j.sigpro.2008.08.003>. URL: [//www.sciencedirect.com/science/article/pii/S0165168408002673](http://www.sciencedirect.com/science/article/pii/S0165168408002673).
- [77] Kiran Joshi et al. "WiDeo: Fine-grained Device-free Motion Tracing Using RF Backscatter". In: *Proceedings of USENIX NSDI*. 2015.
- [78] S. Kassing et al. "Beyond fat-trees without antennae, mirrors, and disco-balls". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [79] P. Kietzmann et al. "The Need for a Name to MAC Address Mapping in NDN: Towards Quantifying the Resource Gain". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.
- [80] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *CoRR abs/1412.6980* (2014). arXiv: [1412.6980](https://arxiv.org/abs/1412.6980). URL: <http://arxiv.org/abs/1412.6980>.
- [81] V. Kotronis et al. "Shortcuts through Colocation Facilities". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [82] M. Krol and I. Psaras. "NFaaS: Named Function as a Service". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.
- [83] Swarun Kumar et al. "Bringing Cross-layer MIMO to Today's Wireless LANs". In: *SIGCOMM Comput. Commun. Rev.* 43.4 (Aug. 2013), pp. 387–398. ISSN: 0146-4833. DOI: [10.1145/2534169.2486034](https://doi.org/10.1145/2534169.2486034). URL: <http://doi.acm.org/10.1145/2534169.2486034>.
- [84] A. Langley et al. "The QUIC Transport Protocol: Design and Internet-Scale Deployment". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.

- [85] W.K. Leong, W. Zixiao, and B. Leong. "TCP Congestion Control Beyond Bandwidth-Delay Product for Mobile Cellular Networks". In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT'17*. Seoul/Icheon, South Korea, 2017.
- [86] F. Li et al. "lib-erate, (n): A library for exposing (traffic-classification) rules and avoiding them efficiently". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [87] Y.-S. Lim et al. "ECF: An MPTCP Path Scheduler to Manage Heterogeneous Paths". In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT'17*. Seoul/Icheon, South Korea, 2017.
- [88] M. Liu et al. "A Unified Data Structure of Name Lookup for NDN Data Plane". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.
- [89] LoRa LoRa. <https://www.semtech.com/uploads/documents/an1200.22.pdf>. Accessed: 2019-05-12.
- [90] M. Luckie and R. Beverly. "The Impact of Router Outages on the AS-level Internet". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [91] Mohamed Naoufal Mahfoudi et al. "Deploy a 5G Network in Less Than 5 Minutes". In: *Proceedings of the SIGCOMM Posters and Demos*. SIGCOMM Posters and Demos '17. Los Angeles, CA, USA: ACM, 2017, pp. 113–115. ISBN: 978-1-4503-5057-0. DOI: [10.1145/3123878.3132005](https://doi.org/10.1145/3123878.3132005). URL: <http://doi.acm.org/10.1145/3123878.3132005>.
- [92] Mohamed Naoufal Mahfoudi et al. "Joint range extension and localization for LPWAN". In: *Internet Technology Letters* 0,ja (). e120 ITL-19-0043.R1, e120. DOI: [10.1002/itl2.120](https://doi.org/10.1002/itl2.120). eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/itl2.120>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/itl2.120>.
- [93] Mohamed Naoufal Mahfoudi et al. "ORION: Orientation Estimation Using Commodity Wi-Fi". In: *Workshop on Advances in Network Localization and Navigation (ANLN)*. Paris, France, May 2017. URL: <https://hal.archives-ouvertes.fr/hal-01424239>.
- [94] H. Mao, R. Netravali, and M. Alizadeh. "Neural Adaptive Video Streaming with Pensieve". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [95] W.M. Mellette et al. "RotorNet: A Scalable, Low-complexity, Optical Datacenter Network". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [96] X. Mi et al. "An Empirical Characterization of IFTTT: Ecosystem, Usage, and Performance". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [97] P. Moll, J. Janda, and H. Hellwagner. "Adaptive Forwarding of Persistent Interests in Named Data Networking". In: *The 4th ACM Conference on Information-Centric Networking – ICN'17*. Berlin, Germany, 2017.

- [98] T. K. Moon. "The expectation-maximization algorithm". In: *IEEE Signal Processing Magazine* 13.6 (1996), pp. 47–60. ISSN: 1053-5888. DOI: [10.1109/79.543975](https://doi.org/10.1109/79.543975).
- [99] M. Müller et al. "Recursives in the Wild: Engineering Authoritative DNS Servers". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [100] Vinod Nair and Geoffrey E. Hinton. "Rectified Linear Units Improve Restricted Boltzmann Machines". In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML'10. Haifa, Israel: Omnipress, 2010, pp. 807–814. ISBN: 978-1-60558-907-7. URL: <http://dl.acm.org/citation.cfm?id=3104322.3104425>.
- [101] J.A. Nanzer. *Microwave and Millimeter-wave Remote Sensing for Security Applications*. Artech House Remote Sensing Library. Artech House, 2012. ISBN: 9781608071722.
- [102] Richard van Nee. "802.11n: The Global Wireless LAN Standard". In: *Globalization of Mobile and Wireless Communications: Today and in 2020*. Ed. by Ramjee Prasad et al. Dordrecht: Springer Netherlands, 2011, pp. 103–118. ISBN: 978-94-007-0107-6. DOI: [10.1007/978-94-007-0107-6_8](https://doi.org/10.1007/978-94-007-0107-6_8). URL: https://doi.org/10.1007/978-94-007-0107-6_8.
- [103] Farzaneh Pakzad et al. "R2Lab Testbed Evaluation for Wireless Mesh Network Experiments". In: *International Telecommunication Networks and Applications Conference - ITNAC '18*. Sydney, Australia, Nov. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01968766>.
- [104] Thierry Parmentelat et al. "Nepi-ng: An Efficient Experiment Control Tool in R2Lab". In: *Proceedings of the 12th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*. WINTECH '18. New Delhi, India: ACM, 2018, pp. 56–65. ISBN: 978-1-4503-5930-6. DOI: [10.1145/3267204.3267213](https://doi.org/10.1145/3267204.3267213). URL: <http://doi.acm.org/10.1145/3267204.3267213>.
- [105] Roger Pierre and Fabris Hoefel. "Ieee 802.11n: on Performance of Channel Estimation Schemes over Ofdm Mimo Spatially- Correlated Frequency Selective Fading Tgn Channels". In: 2012.
- [106] Nico Podevijn et al. "TDoA-Based Outdoor Positioning with Tracking Algorithm in a Public LoRa Network". In: *Wireless Communications and Mobile Computing*. Hindawi, 2018. DOI: [10.1155/2018/1864209](https://doi.org/10.1155/2018/1864209).
- [107] Kun Qian et al. "Widar: Decimeter-Level Passive Tracking via Velocity Monitoring with Commodity Wi-Fi". In: *Proceedings of ACM MobiHoc*. 2017.
- [108] Kun Qian et al. "Widar2.0: Passive Human Tracking with a Single Wi-Fi Link". In: *Proceedings of ACM MobiSys*. 2018.
- [109] Alec Radford, Luke Metz, and Soumith Chintala. "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks". In: *CoRR abs/1511.06434* (2015). arXiv: [1511.06434](https://arxiv.org/abs/1511.06434). URL: <http://arxiv.org/abs/1511.06434>.

- [110] U. Raza, P. Kulkarni, and M. Sooriyabandara. “Low Power Wide Area Networks: An Overview”. In: *IEEE Communications Surveys Tutorials* 19.2 (2017), pp. 855–873. ISSN: 1553-877X. DOI: [10.1109/COMST.2017.2652320](https://doi.org/10.1109/COMST.2017.2652320).
- [111] A. Razaghpanah et al. “Studying TLS Usage in Android Apps”. In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17*. Seoul/Icheon, South Korea, 2017.
- [112] Ramon dos Reis Fontes et al. “How far can we go? Towards Realistic Software-Defined Wireless Networking Experiments”. In: *The Computer Journal (Oxford)* (June 2017). URL: <https://hal.inria.fr/hal-01480973>.
- [113] Ettus Research. *Synchronization and MIMO Capability with USRP Devices*. https://kb.ettus.com/Synchronization_and_MIMO_Capability_with_USRP_Devices. Accessed June 12, 2019. 2016.
- [114] RFLocDeveloper. *RFLoc Dataset*, <https://drive.google.com/file/d/1mdmxuUrR07TGljsgllucPF5fuyqD0rvb/view?usp=sharing>.
- [115] C. Rohrig and F. Kunemund. “Estimation of position and orientation of mobile systems in a wireless LAN”. In: *IEEE Decision Control*, 2007. DOI: [10.1109/CDC.2007.4434256](https://doi.org/10.1109/CDC.2007.4434256).
- [116] R. Roy and T. Kailath. “ESPRIT-estimation of signal parameters via rotational invariance techniques”. In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 37.7 (1989), pp. 984–995. ISSN: 0096-3518. DOI: [10.1109/29.32276](https://doi.org/10.1109/29.32276).
- [117] Jan R uth, Christian Bormann, and Oliver Hohlfeld. “Large-scale Scanning of TCP’s Initial Window”. In: *Proceedings of the 2017 Internet Measurement Conference*. IMC ’17. London, United Kingdom: ACM, 2017, pp. 304–310. ISBN: 978-1-4503-5118-8. DOI: [10.1145/3131365.3131370](https://doi.org/10.1145/3131365.3131370). URL: <http://doi.acm.org/10.1145/3131365.3131370>.
- [118] R.R. Sambasivan et al. “Bootstrapping evolvability for inter-domain routing with D-BGP”. In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [119] T. Sattler, B. Leibe, and L. Kobbelt. “Fast Image-Based Localization Using Direct 2D-to-3D Matching”. In: *International Conference on Computer Vision*. 2011.
- [120] Damien Saucez and Luigi Iannone. “Thoughts and Recommendations from the ACM SIGCOMM 2017 Reproducibility Workshop”. In: *SIGCOMM Comput. Commun. Rev.* 48.1 (Apr. 2018), pp. 70–74. ISSN: 0146-4833. DOI: [10.1145/3211852.3211863](https://doi.org/10.1145/3211852.3211863). URL: <http://doi.acm.org/10.1145/3211852.3211863>.
- [121] N. Scarmeas et al. “Motor Signs During the Course of Alzheimer Disease”. In: *Neurology* 63 (6 2004).
- [122] C Scherb, M Sifalakis, and C Tschudin. *CCN-lite*. 2013.
- [123] R. Schmidt. “Multiple emitter location and signal parameter estimation”. In: *IEEE TAP* 34.3 (1986), pp. 276–280. ISSN: 0018-926X. DOI: [10.1109/TAP.1986.1143830](https://doi.org/10.1109/TAP.1986.1143830).

- [124] R. Schmidt. "Multiple Emitter Location and Signal Parameter Estimation". In: *IEEE Transactions on Antennas and Propagation* 34.3 (1986).
- [125] SM Series. "Comparison of Time-Difference-of-Arrival and Angle-of-Arrival Methods of Signal Geolocation". In: (2011).
- [126] Yushi Shen and Ed Martinez. "Channel estimation in OFDM systems". In: *Freescale Semiconductor Application Note* (2006), pp. 1–15.
- [127] Shuyu Shi et al. "Accurate Location Tracking from CSI-based Passive Device-free Probabilistic Fingerprinting". In: *IEEE Transactions on Vehicular Technology* (2018).
- [128] Da-Shan Shiu et al. "Fading correlation and its effect on the capacity of multi-element antenna systems". In: *Universal Personal Communications, 1998. ICUPC '98. IEEE 1998 International Conference on*. Vol. 1. 1998, 429–433 vol.1. DOI: [10.1109/ICUPC.1998.733015](https://doi.org/10.1109/ICUPC.1998.733015).
- [129] *SIGCOMM 2017 Artifacts Survey data*. <https://team.inria.fr/diana/reproducibility/>.
- [130] P. Snyder et al. "Fifteen Minutes of Unwanted Fame: Detecting and Characterizing Doxing". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [131] J. Sommers, R. Durairajan, and P. Barford. "Automatic Metadata Generation for Active Measurement". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [132] RN Staff. "RIPE Atlas: A Global Internet Measurement Network". In: *Internet Protocol Journal* 18.3 (2015).
- [133] D. Steinmetzer et al. "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices". In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT'17*. Seoul/Icheon, South Korea, 2017.
- [134] Yaniv Taigman, Adam Polyak, and Lior Wolf. "Unsupervised Cross-Domain Image Generation". In: *CoRR* abs/1611.02200 (2016). arXiv: [1611.02200](https://arxiv.org/abs/1611.02200). URL: <http://arxiv.org/abs/1611.02200>.
- [135] Cristian Tala et al. "Guidelines for the Accurate Design of Empirical Studies in Wireless Networks". In: *Testbeds and Research Infrastructure. Development of Networks and Communities: 7th International ICST Conference, TridentCom 2011, Shanghai, China, April 17-19, 2011, Revised Selected Papers*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 208–222. ISBN: 978-3-642-29273-6. DOI: [10.1007/978-3-642-29273-6_17](https://doi.org/10.1007/978-3-642-29273-6_17). URL: http://dx.doi.org/10.1007/978-3-642-29273-6_17.
- [136] CAIDA Team. *Center for Applied Internet Data Analysis*. 2018.
- [137] MathWorks Team. *MATLAB - MathWorks*. 2018.
- [138] B.D. Van Veen and K.M Buckley. "Beamforming: a versatile approach to spatial filtering". In: *IEEE ASSP* 5.2 (1988), pp. 4–24. DOI: [10.1109/53.665](https://doi.org/10.1109/53.665). URL: <https://doi.org/10.1109%2F53.665>.

- [139] Y. Vanaubel et al. "Through the Wormhole: Tracking Invisible MPLS Tunnels". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [140] C. Villani. *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2016. ISBN: 9783662501801. URL: <https://books.google.com/books?id=5p8SDAEACAAJ>.
- [141] W. B. de Vries et al. "Broad and Load-Aware Anycast Mapping using Verfloeter". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [142] Fei Wang et al. "Residual Attention Network for Image Classification". In: *CoRR abs/1704.06904* (2017). arXiv: 1704.06904. URL: <http://arxiv.org/abs/1704.06904>.
- [143] Ju Wang et al. "LiFS: Low Human-effort, Device-free Localization with Fine-grained Subcarrier Information". In: *Proceedings of ACM MobiCom*. 2016.
- [144] X. Wang, L. Gao, and S. Mao. "CSI Phase Fingerprinting for Indoor Localization With a Deep Learning Approach". In: *IEEE Internet of Things Journal* 3.6 (2016), pp. 1113–1123. ISSN: 2327-4662. DOI: 10.1109/JIOT.2016.2558659.
- [145] X. Wang et al. "CSI-Based Fingerprinting for Indoor Localization: A Deep Learning Approach". In: *IEEE Transactions on Vehicular Technology* 66.1 (2017).
- [146] Z. Wang et al. "Your State is Not Mine: A Closer Look at Evading Stateful Internet Censorship". In: *ACM Internet Measurement Conference 2017 – IMC'17*. London, UK, 2017.
- [147] J. Wilson and N. Patwari. "See-Through Walls: Motion Tracking Using Variance-Based Radio Tomography Networks". In: *IEEE Transactions on Mobile Computing* 10.5 (2011).
- [148] J. Xiao et al. "Pilot: Passive Device-Free Indoor Localization Using Channel State Information". In: *IEEE International Conference on Distributed Computing Systems*. 2013.
- [149] Y. Xie et al. "mD-Track: Leveraging Multi-Dimensionality for Passive Indoor Wi-Fi Tracking". In: *Proceedings of ACM MobiCom*. 2019.
- [150] Jie Xiong and Kyle Jamieson. "ArrayTrack: A Fine-Grained Indoor Location System". In: *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*. Lombard, IL: USENIX, 2013, pp. 71–84. ISBN: 978-1-931971-00-3. URL: <https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/xiong>.
- [151] Zheng Yang, Zimu Zhou, and Yunhao Liu. "From RSSI to CSI: Indoor Localization via Channel Response". In: *ACM Comput. Surv.* 46.2 (2013).
- [152] A. Zaostrovnykh et al. "A Formally Verified NAT". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.
- [153] P. Zave et al. "Dynamic Service Chaining with Dysco". In: *ACM SIGCOMM 2017*. Los Angeles, USA, 2017.

- [154] C. Zhang, X. Ouyang, and P. Patras. “ZipNet-GAN: Inferring Fine-grained Mobile Traffic Patterns via a Generative Adversarial Neural Network”. In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17*. Seoul/Icheon, South Korea, 2017.
- [155] Jie Zhang et al. “CrossSense: Towards Cross-Site and Large-Scale WiFi Sensing”. In: *Proceedings of ACM MobiCom*. 2018.
- [156] Lixia Zhang et al. “Named data networking (ndn) project”. In: *Relatório Técnico NDN-0001, Xerox Palo Alto Research Center-PARC* (2010).
- [157] Q. Zhang et al. “High-Resolution Measurement of Data Center Microbursts”. In: *ACM Internet Measurement Conference 2017 – IMC’17*. London, UK, 2017.
- [158] R. Zhang and B. Preneel. “On the Necessity of a Prescribed Block Validity Consensus: Analyzing Bitcoin Unlimited Mining Protocol”. In: *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17*. Seoul/Icheon, South Korea, 2017.
- [159] Z. Zhang. “Microsoft Kinect Sensor and Its Effect”. In: *IEEE MultiMedia* 19.2 (2012).
- [160] Mingmin Zhao et al. “Through-Wall Human Pose Estimation Using Radio Signals”. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018.
- [161] Jun-Yan Zhu et al. “Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks”. In: *CoRR* abs/1703.10593 (2017). arXiv: 1703.10593. URL: <http://arxiv.org/abs/1703.10593>.
- [162] P. Zuraniewski et al. “Facilitating ICN Deployment with an Extended Open-Flow Protocol”. In: *The 4th ACM Conference on Information-Centric Networking – ICN’17*. Berlin, Germany, 2017.