Grant Agreement No: 101096307

Full Title: THz Industrial Mesh Networks in Smart Sensing and Propagation Environments
Start date: 01/01/2023
End date: 31/12/2025
Duration: 36 Months

Deliverable D4.4

Deliverable on multi-goal mesh network optimization and exploitation of smart propagation environments

Document Type	Deliverable
Title	D4.4 - Deliverable on multi-goal mesh network optimization and exploitation of smart propagation environments
Contractual due date	30/09/2025
Actual submission date	03/10/2025
Nature	Report
Dissemination Level	PUB
Lead Beneficiary	CNIT
Responsible Author	Giampaolo Cuozzo (CNIT)
Contributions from	Giacomo Bacci (CNIT), Chiara Buratti (CNIT), Sara Cavallero (CNIT), Giampaolo Cuozzo (CNIT), Davide Dardari (CNIT), Mattia Fabiani (CNIT), Wen-Xuan Long (CNIT), Francesco Miccoli (CNIT), Marco Moretti (CNIT), Andrea Pumilia (CNIT), Alessia Tarozzi (CNIT), Giulia Torcolacci (CNIT), Lutfi Samara (HWDU), Marco Skocaj (HWDU), Malte Schellmann (HWDU), Tommaso Zugno (HWDU), Steffen Kroos (TUBS)

Revision history

Version	Issue Date	Changes	Contributor(s)
v0.1	09/09/2025	Version for the WPL and Reviewers	All authors
v0.2	26/09/2025	Review comments	Luca Sanguinetti (CNIT), Thomas Kürner (TUBS)
v1.0	02/10/2025	Final Version	Flaminia Saratti (CNIT), Luca Sanguinetti (CNIT)

Disclaimer

The content of the publication herein is the sole responsibility of the publishers, and it does not necessarily represent the views expressed by the European Commission or its services. While the information contained in the documents is believed to be accurate, the authors(s) or any other participant in the TIMES consortium make no warranty of any kind with regard to this material including, but not limited to the implied warranties of merchantability and fitness for a particular purpose. Neither the TIMES Consortium nor any of its members, their officers, employees or agents shall be responsible or liable in negligence or otherwise howsoever in respect of any inaccuracy or omission herein. Without derogating from the generality of the foregoing neither the TIMES Consortium nor any of its members, their officers, employees or agents shall be liable for any direct or indirect or consequential loss or damage caused by or arising from any information, advice, inaccuracy, or omission herein.

Copyright Message

© TIMES Consortium, 2022-2025. This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation, or both. Reproduction is authorised provided the source is acknowledged.

Table of Contents

List of Abbreviations

3DQN Dueling Double Deep Q Network

6G sixth-generation

A2C Advantage Actor-Critic

ACK Acknowledgment

AI Artificial Intelligence

AM-AM Amplitude Modulation to Amplitude Modulation

AM-PM Amplitude Modulation to Phase Modulation

AoA Angle of Arrival

AWGN Additive White Gaussian Noise

BER Bit Error Rate

BO Back-Off

BS Base Station

CIR Channel Impulse Response

CRLB Cramér-Rao Lower Bound

CSI Channel State Information

CSMA Carrier Sense Multiple Access

CSMA/CA Carrier Sense Multiple Access/Collision Avoidance

CW Congestion Window

DDQN Double Deep Q Network

DNN Deep Neural Network

DQL Deep Q-Learning

DQN Deep Q Network

DRL Deep Reinforcement Learning

ELAA Electrically Large Aperture Array

EM electromagnetic

FEC Forward Error Correction

FFT Fast Fourier Transform

FIM Fisher Information Matrix

HPBW Half Power Beam Width

HP high priority

IF Intermediate Frequency

KPI Key Performance Indicator

LOS Line of Sight

LP low priority

MAC Medium Access Control

MADRL Multi-Agent Deep Reinforcement Learning

MCS Modulation and Coding Scheme

MDP Markov Decision Process

MIMO Multiple-Input-Multiple-Output

MISO Multiple-Input-Single-Output

SISO Single-Input-Single-Output

ML Maximum Likelihood

MPC Multipath Component

MSE Mean Square Error

MMSE Minimum Mean Square Error

LMMSE Linear MMSE

NET NETwork

NLOS Non Line of Sight

NMSE Normalized Mean Square Error

OFDM Orthogonal Frequency Division Multiplexing

PA Power Amplifier

PHY physical

PLC Programmable Logic Controller

PoC Proof-of-Concept

QAM Quadrature Amplitude Modulation

RCS Radar Cross Section

RF Radio Frequency

RIS Reconfigurable Intelligent Surface

RL Reinforcement Learning

RS-LS Reduced-Space Least-Squares

RX Receiver

SCS subcarrier spacing

SIMO Single Input Multiple Output

SiMoNe Simulator for Mobile Networks

SIR Signal-to-Interference Ratio

SNR Signal-to-Noise-Ratio

SINR Signal-to-Interference-plus-Noise Ratio

THz Terahertz

TX Transmitter

UE User Equipment

WP Work Package

BS base station

FF far-field

LoS line-of-sight

NF near-field

NLoS non-line-of-sight

RIS Reflective Intelligent Surface

UE User Equipment

UPA Uniform Planar Array

ULA uniform linear array

TB Table-Based

TL Table-Less

FMCW Frequency-Modulated Continuous-Wave

FFT Fast Fourier Transform

DBSCAN Density-Based Spatial Clustering of Applications with Noise

MRC Maximum Ratio Combining

1 Executive Summary

This deliverable reports on the activities carried out until September 2025 under Tasks T4.3 and T4.4 of Work Package (WP)4 within the TIMES project.

- Task T4.3, titled "THz communications in smart propagation environments", focuses on the investigation of smart radio environments to enable physics-aware algorithms and procedures for Terahertz (THz) networks;
- Task T4.4, titled "THz-tailored protocol design and optimization for intelligent multi-goal mesh networks", aims to develop Medium Access Control (MAC) and NETwork (NET) layer protocols specifically tailored to THz communications in intelligent, multi-goal mesh networks.

The activities reported in this document make use of inputs from other WPs of the project, including WP2 for the definition of simulation scenarios and target Key Performance Indicators (KPIs), WP3 for THz propagation models and Reflective Intelligent Surface (RIS) characterization, and WP5 for simulation parameters. Part of the results hereby presented will be used by WP6 for the implementation of the target Proof Of Concepts (POCs).

The document is thus structured as follows:

- Section 2 provides a brief recap of the main objectives of the TIMES project and explains how this deliverable contributes to them;
- Section 3 presents a concise overview of the inputs collected from other WPs (specifically WP2, WP3 and WP5) and the literature to define the system models adopted for the proposed solutions. These primarily include physical (PHY) layer aspects such as Radio Frequency (RF) impairments, channel and RIS models;
- Section 4 details the outcomes of Task 4.3, which include the development of enhanced algorithms and procedures for exploiting smart propagation environments in THz-based industrial scenarios. The section covers channel estimation, beam focusing, mitigation of RF impairments in the presence of RISs, as well as an analysis of theoretical performance enabled by sensing, RIS and Multiple-Input-Multiple-Output (MIMO) technologies, both in far and near fields;
- Section 5 outlines the contributions of Task 4.4, where upper-layer solutions for intelligent THz networks have been developed. This includes both model-based and data-driven protocol designs and network optimizations for multi-goal industrial mesh networks, as well as predictive algorithms for blockage detection and proactive resource allocation.

It is worth noting that the results presented in this deliverable build upon the foundations laid in Deliverable D4.2 [1], providing further advancements and insights into the ongoing development of the TIMES project.

2 Introduction

The development of next-generation wireless systems is expected to unlock the potential of THz communications for a wide range of applications, among which industrial automation stands as one of the most promising and demanding. The adoption of THz frequencies promises ultra-high data rates and ultra-low latency, essential to meet the stringent requirements of next-generation industrial environments. However, the inherent propagation challenges of THz signals, such as high path loss, sensitivity to blockages, and limited scattering, necessitate a fundamental rethinking of the wireless communication stack, from the physical layer to network control and management.

In this context, the TIMES project introduces two foundational innovations tailored to the needs of THz-enabled industrial systems: the use of smart propagation environments and the design of intelligent multi-goal mesh networks. These two pillars aim to jointly improve both the physical robustness and the network-level adaptability of THz communication systems in complex, dynamic environments.

Smart propagation environments rely on programmable surfaces, such as RISs and metasurfaces, to manipulate the wireless channel in real time, enhancing functionalities such as beamforming/beamfocusing, and improving reliability in scenarios traditionally considered hostile to high-frequency transmission. Coupled with the opportunities given by near-field propagation, such as enhanced sensing capabilities and increased spatial multiplexing gains, these technologies represent a crucial component in addressing the spatial and spectral characteristics of THz links in complex industrial settings.

At the network level, the TIMES project introduces intelligent mesh architectures designed for multi-goal operation, where communication, sensing, and heterogeneous tasks coexist and dynamically compete for network resources. These networks leverage Artificial Intelligence (AI)-based mechanisms for adaptive MAC protocols, routing, predictive maintenance, and proactive resource allocation, ensuring resilience and performance under highly variable environmental and operational conditions.

This deliverable presents the technical results achieved within this framework, with contributions spanning both the physical and upper layers and referring to both pillars. On the physical layer side, it includes novel procedures, along with theoretical characterization, for channel estimation, RF impairment mitigation, beam focusing, and sensing in the presence of smart surfaces. On the network layer side, it introduces data-driven and model-based designs for THz mesh networks, as well as predictive and AI-assisted solutions for MAC protocols, routing, handover and network management.

Altogether, the findings presented in this deliverable contribute to the realization of a highly adaptive, high-capacity, and context-aware wireless infrastructure for future industrial applications operating at THz frequencies. By combining advanced physical-layer techniques with intelligent network-level coordination, the innovation of this deliverable sets a solid foundation for the deployment of THz technologies in realistic, performance-critical industrial environments. Moreover, the outcomes presented in this deliverable will provide guidelines for future activities of the project, including the realization of POCs demonstrating the potential of developed solutions.

3 Reference models

In this section, we present the reference models that have been adopted for the modeling of RF impairments, propagation channels, and RISs.

3.1 RF impairments

In Deliverable 4.2 [1], we mainly investigated the effect of phase shifting errors at the RIS on the performance of a far-field communication link. The performance evaluation was carried out by computing the average Signal-to-Noise-Ratio (SNR) and a bound on the spectral efficiency, where the computed spectral efficiency has led to recommendations of the adoption of Quadrature Amplitude Modulation (QAM) modulation orders as a function of the severity of the RIS impairments. The impairment that was considered was a phase shifting error caused by the imperfection of the phase shifting operation at each of the RIS elements.

In this deliverable, we jointly consider two types of impairments. The first is related to the RIS, manifested through phase-shifting errors and thus loss in beamforming gain. These impairments were introduced in Deliverable 4.2 [1] and are based on the formulation in [2]. The second source of impairment is one that is prevalent in THz systems, which is transmission Power Amplifier (PA) nonlinearities that causes waveform distortions, and thus increased error floors at the receiver. The PA nonlinearity model is based on the analysis and parameter derivations carried out in Deliverable 4.3 [3]. It is worth mentioning that the previous deliverables (in particular Deliverables 4.1 [4] and 4.3 [3]) consider the measurements and modeling of several other limiting RF impairments such as oscillator phase noise, in-phase and quadrature imbalance and local oscillator chain harmonics. In this deliverable, we focus on RF impairments that are most relevant to achieving an improved communication coverage, which is the main motivation for deploying RISs. Toward this goal, the two main identified limitations are: i) limitations on Tx power due to PA nonlinearities and ii) reduction in RIS beamforming accuracy arising from phase shifting errors at the RIS elements.

3.1.1 RF impairments at the RIS

Herein, we start by presenting the phase error model at each reflecting element in the RIS. Let $h_{t,n} = e^{j\phi_{t,n}}$, and $h_{r,n} = e^{j\phi_{n,r}}$. Therefore, each phase at the RIS element, which we denote by θ_n, should be adjusted such that such that

$$\theta_n = -(\phi_{n,t} + \phi_{n,r}). \tag{1}$$

However, RIS elements can manifest imperfections in achieving a desired phase shift, and thus, we get

$$\theta_n + \phi_{n,t} + \phi_{n,r} = \epsilon_n, \tag{2}$$

where ϵ_n is the phase error.

The phase errors at the RIS are primarily attributed to two factors: intrinsic effects and limitations in the RIS design. Below, we provide a detailed explanation of the adopted RF impairment models at the RIS.

RIS phase shifting errors To model the phase error, a statistical approach using a generalized distribution is adopted. The Von-Mises distribution is adopted to model $\epsilon_n \ \forall n \in \{0, \dots, Q-1\}$, where it is modeled as an i.i.d zero-mean Von-Mises random variable with concentration parameter κ and a probability density function that may be written as [5]

$$f_\epsilon(y) = \frac{e^{\kappa \cos(y)}}{2\pi I_0(\kappa)}, \quad -\pi < y < \pi \tag{3}$$

where $I_0(.)$ is the zero-th order modified Bessel function of the first kind defined in [6, 9.6.19]. To demonstrate the effect of the concentration parameter κ, Figure 1 shows the effect of varying κ on the probability density function (PDF) defined in (3). It is shown that as κ is increased, the distribution of the Von Mises random variable becomes more concentrated around zero.

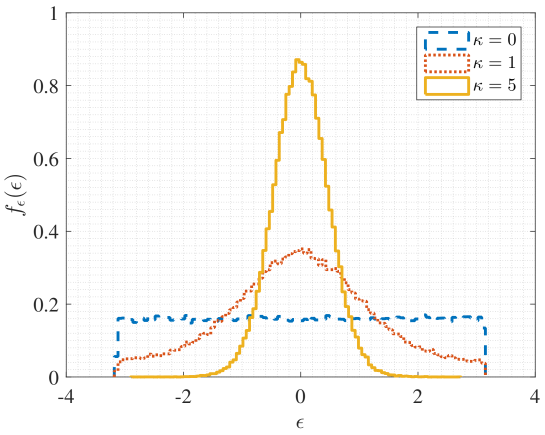

Figure 1: Von Mises PDFs for different values of the concentration parameter κ.

RIS phase quantization effects When RISs are equipped with a large number of reflective elements, and when operating at higher frequencies, controlling the phase shift of each element becomes a complex task. Thus, in practical RIS setups, each RIS element is equipped with the capability to vary the phase for only a limited number of values, i.e. quantized phase shifts. It is then crucial to incorporate the impact of quantization on the performance of RIS assisted links.

To model the impact of quantization errors in RIS assisted links, we resort to the method proposed in [2], where the distribution of the error resulting from the quantization procedure at each RIS element follow a uniform distribution with varying limits, i.e.

$$f_\epsilon(y) = \mathcal{U}\left[-\frac{\pi}{L}, \frac{\pi}{L}\right], \tag{4}$$

where L is the number of quantization levels. Typical values used in the literature for L are 2, 4 and 8 [2].

3.1.2 Transmitter PA nonlinearity

High antenna gain and high transmission powers are needed to meet link budget constraints when communicating at the THz band. If high transmission power is relied upon to meet the link budget, it is important to consider the nonlinearities generated as a result of adopting high transmission powers due to the nonlinearities induced by the PA.

PA nonlinearity causes saturation in the transmission power and unwanted phase shifts that has a detrimental effect on communication performance. Many different approaches exist to model the PA nonlinearities. We have surveyed the different methods to model PA nonlinearities in [4], and proposed new model parameters in [3] in the quasi-memoryless PA nonlinearity modeling framework. Amplitude Modulation to Amplitude Modulation (AM-AM) and Amplitude Modulation to Phase Modulation (AM-PM) models can be derived by using any modeling approach that is available in the literature. In this deliverable, we use the models extracted from the results reported in Deliverable 4.3 [3], where we have derived different model parameters related to a PA operating at 300 GHz. The derived models were a polynomial model, a modified Rapp model, a Ghorbani and a Saleh model. The models vary in terms of accurately predicting the performance of the PA and the simplicity of the formula. Let us, for completeness, report the equations related to the modified Rapp model, that may be realized by applying

$$F_{\text{AMAM}}(v_{in}) = \frac{Gx}{(1 + |\frac{Gv_{in}}{V_{sat}}|^{2p})^{(\frac{1}{2p})}}, \tag{5}$$

$$F_{\text{AMPM}}(v_{in}) = \frac{(Av_{in}^{q_1})}{(1 + |\frac{v_{in}}{B}|^{q_2})}, \tag{6}$$

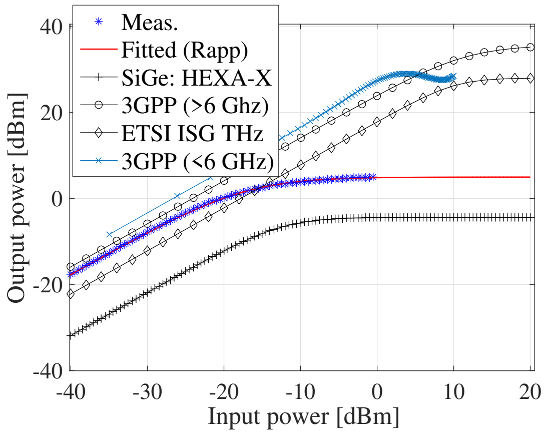

Figure 2: Comparisons between different AM-AM models.

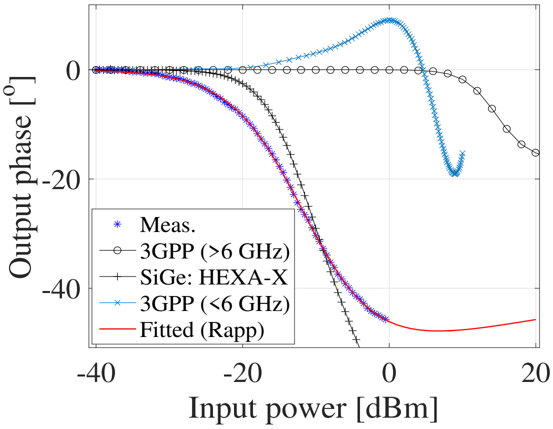

Figure 3: Comparisons between different AM-PM models.

where v_{in} is the input voltage to the PA, G represents the small signal gain (see the related definition in [7]), p denotes the smoothness factor, V_{Sat} is the saturation voltage coefficient and the coefficients A, B, q are AM-PM distortion curve parameters. This model is referred to in the literature as the modified-Rapp model [7], which is a widely adopted quasi-memory-less model [8] that is intended to model the AM-AM and AM-PM distortion relations. At such higher frequencies, the linear mode of operation of a power amplifier is limited. To give the reader an idea on the limitations imposed on operating a PA at such high frequencies, we next compare the derived model with models available in the literature as seen in Figures 2 and 3. The models that are selected to be compared with are:

- 3GPP (< 6 GHz): This is a polynomial-fitted model for both AM-AM and AM-PM, which was proposed in [9] for the purpose of evaluating waveforms operating below 6 GHz.
- 3GPP (> 6 GHz): This model is a modified Rapp-fitted model, which was proposed in [9] to evaluate waveforms operating above 6 GHz, namely 30 and 70 GHz.
- SiGe-HEXA-X: SiGe-based modified Rapp model, which was proposed in [10, 11] for an operating frequency of 290 GHz.
- ETSI ISG THz-based model, which is a modified Rapp model proposed in [12] for operating frequencies in the 100-200 GHz band, extracted by averaging out the modified Rapp model parameters in the

specified frequency range solely targeting AM-AM behavior. Parameters for AM-PM behavior were left unaddressed.

It is worth noting that the 3GPP (< 6 GHz) model is strictly evaluated in the range of input powers $[-35, 9]$ dBm, which is the range in which the polynomial was fitted, otherwise, the behavior of the model is unreliable. This hints to why the modified Rapp model is preferred over the fitted polynomial model, as the ranges of input power that are fed to the modified Rapp models exceed the range that it was fitted within while still producing an expectable behavioral performance. One notable conclusion of this comparison is the over-estimation of the linear region of the PA models proposed by 3GPP (both higher and lower than 6 GHz) and ETSI ISG THz-based models when compared to our derived model. For instance, at an input power value of -10 dBm, the 3GPP (both higher and lower than 6 GHz) and ETSI ISG THz-based models still operate in the linear region for the AM-AM case and introduce low phase shift values, while the SiGe-based MMIC and the modeled amplifier produce around -30^o phase shift, and are clearly in the compression region, which is detrimental for even low-order QAM-based modulation.

3.2 Channel characterization

3.2.1 Far-field channel model

In far-field conditions, the channel model considered in this deliverable has been taken from Deliverable D3.1 (Sec. 5.2.1) [13]. The latter has been derived experimentally from channel measurements and is described using a single-band model, where the path loss PL in dB can be formulated as follows:

$$\mathrm{PL[dB]} = 10 \cdot \mathrm{PLE} \cdot \log_{10}\left(\frac{d}{d_0}\right) + \mathrm{FSPL}(d_0) + X_\sigma \tag{7}$$

where d is the Tx-Rx Euclidean distance, and $d_0 = 1\,\mathrm{m}$ in this work, represents the reference distance. PLE is the path loss exponent, and X is a zero-mean Gaussian random variable with standard deviation σ. Free-Space Path-Loss (FSPL) is given by the Friis' law as:

$$\mathrm{FSPL\ [dB]} = 20 \cdot \log_{10}(d) + 32.4 + 20 \cdot \log_{10}(f_c) \tag{8}$$

with f_c as the carrier frequency.

The path loss parameters were estimated for the Line of Sight (LOS) and Non Line of Sight (NLOS) conditions, and we took from Table 4 in Deliverable D3.1 [13], the corresponding values for communication at THz:

Table 1: Single-band PLE and standard deviation σ (in dB)

Scenario	**Band**	**PLE**	σ
LOS	THz	2.26	1.51
NLOS	THz	3.27	7.32

The power received from the useful receiver, P_{RX} can be then expressed as:

$$P_{\mathrm{RX}}\ [dBW] = P_{\mathrm{TX}}\ [dBW] + \eta_{\mathrm{TX}}\ [dB] + \eta_{\mathrm{RX}}\ [dB] + G_{\mathrm{TX}}\ [dB] + G_{\mathrm{RX}}\ [dB] - PL\ [dB], \tag{9}$$

where P_{TX} is the transmitted power, η_{TX}, η_{RX} are the transmitter and receiver antenna efficiencies, respectively, G_{TX}, G_{RX} are the transmitter and receiver gains.

3.2.2 Near-field channel model

To illustrate the impact of near-field propagation on channel modeling, the Uniform Planar Array (UPA) model taken from Deliverable D4.2 (Sec. 2.1) [1] and depicted in Figure 4 is considered. The array is placed in the $yz-$plane of a three-dimensional space, where a spherical coordinate system is defined, with φ being the azimuth angle, θ the elevation angle and r the distance. The array has N_H elements in each row and N_V elements in each column, resulting in a total of $N = N_H \times N_V$ elements. When the UPA elements are indexed row-by-row by $n \in [1, N]$, the location of the n-th array element is given by

$$\mathbf{v}_n = [0, y_n, z_n]^{\mathrm{T}} = \left[0, \mathrm{mod}(n-1, N_H) d_N, \left\lfloor \frac{n-1}{N_H} \right\rfloor d_N \right]^{\mathrm{T}}, \tag{10}$$

where $\mathrm{mod}(\cdot)$ denotes the modulus operation, $\lfloor \cdot \rfloor$ truncates the argument, and d_N is the horizontal and vertical inter-element spacing of the UPA.

The Fraunhofer distance of the array can be computed as [14]

$$R = \frac{2D^2}{\lambda}, \tag{11}$$

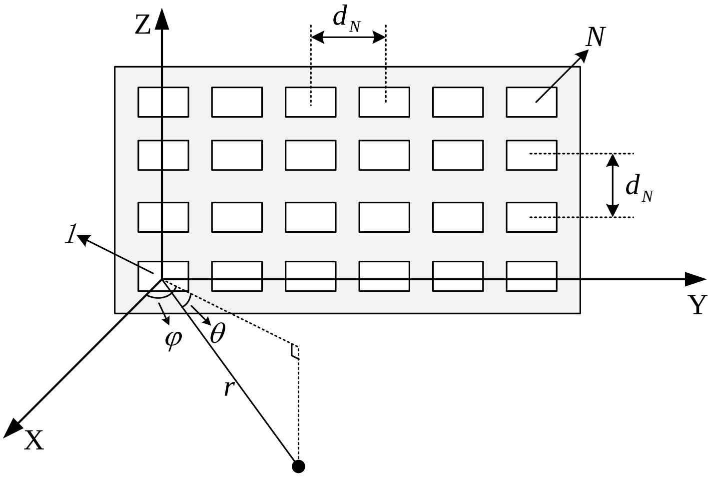

Figure 4: The array-based system model.

where $D = \sqrt{(N_H^2 + N_V^2)}d_N$ is the array aperture length and λ is the wavelength. This distance is used to distinguish between the far-field and radiative near-field regions of the array. When the distance between the array and receiver is smaller than R, the receivers would be located within the relative near-field region of the array. Moreover, for a Extra-Large RIS (XL-RIS)-aided system, when the distance r_1 between the XL-RIS and the BS and the distance r_2 between the XL-RIS and the User Equipment (UE) satisfy [14]

$$\frac{r_1 r_2}{r_1 + r_2} \leq R = \frac{2D^2}{\lambda}, \tag{12}$$

the RIS-aided cascaded channel should be modeled as the near-field channel. As the increase in the array size and the decrease in the wavelength, the Fraunhofer distance of the XL-RIS-aided system in the THz band is likely to reach hundreds of meters. As an example, consider an array with a size of 0.2×0.1 m^2 and an inter-element spacing $d_N = 2\lambda$, operating at $0.3\,\mathrm{THz}$ (i.e., $\lambda = 0.001\,\mathrm{m}$). This array consists of 5000 elements, arranged in a 100×50 configuration, yielding a Fraunhofer distance of $R = 100\,\mathrm{m}$. This implies that in the near-field region, where $\{r_1, r_2\} \leq R$, the far-field approximation cannot be used, and the exact propagation model for both the UE-RIS channel and RIS-base station (BS) must be considered instead. This model is introduced next.

Based on the power delay profile measurement results in [15], a reasonable assumption in the THz band is that the signal transmitted by each UE arrives at the array within a small solid angle centered around the LoS path. We focus on a specific UE and define the LoS angles from the UE to the array as (φ_h, θ_h). Assuming a conventional correlated Rayleigh fading channel, the channel vector $\mathbf{h} \in \mathbb{C}^N$ can be modeled as

$$\mathbf{h} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}_N, \mathbf{R}_h) \tag{13}$$

which is fully characterized by the spatial correlation matrix

$$\mathbf{R}_h = \beta_h \int_{r_h - \Delta_{r_h}}^{r_h + \Delta_{r_h}} \int_{\varphi_h - \Delta_{\varphi_h}}^{\varphi_h + \Delta_{\varphi_h}} \int_{\theta_h - \Delta_{\theta_h}}^{\theta_h + \Delta_{\theta_h}} f_h(\tilde{r}, \tilde{\varphi}, \tilde{\theta}) \mathbf{a}(\tilde{r}, \tilde{\varphi}, \tilde{\theta}) \mathbf{a}^{\mathrm{H}}(\tilde{r}, \tilde{\varphi}, \tilde{\theta}) d\tilde{\theta} d\tilde{\varphi} d\tilde{r}, \tag{14}$$

where $\beta_h = \frac{1}{N}\mathrm{tr}\{\mathbf{R}\}$ is the average channel power (capturing pathloss and shadowing), $(r_h, \varphi_h, \theta_h)$ are the distance, the azimuth and elevation angles of the considered UE, while the triplet $(\Delta_{r_h}, \Delta_{\varphi_h}, \Delta_{\theta_h})$ accounts for the corresponding distance and angular spreads. Also, $\mathbf{a}(\tilde{r}, \tilde{\varphi}, \tilde{\theta})$ is the array response vector [16]:

$$\mathbf{a}(\tilde{r}, \tilde{\varphi}, \tilde{\theta}) = \left[1, \cdots, e^{j\frac{2\pi}{\lambda}(\tilde{r}_n - \tilde{r})}, \ldots, e^{j\frac{2\pi}{\lambda}(\tilde{r}_N - \tilde{r})}\right]^{\mathrm{T}} \tag{15}$$

where $\tilde{r}_n$ and $\tilde{r}$ are the distances from the n-th array element and the reference element to a point within the spatially spread region. The distance $\tilde{r}_n$ can be computed as [16]

$$\tilde{r}_n = \tilde{r}\sqrt{1 - \frac{2\mathbf{k}^{\mathrm{T}}(\tilde{\varphi}, \tilde{\theta})\mathbf{v}_n}{\tilde{r}} + \frac{\|\mathbf{v}_n\|^2}{\tilde{r}^2}} = \tilde{r}\sqrt{1 - \frac{2(y_n \cos\tilde{\theta} \sin\tilde{\varphi} + z_n \sin\tilde{\theta})}{\tilde{r}} + \frac{y_n^2 + z_n^2}{\tilde{r}^2}}, \tag{16}$$

where $\mathbf{k}(\tilde{\varphi}, \tilde{\theta}) = [\cos\tilde{\theta}\cos\tilde{\varphi}, \cos\tilde{\theta}\sin\tilde{\varphi}, \sin\tilde{\theta}]^{\mathsf{T}}$ is the radiation direction from the point to the array. Finally, $f(\cdot)$ is the normalized spatial scattering function [17]. Based on (16), the near-field array response vector in (15) is influenced not only by $\tilde{\varphi}$ and $\tilde{\theta}$, but also by the distance $\tilde{r}$.

3.2.3 Near-field electromagnetic interference model

In most existing RIS-aided communication systems, such as the ones investigated in Deliverable D4.2 [1], attention is primarily given to the signals generated by the system, while the impact of Electro-Magnetic Interference (EMI) or ‘noise’ (sometimes called ‘pollution’) present in the environment is often overlooked. However, as the size of the RIS increases, it becomes more vulnerable to EMI. This interference can arise from a variety of sources, including man-made devices, natural background radiation, and both intentional and unintentional factors.

We define $\mathbf{e}(i) \in \mathbb{C}^N$ as the vector collecting the EMI during the i-th channel use, capturing any uncontrollable factors (e.g., electromagnetic or hardware-related disturbances) affecting the incoming signals at the XL-RIS. The EMI is modeled as $\mathbf{e}(i) \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}_N, \sigma_e^2\mathbf{R}_e)$, and its realizations are assumed to be independent across channel uses, i.e., $\mathbb{E}\{\mathbf{e}(i)\mathbf{e}(i')^{\mathsf{H}}\} = 0$ for $i \neq i'$. The normalized spatial correlation matrix $\mathbf{R}_e \in \mathbb{C}^{N \times N}$ follows the same form as in (14), i.e.,

$$\mathbf{R}_e = \int_{r_e-\triangle_{r_e}}^{r_e+\Delta_{r_e}} \int_{\varphi_e-\triangle_{\varphi_e}}^{\varphi_e+\Delta_{\varphi_e}} \int_{\theta_e-\Delta_{\theta_e}}^{\theta_e+\Delta_{\theta_e}} f_e(\tilde{r}, \tilde{\varphi}, \tilde{\theta})\mathbf{a}(\tilde{r}, \tilde{\varphi}, \tilde{\theta})\mathbf{a}^{\mathsf{H}}(\tilde{r}, \tilde{\varphi}, \tilde{\theta}) d\tilde{\theta} d\tilde{\varphi} d\tilde{r}, \tag{17}$$

but with a different spatial scattering function $f_e(\tilde{r}, \tilde{\varphi}, \tilde{\theta})$.

3.2.4 Raytracing-based channel modeling

In this section, we present two ray tracing tools employed for the deterministic modeling of wireless channels. Unlike the stochastic models discussed previously, ray tracing enables precise modeling of signal propagation in specific environments. This enables us to carry out simulations that account for the unique characteristics of the industrial scenarios targeted by this project. As detailed in Deliverable 2.2 [18], we evaluated several ray tracing tools and identified SiMoNe and Sionna as the most appropriate for the project's requirements. In the following subsections, we provide an overview of these tools and outline the rationale behind this choice.

SiMoNe The Simulator for Mobile Networks (SiMoNe) ray tracer models the propagation of electromagnetic waves in complex environments by tracing rays between Transmitter (TX) and Receiver (RX) through a geometric representation of the scene. The algorithm supports four types of deterministic wave interactions with the environment: reflection, scattering, transmission, and diffraction. The spatial resolution of the scene is defined by polygonal meshes representing walls, objects, and other obstacles. Based on the physical interaction of rays with these surfaces, the path gain of individual rays is computed. It is important to note that, as with conventional ray-tracing techniques, the ray-based propagation model in SiMoNe is only valid under far-field conditions. Near-field effects such as reactive coupling or evanescent waves are not considered within the ray tracing framework. However, first measurements involving RIS at 300 GHz carried out in [19] have shown, that far-field solutions may be valid also for distances significantly smaller than the Fraunhofer distance.

Reflection is modeled based on the angle of incidence and the local surface normal. The strength of the reflected component depends on the material properties of the surface and the incident angle, following standard Fresnel reflection models. Specular reflection is assumed, and only single-bounce or multi-bounce paths that maintain geometric consistency are traced.

Scattering is treated as a diffuse process, typically occurring at surfaces with rough structures or non-ideal geometries. In SiMoNe, this is modeled as an angular dispersion of energy away from the specular direction, weighted according to empirical or material-specific scattering coefficients.

Refraction describes the partial penetration of electromagnetic energy through obstacles. In SiMoNe, this is modeled using material-specific transmission losses that attenuate the signal power as it passes through a surface. The loss values are derived from simplified models and are typically assigned based on empirical data or assumed material parameters. This modeling approach enables an efficient yet realistic approximation of signal attenuation through semi-transparent materials such as glass or plastic.

Diffraction is modeled to account for wave bending around sharp edges or corners. In SiMoNe, diffraction is implemented based on geometric theory, allowing rays to reach regions not directly accessible through line-of-sight or simple reflections. This mechanism is particularly relevant in environments with frequent occlusions or narrow passageways.

The direct line-of-sight (LoS) component or rays that propagate without interaction are modeled using the free-space pathloss model, based on the Friis transmission equation:

$$L_{\mathrm{FS}} = \left(\frac{4\pi df}{c}\right)^2 \tag{18}$$

where L_{FS} is the free-space pathloss, d is the distance between transmitter and receiver, f is the carrier frequency, and c is the speed of light. The received power is then computed by multiplying the transmitted power with the gains of the transmit and receive antennas and the inverse of the pathloss.

To efficiently identify valid ray paths, SiMoNe employs a spatial subdivision algorithm as described in [20], which accelerates the intersection tests between rays and scene geometry. This enables scalable simulation of large and complex industrial environments.

In order to enable RIS-assisted communication modeling, the ray tracing process in SiMoNe has been extended to support two-segment propagation paths. Instead of treating the RIS as a conventional reflecting surface, the interaction is split into two independent ray tracing steps: one from the TX to the RIS, and another from the RIS to the RX. This approach avoids the need to model non-specular reflection directly within the ray tracer and instead allows for postprocessing-based recombination of rays according to a programmable RIS model [21].

To achieve this, the RIS is treated as both a virtual receiver (during the TX–RIS phase) and a virtual transmitter (during the RIS–RX phase). After the individual ray segments are computed, they are combined based on geometric consistency and incident/departure angles. The resulting two-hop paths are then evaluated using an external gain model, which maps angle-of-arrival (AoA) and angle-of-departure (AoD) pairs to corresponding reflection coefficients or gain values. This modular approach allows flexible modeling of different RIS behaviors, such as ideal beam steering or realistic angular-dependent scattering.

As a result of the RIS interaction, the set of possible propagation paths is obtained as the product of all rays impinging on the RIS and all rays departing from it. This combinatorial construction can lead to a large number of potential RIS-assisted paths. To maintain computational efficiency and physical relevance, postprocessing filters are applied to discard paths with excessive pathloss or geometrical inconsistencies.

In SiMoNe, rays interacting with a RIS are processed by separately computing the propagation paths from the TX to the RIS and from the RIS to the RX. Unlike conventional reflections, the RIS is modeled as a passive programmable surface, whose contribution to the overall path gain is computed using a compensation factor derived from the radar equation, as proposed in [22].

The total free-space path gain of the two-hop link without considering the RIS properties is given by:

$$G_{\text{FSPG, dB}} = 40\log_{10}(\lambda) - 40\log_{10}(4\pi) - 20\log_{10}(d_1) - 20\log_{10}(d_2), \tag{19}$$

where d_1 and d_2 are the distances from the TX to the RIS and from the RIS to the RX, respectively, and λ is the carrier wavelength.

To include the effect of the RIS, the radar equation is applied:

$$G_{\text{Radar, dB}} = 20\log_{10}(\lambda) + 10\log_{10}(\sigma) - 30\log_{10}(4\pi) - 20\log_{10}(d_1) - 20\log_{10}(d_2), \tag{20}$$

where σ denotes the effective Radar Cross Section (RCS) of the RIS in m^2. The difference between both expressions leads to a correction term that accounts for the influence of the RIS:

$$\alpha_{\text{dB}} = G_{\text{Radar, dB}} - G_{\text{FSPG, dB}} = -20\log_{10}(\lambda) + 10\log_{10}(4\pi) + 10\log_{10}(\sigma). \tag{21}$$

This correction factor α_{dB} is added to the path gain of all rays that are reflected via the RIS. The resulting power contribution is then coherently added to all other rays at the receiver, taking both amplitude and phase into account. This methodology enables the accurate inclusion of the RIS-induced reflection gain into the overall received signal power.

Sionna Sionna is an open source python library developed by Nvidia for the simulation of wireless systems [23]. It includes a differentiable ray tracing tool for radio propagation modeling, which makes use of the just-in-time compiler Dr.Jit and the Mitsuba renderer. The ray tracing algorithm combines the Shooting and Bouncing of Rays (SBR) method to generate candidate propagation paths and the image method to determine which of the candidates are valid. The valid paths are then taken into account to obtain the channel impulse response. To model the interaction of EM waves with objects in the scene, multiple mechanisms are supported, including specular reflections, diffuse reflections, and refraction. The dielectric behavior of materials is taken into account by adopting the ITU-R recommendation P.2040-3, with the possibility of specifying custom properties for unsupported materials. Similarly to SiMoNe and other conventional ray tracing tool, the ray-based propagation model implemented in Sionna is valid in far-field conditions. In the case of large arrays or RIS, near-field effects such as spherical wave propagation can be taken into account by modeling the overall channel as the superposition of links between each transmit and receive antenna elements. The tool is optimized to run on GPUs, enabling a fast computation of propagation paths even in presence of multiple transmitters and receivers. While not primarily intended for channel modeling at THz frequencies, this tool has been validated through calibration with real THz channel measurements in industrial environments [24, 25].

Although SiMoNe offers extensive support for ray tracing simulations at THz frequencies, in the work presented in Sec. 5.4 we made use of Sionna. This choice was motivated by the fact that Sionna is open source, written in Python, and fully-differentiable, thus providing seamless compatibility with popular deep learning frameworks and enabling an easier integration into our workflow.

3.3 RIS modeling

The model of a THz-band XL-RIS-aided system from a given UE with a single antenna to a BS, as shown in Figure 5, is taken from Deliverable D4.2 (Sec. 2.1) [1]. The XL-RIS is composed of N passive reconfigurable elements, organized in a UPA with N_V rows and N_H columns, such as the one introduced in Sect. 3.2, where the horizontal and vertical inter-element spacing is d_N. Similarly, the BS is equipped with M antennas, arranged in a UPA with M_V rows and M_H columns, with inter-element spacing d_M in both the horizontal and vertical directions. As can be seen in Figure 5, the architecture investigated in this deliverable differs from the one used in [1] due to the presence of EMI $\mathbf{e}(i)$.

In addition, in the RIS-aided system, we assume a block-fading model where each channel takes one realization in a coherence block of τ_c channel uses and independent realizations across blocks, as presented in D3.1 (Sec. 5.1.1) [13]. The channel from the UE to the RIS is called $\mathbf{h}$ and modelled as the correlated Rayleigh fading channel $\mathbf{h} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}_N, \mathbf{R}_h)$, as shown in (13). The narrowband channel from the n-th RIS element to the m-th BS element is indicated as $g_{m,n}$. We call $\mathbf{g}_m = [g_{m,1}, g_{m,2}, \dots, g_{m,N}]^{\mathsf{T}} \in \mathbb{C}^N$ the channel vector from the RIS to the m-th BS element, while $\mathbf{g}'_n = [g_{1,n}, g_{2,n}, \cdots, g_{M,n}]^{\mathsf{T}} \in \mathbb{C}^M$ denotes the channel vector from the n-th RIS element to the BS array. Using the Kronecker model, we have that

$$\mathbf{g}_m \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}_N, [\mathbf{R}_{g'_n}]_{m,m}\mathbf{R}_{g_m}) \tag{22}$$

$$\mathbf{g}'_n \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}_M, [\mathbf{R}_{g_m}]_{n,n}\mathbf{R}_{g'_n}) \tag{23}$$

where $[\mathbf{R}_{g'_n}]_{m,m}\mathbf{R}_{g_m}$ and $[\mathbf{R}_{g_m}]_{n,n}\mathbf{R}_{g'_n}$ are the spatial correlation matrices of $\mathbf{g}_m$ and $\mathbf{g}'_n$, respectively. Both $\mathbf{R}_{g_m}$ and $\mathbf{R}_{g'_n}$ have the same form as (14) but must be computed using the specific spatial scattering function and channel gain corresponding to $\mathbf{g}_m$ and $\mathbf{g}'_n$. We assume the channels $\mathbf{g}_m$ (and thus $\mathbf{g}'_n$) and $\mathbf{h}$ are independent of each other. Thus, the cascaded channel between the UE and the m-th element of the BS is given by the N-dimensional vector

$$\mathbf{x}_m = \mathbf{g}_m \odot \mathbf{h}. \tag{24}$$

3.3.1 Scalable Diagonal Scattering Matrix Model for Large RIS

The model illustrated in this section is derived from [26] and [27]. In particular, the RIS is constituted by $N = N_\mathrm{x} \times N_\mathrm{y}$ elements and is assumed to be passive. Moreover, it follows a single-connected architecture, where each RIS element is connected to a single tunable complex impedance. Under this configuration, the anomalous reflection behavior of the RIS can be mathematically described by a diagonal scattering matrix $\mathbf{\Psi} \in \mathbb{C}^{N \times N}$, defined:

$$\mathbf{\Psi} = \mathrm{diag}(e^{-\mathrm{j}\psi_1}, e^{-\mathrm{j}\psi_2}, \dots, e^{-\mathrm{j}\psi_N}), \tag{25}$$

where a generic ψ_n represents the phase shift introduced by the n-th element of the RIS. These phase shifts are designed such that the reflected signals from all elements combine coherently in the desired direction, thus enhancing the beamforming performance.

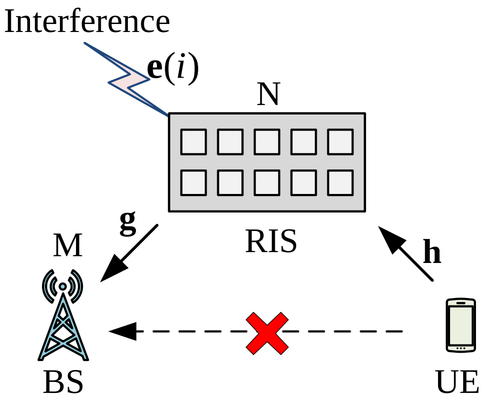

Figure 5: The XL-RIS-aided communication system.

Following the model exploited in [26], and assuming the RIS elements distributed along the x-y plane, the phase shift of an element located at position $(n_{\mathrm{x}}, n_{\mathrm{y}})$ can be computed as:

$$\psi_{n_{\mathrm{x}},n_{\mathrm{y}}}(\Theta_{\mathrm{inc}}, \Theta_{\mathrm{ref}}) = -\frac{2\pi n_{\mathrm{x}} d_{\mathrm{x}}}{\lambda} + (u_x(\Theta_{\mathrm{inc}}) + u_x(\Theta_{\mathrm{ref}})) - \frac{2\pi n_{\mathrm{y}} d_{\mathrm{y}}}{\lambda}(u_y(\Theta_{\mathrm{inc}}) + u_y(\Theta_{\mathrm{ref}})), \tag{26}$$

where $\Theta_{\mathrm{inc}} = (\theta_{\mathrm{inc}}, \phi_{\mathrm{inc}})$ and $\Theta_{\mathrm{ref}} = (\theta_{\mathrm{ref}}, \phi_{\mathrm{ref}})^1$ denote the angles of incidence and reflection, respectively. The parameters d_x and d_y indicate the inter-element spacing along the x and y directions, λ the wavelength, and for convenience, the parameters $u(\Theta)$ are defined as:

- $u_x(\Theta) = sin(\theta)cos(\phi)$;
- $u_y(\Theta) = sin(\theta)sin(\phi)$.

Based on this phase-shift model, which depends on the angles of incidence and reflection, it is possible to precompute a codebook containing vectors of phase shifts. The RIS can then be configured in real time by selecting the appropriate phase shift vector from the codebook, corresponding to the desired steering directions.

However, the number of RIS elements, N, is typically in the order of thousands to mitigate the severe path loss typically associated with THz frequency propagation. This requirement imposes a significant computational burden due to the complexity of simulating RISs with thousands of elements. To address this challenge, the RIS model described above can be extended using the approximation proposed in [27], and this allows to balance modeling accuracy and computational efficiency. Specifically, the RIS surface is partitioned into tiles. Each tile aggregates a group of multiple adjacent elements, and its overall reflection characteristics, both amplitude and phase, are modeled based on the number of constituent elements. This physics-based approximation captures the collective behavior of the elements, enabling scalable simulations of the real-time RIS-assisted THz communication systems, as discussed in Section 5.3, while maintaining a reasonable trade-off between simulation model fidelity and complexity.

---

[1] Assuming the conventional spherical coordinate system, the azimuth and elevation angles are defined as: $\phi \in [0, 2\pi)$ and $\theta \in [0, \pi)$, respectively.

4 Enhanced procedures for exploitation of smart propagation environments in industrial ecosystems

In this section, we present enhanced algorithms and procedures for exploiting smart propagation environments in THz-based industrial scenarios, including channel estimation, beam focusing, mitigation of RF impairments in the presence of RISs, as well as an analysis of theoretical performance enabled by sensing, RIS and MIMO technologies, both in far and near fields.

4.1 Estimation schemes for RIS-based THz channels

As the RIS size increases and the wavelength decreases, the Fraunhofer distance [28], which serves as the boundary between the radiative near-field (NF) and far-field (FF) regions, significantly expands in RIS-aided THz communication systems. As a result, both the UE and the BS are highly likely to reside within the radiative NF region of the RIS. In this region, the electromagnetic wavefront exhibits a spherical curvature, leading to spherical phase variations across the elements of RIS. These variations must be characterized jointly by the angle and distance between the RIS and the point source, fundamentally altering the channel propagation characteristics. Therefore, it is necessary to re-evaluate the applicability of conventional FF channel estimation methods in NF scenarios.

More critically, as the physical dimensions of the RIS grow, the system becomes more vulnerable to EMI. Such interference may originate from natural phenomena, such as atmospheric noise [29], or be caused by intentional or incidental human-generated sources, including signals directed to other users [30] or radiation from power lines [29]. However, in existing studies on RIS-aided NF channel estimation [31–33], the impact of such inevitable EMI is often overlooked.

To address the above challenges in RIS-aided THz wireless communications, this part extends the Linear MMSE (LMMSE) estimator and the Alternating Optimization (AO) strategy originally proposed in [34] to a scenario where the UE, BS, and an interference source are all located within the NF region of the RIS. In this setting, the UE signal reaches the RIS via a NF LoS path with a narrow solid angular spread, and is reflected to the BS through another similarly NF LoS path, while EMI at the RIS is also reflected toward the BS. The UE-RIS and RIS-BS channels, along with the EMI, are jointly characterized by the distances and Direction of Arrivals (DoAs) from the UE, BS, and interference source to the RIS. Distinct from the FF assumptions in prior work [34], we demonstrate that the aforementioned algorithms can only operate effectively under NF conditions when accurate NF channel and EMI statistics are incorporated. Based on these NF statistical properties, we estimate the uplink cascaded channel using the LMMSE estimator from [34], and optimize the RIS phase-shifts via the corresponding AO-based iterative algorithm. Simulation results in the THz band confirm the necessity of employing NF statistical knowledge in this setting and show that the proposed NF-aware LMMSE approach significantly outperforms existing benchmark schemes in terms of Normalized Mean Square Error (NMSE).

4.1.1 Pilot transmission

We consider a system operating under a communication protocol that allocates τ channel uses for uplink channel estimation before data transmission begins. The training sequence consists entirely of ones. Let $\phi(i) \in \mathbb{C}^N$ denote the vector containing the adjustable phase-shifts $\{\phi_n(i) \in [0, 2\pi); n = 1, \dots, N\}$ applied by the XL-RIS during the i-th channel use. Consequently, the training signals at the XL-RIS can be expressed as

$$\mathbf{y}_{\mathrm{RIS}}^{\mathrm{tr}}(i) = \sqrt{\rho^{\mathrm{tr}}}\mathbf{h} + \mathbf{e}(i), \tag{27}$$

where ρ^{tr} is the power of the training signal, and $\mathbf{e}(i)$ is the additive EMI to $\mathbf{h}$. After RIS reflection, the training signal received by the m-th BS element is expressed as

$$y_m^{\text{tr}}(i) = \sqrt{\rho^{\text{tr}}}\phi(i)^{\mathsf{T}}\mathbf{x}_m + w_m(i) + z_m(i), \tag{28}$$

where

$$w_m(i) = \phi(i)^{\mathsf{T}}\left(\mathbf{g}_m \odot \mathbf{e}(i)\right) \tag{29}$$

is the EMI reflected from the RIS to the BS, and $z_m(i) \sim \mathcal{N}_{\mathbb{C}}(0, \sigma_z^2)$ is the additive white Gaussian noise. It is important to note that the term $w_m(i)$ is influenced by the RIS configuration.

By collecting all the training signals received at the m-th BS element across τ channel uses during the training phase, we construct the vector $\mathbf{y}_m^{\text{tr}} = [y_m^{\text{tr}}(1), \dots, y_m^{\text{tr}}(\tau)]^{\mathsf{T}} \in \mathbb{C}^{\tau}$, which can be expressed as

$$\mathbf{y}_m^{\text{tr}} = \sqrt{\rho^{\text{tr}}}\mathbf{\Phi}_\tau\mathbf{x}_m + \mathbf{w}_m^{\text{tr}} + \mathbf{z}_m^{\text{tr}}, \tag{30}$$

where

$$\mathbf{\Phi}_\tau = [\phi(1), \phi(2), \cdots, \phi(\tau)]^{\mathsf{T}} \in \mathbb{C}^{\tau\times N}, \tag{31}$$

$\mathbf{w}_m^{\text{tr}} = [w_m(1), \dots, w_m(\tau)]^{\mathsf{T}}$ and $\mathbf{z}_m^{\text{tr}} = [z_m(1), \dots, z_m(\tau)]^{\mathsf{T}}$. Accordingly, the vector $\mathbf{y}_{\text{BS}}^{\text{tr}} = [\mathbf{y}_1^{\text{tr}\,\mathsf{T}}, \dots, \mathbf{y}_M^{\text{tr}\,\mathsf{T}}]^{\mathsf{T}} \in \mathbb{C}^{M\tau}$, which collects the signals received at the BS array during the training phase, can be expressed as

$$\mathbf{y}_{\text{BS}}^{\text{tr}} = \sqrt{\rho^{\text{tr}}}\mathbf{\Phi}_{M\tau}\mathbf{x} + \mathbf{w}^{\text{tr}} + \mathbf{z}^{\text{tr}} \tag{32}$$

with

$$\mathbf{\Phi}_{M\tau} = \mathbf{I}_M \otimes \mathbf{\Phi}_\tau, \tag{33}$$

and $\mathbf{x} = [\mathbf{x}_1^{\mathsf{T}}, \dots, \mathbf{x}_M^{\mathsf{T}}]^{\mathsf{T}} \in \mathbb{C}^{MN}$, $\mathbf{w}^{\text{tr}} = [\mathbf{w}_1^{\text{tr}\,\mathsf{T}}, \dots, \mathbf{w}_M^{\text{tr}\,\mathsf{T}}]^{\mathsf{T}} \in \mathbb{C}^{M\tau}$ and $\mathbf{z}^{\text{tr}} = [\mathbf{z}_1^{\text{tr}\,\mathsf{T}}, \dots, \mathbf{z}_M^{\text{tr}\,\mathsf{T}}]^{\mathsf{T}} \in \mathbb{C}^{M\tau}$.

Considering that the RIS can acquire statistical information of the channels and EMI through its sensing mode [35], we assume that the BS has knowledge of the correlation matrices $\mathbf{R}_x = \mathbb{E}\{\mathbf{x}\mathbf{x}^{\mathsf{H}}\}$ and $\mathbf{R}_w^{\text{tr}} = \frac{1}{\sigma_e^2}\mathbb{E}\{\mathbf{w}^{\text{tr}}(\mathbf{w}^{\text{tr}})^{\mathsf{H}}\}$. These are given by [34]

$$\mathbf{R}_x = \mathbf{R}_{g_n'} \otimes (\mathbf{R}_{g_m} \odot \mathbf{R}_h) = \mathbf{R}_{g_n'} \otimes \mathbf{R}_c \tag{34}$$

and

$$\mathbf{R}_w^{\text{tr}} = \mathbf{R}_{g_n'} \otimes \left(\left(\mathbf{\Phi}_\tau\mathbf{R}_q\mathbf{\Phi}_\tau^{\mathsf{H}}\right) \odot \mathbf{I}_\tau\right), \tag{35}$$

where we define

$$\mathbf{R}_c = \mathbf{R}_{g_m} \odot \mathbf{R}_h \in \mathbb{C}^{N\times N} \tag{36}$$

and

$$\mathbf{R}_q = \mathbf{R}_{g_m} \odot \mathbf{R}_e \in \mathbb{C}^{N\times N} \tag{37}$$

for later use.

4.1.2 Linear MMSE estimator

The LMMSE estimate of $\mathbf{x}$ based on the observation $\mathbf{y}_{\text{BS}}^{\text{tr}}$ can be computed using the statistics provided in (34) and (35), as expressed by

$$\hat{\mathbf{x}} = \frac{1}{\sqrt{\rho^{\text{tr}}}}\mathbf{R}_x\mathbf{\Phi}_{M\tau}^{\mathsf{H}}\left(\mathbf{R}_y^{\text{tr}}\right)^{-1}\mathbf{y}_{\text{BS}}^{\text{tr}}, \tag{38}$$

where $\mathbf{R}_y^{\text{tr}} = \frac{1}{\rho^{\text{tr}}}\mathbb{E}\{\mathbf{y}_{\text{BS}}^{\text{tr}}(\mathbf{y}_{\text{BS}}^{\text{tr}})^{\text{H}}\}$ is given by

$$\mathbf{R}_y^{\text{tr}} = \boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x\boldsymbol{\Phi}_{M\mathcal{T}}^{\text{H}} + \frac{\sigma_e^2}{\rho^{\text{tr}}}\mathbf{R}_w^{\text{tr}} + \frac{\sigma_Z^2}{\rho^{\text{tr}}}\mathbf{I}_{M\mathcal{T}}. \tag{39}$$

Then, the correlation matrix of the estimation error $\tilde{\mathbf{x}} = \mathbf{x} - \hat{\mathbf{x}}$ is given by

$$\mathbf{R}_{\tilde{x}} = \mathbb{E}\{\tilde{\mathbf{x}}\tilde{\mathbf{x}}^{\text{H}}\} = \mathbf{R}_x - \mathbf{R}_x\boldsymbol{\Phi}_{M\mathcal{T}}^{\text{H}}\left(\mathbf{R}_y^{\text{tr}}\right)^{-1}\boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x. \tag{40}$$

Derived from (40), the Mean Square Error (MSE) of the LMMSE estimator can be expressed as

$$\mathcal{E}_x(\boldsymbol{\Phi}_\tau) = \text{tr}\left\{\mathbf{R}_{\tilde{x}}\right\} = \text{tr}\left\{\mathbf{R}_x - \mathbf{R}_x\boldsymbol{\Phi}_{M\mathcal{T}}^{\text{H}}\left(\mathbf{R}_y^{\text{tr}}\right)^{-1}\boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x\right\}, \tag{41}$$

which solely depends on $\boldsymbol{\Phi}_\tau$, as the channel statistics are assumed to be constant and provided.

4.1.3 XL-RIS phase-shift optimization

From (41), the MSE is solely determined by $\boldsymbol{\Phi}\mathcal{T}$. To minimize the estimation error, we can leverage its dependence on $\boldsymbol{\Phi}\mathcal{T}$ and optimize the MSE by solving the following problem

$$\min_{\boldsymbol{\Phi}_\tau \in \mathcal{F}} \ \mathcal{E}_x(\boldsymbol{\Phi}_\tau) \tag{42}$$

where the feasible set is defined as

$$\mathcal{F} = \{\boldsymbol{\Phi}_\tau \in \mathbb{C}^{\tau \times N} | \ |[\boldsymbol{\Phi}_\tau]_{i,n}| = 1; \forall i, n\}, \tag{43}$$

which reflects the fact that the RIS is a passive element and thus its coefficients must have unit modulus.

The optimization problem in (42) is not convex in $\boldsymbol{\Phi}_\tau$, and solving it requires optimization over large matrices. To address this, we reformulate (41) as

$$\mathcal{E}_x(\boldsymbol{\Phi}_\tau) = \text{tr}\left\{\mathbf{R}_x - \boldsymbol{\Lambda}(\boldsymbol{\Phi}_\tau)\boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x\right\}, \tag{44}$$

where $\boldsymbol{\Lambda}(\boldsymbol{\Phi}_\tau)$ is defined as

$$\boldsymbol{\Lambda}(\boldsymbol{\Phi}_\tau) = \mathbf{R}_x\boldsymbol{\Phi}_{M\mathcal{T}}^{\text{H}}\left(\mathbf{R}_y^{\text{tr}}\right)^{-1}. \tag{45}$$

Moreover, we introduce a simplifying assumption by neglecting the dependence of $\boldsymbol{\Lambda}(\boldsymbol{\Phi}_\tau)$ on $\boldsymbol{\Phi}_\tau$. Consequently, the objective function in (42) can be rewritten as

$$\mathcal{E}_x(\boldsymbol{\Phi}_\tau) = \mathcal{E}_x(\boldsymbol{\Lambda}, \boldsymbol{\Phi}_\tau), \tag{46}$$

treating $\boldsymbol{\Lambda}$ and $\boldsymbol{\Phi}_\tau$ as independent variables. Under this assumption, we can solve (42) by employing an AO approach, which iteratively optimizes $\boldsymbol{\Lambda}$ and $\boldsymbol{\Phi}_\tau$ in an alternating fashion.

Specifically, we propose a two-step iterative algorithm in which $\mathcal{E}_x(\boldsymbol{\Lambda}, \boldsymbol{\Phi}_\tau)$ is optimized alternately with respect to $\boldsymbol{\Lambda}$ and $\boldsymbol{\Phi}_\tau$. Let $\boldsymbol{\Lambda}^{(k)}$ and $\boldsymbol{\Phi}_\tau^{(k)}$ represent the values obtained at iteration k. At iteration $k+1$, the following steps are taken:

1. Having fixed the value of $\boldsymbol{\Phi}_\tau = \boldsymbol{\Phi}_\tau^{(k)}$, we minimize (46) by computing $\boldsymbol{\Lambda}^{(k+1)}$ as $\boldsymbol{\Lambda}(\boldsymbol{\Phi}_\tau)$ in (45). This step involves unconstrained optimization, directly applying the MSE minimization;
2. Fixing $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}^{(k+1)}$, the MSE is now expressed as

$$\begin{aligned}\mathcal{E}_x\left(\boldsymbol{\Lambda}, \boldsymbol{\Phi}_\tau\right) &= \text{tr}\left\{\boldsymbol{\Lambda}\mathbf{R}_y^{\text{tr}}\boldsymbol{\Lambda}^{\text{H}} - 2\Re\left[\boldsymbol{\Lambda}\boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x\right] + \mathbf{R}_x\right\} \\ &= \text{tr}\left\{\boldsymbol{\Lambda}\left(\boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x\boldsymbol{\Phi}_{M\mathcal{T}}^{\text{H}} + \frac{\sigma_e^2}{\rho^{\text{tr}}}\mathbf{R}_w^{\text{tr}} + \frac{\sigma_z^2}{\rho^{\text{tr}}}\mathbf{I}_{M\mathcal{T}}\right)\boldsymbol{\Lambda}^{\text{H}} - \boldsymbol{\Lambda}\boldsymbol{\Phi}_{M\mathcal{T}}\mathbf{R}_x - \mathbf{R}_x\boldsymbol{\Phi}_{M\mathcal{T}}^{\text{H}}\boldsymbol{\Lambda}^{\text{H}} + \mathbf{R}_x\right\}.\end{aligned} \tag{47}$$

   In (47), $\mathcal{E}_x(\boldsymbol{\Lambda}, \boldsymbol{\Phi}_\tau)$ becomes a convex function with respect to $\boldsymbol{\Phi}_\tau$. Thus, the RIS phase-shift matrix can be computed by solving the following minimization problem

$$\boldsymbol{\Phi}_\tau^{(k+1)} = \arg\min_{\boldsymbol{\Phi}_\tau \in \mathcal{F}} \mathcal{E}_x\left(\boldsymbol{\Lambda}^{(k+1)}, \boldsymbol{\Phi}_\tau\right). \tag{48}$$

Concerning the convergence of this AO algorithm, it can be observed that since the MSE is minimized in both steps, the MSE either decreases or stabilizes at each iteration. As the MSE is always a positive quantity, this iterative process will inevitably converge to a local optimum.

Although the objective function in (48) is convex with respect to $\mathbf{\Phi}_{\tau}$, the optimization problem remains non-convex due to the unitary modulus constraint. Since projecting any solution onto the feasible set $\mathcal{F}$ is straightforward, this type of constraint naturally leads to the use of the Projected Gradient (PG) method. In doing this, an additional auxiliary loop variable, $\mathbf{\Psi}^{(s)}$, is introduced to represent the intermediate RIS coefficients during the gradient descent process. To apply the PG method, it is necessary to compute the gradient and Hessian of (47) with respect to $\mathbf{\Phi}_{\tau}$, which are derived and detailed in [34]. Subsequently, the PG algorithm is initialized by setting $\mathbf{\Psi}^{(0)} = \mathbf{\Phi}_{\tau}^{(k)}$, which corresponds to the last solution obtained from the AO algorithm. The two steps involved in applying the PG method are:

1. The unconstrained RIS coefficient matrix is obtained by solving the unconstrained optimization problem using gradient descent. By utilizing the gradient $\nabla_{\mathbf{\Psi}}\mathcal{E}_x\left(\mathbf{\Lambda}, \mathbf{\Psi}^{(s)}\right)$, the update for $\mathbf{\Psi}^{(s)}$ is performed as follows:
$$\left[\mathbf{\Psi}^{(s+1)}\right]_{i,n} = \left[\mathbf{\Psi}^{(s)}\right]_{i,n} - \alpha\left[\mathbf{D}^{(s)}\nabla_{\mathbf{\Psi}}\mathcal{E}_x\left(\mathbf{\Lambda}, \mathbf{\Psi}^{(s)}\right)\right]_{z}, \tag{49}$$
where $\mathbf{D}^{(s)} \in \mathbb{C}^{\tau N \times \tau N}$ is an Hermitian positive definite matrix, and the value of $\mathbf{\Lambda}$ is derived from the outcome of the preceding AO solution. The matrix obtained upon convergence is denoted as $\mathbf{\Psi}^{\star}$.
2. Project $\mathbf{\Psi}^{\star}$ onto $\mathcal{F}$ by normalizing the amplitude of each entry to unity, that is,
$$\mathbf{\Phi}_{\tau}^{(k+1)} = e^{j\angle\mathbf{\Psi}^{\star}}. \tag{50}$$

Gradient descent is inherently iterative, and (49) may require multiple iterations to converge. Consequently, both the convergence rate and the selection of the step size α are crucial factors in the performance of such iterative methods. Specifically, when the Hessian of the objective function is available, we can select

$$\mathbf{D}^{(s)} = \left(\nabla^2_{\mathbf{\Psi}}\mathcal{E}_x\left(\mathbf{\Lambda}, \mathbf{\Psi}^{(s)}\right)\right)^{-1}. \tag{51}$$

In this case, the iterative algorithm is referred to as Newton's method, and has the great advantage of being able to find the minimum of a quadratic function as (47) with very few iterations [36]. However, considering the potentially high computational complexity associated with inverting the $\tau N \times \tau N$ Hessian matrix, a simplified variant of Newton's method can be employed. This approximation is valid when the diagonal elements of the Hessian are all strictly positive. In this simplified approach, the Hessian is approximated by the elements on its main diagonal, resulting in

$$\mathbf{D}^{(s)} = \left(\nabla^2_{\mathbf{\Psi}}\mathcal{E}_x\left(\mathbf{\Lambda}, \mathbf{\Psi}^{(s)}\right) \odot \mathbf{I}_{\tau N}\right)^{-1}. \tag{52}$$

In this specific case, the update rule for the diagonally scaled steepest descent method takes the form

$$\left[\mathbf{\Psi}^{(s+1)}\right]_{i,n} = \left[\mathbf{\Psi}^{(s)}\right]_{i,n} - \frac{\alpha\left[\nabla_{\mathbf{\Psi}}\mathcal{E}_x\right]_{(n-1)\tau+i}}{\left[\nabla^2_{\mathbf{\Psi}}\mathcal{E}_x\right]_{(n-1)\tau+i,(n-1)\tau+i}}. \tag{53}$$

By applying the AO algorithm in the outer loop and the PG algorithm in the inner loop, the optimal $\mathbf{\Phi}_{\text{opt}}$ with the minimum estimation error can be obtained. This procedure relies solely on the known spatial correlation matrices $\mathbf{R}_x$ and $\mathbf{R}_w^{\text{tr}}$, without requiring any online information. Therefore, the RIS phase-shifts for the training phase can be designed offline without consuming any channel uses.

4.1.4 Near-field channel estimation performance analysis

We now evaluate the performance of the proposed LMMSE estimator with AO in terms of the NMSE. For simplicity, we consider an RIS-aided Single-Input-Single-Output (SISO) system, i.e., $M_H = M_V = 1$, except for Figure 6 (b). Unless stated otherwise, the proposed system operates at $f_0 = 0.3\,\text{THz}$, corresponding to

Table 2: The parameter used for the near-field simulations.

Wavelength	$\lambda_c = 0.001$m ($f_c = 0.3$ THz)		
Number of RIS elements	$N = 12 \times 2$		
Inter-element spacing	$d_N = 10\lambda$		
SNR	0 dB		
SIR	5 dB		
Pilot length	$\tau = \text{rank}\{\mathbf{R}_{g_m} \odot \mathbf{R}_h\} = 15$		
Location of UE relative to XL-RIS	$r_h = 15.0\,\text{m}$	$\varphi_h = 70.0^\circ$	$\theta_h = -20.0^\circ$
Location of BS relative to XL-RIS	$r_g = 20.0\,\text{m}$	$\varphi_h = -60.0^\circ$	$\theta_h = -30.0^\circ$
Location of EMI relative to XL-RIS	$r_e = 25.0\,\text{m}$	$\varphi_e = -10.0^\circ$	$\theta_e = 20.0^\circ$
Elevation angular spread of **h**	$\Delta_{\theta_h} = 1^\circ$		
Elevation angular spread of **g**	$\Delta_{\theta_g} = 1^\circ$		
Elevation angular spread of EMI	$\Delta_{\theta_e} = 3^\circ$		

$\lambda = 1\,\text{mm}$. The RIS is equipped with a 12×2 UPA with $d_N = 10\lambda$. This configuration results in a Fraunhofer distance $R = 29.6\,\text{m}$. The UE and BS are positioned at $(15\,\text{m}, 70^\circ, -20^\circ)$ and $(20\,\text{m}, -60^\circ, -30^\circ)$ relative to the RIS, respectively, and both are located within the NF region of the RIS. The random EMI is generated at $(25\,\text{m}, -10^\circ, 20^\circ)$ relative to the RIS. Other simulation parameters are provided in Table 2.

In this part, we compare three different methods:

1. The first baseline method is the Reduced-Space Least-Squares (RS-LS) estimator given by [37, Eq. (13)], which is labeled as ‘RS-LS’ in the figures. The optimal RIS phase-shift for this estimator can be designed according to [38, Eq. (26)]. This phase-shift is also used as the initial RIS configuration $\mathbf{\Phi}_\tau^{(0)}$ in AO algorithm;
2. The second is the proposed LMMSE estimator with the AO strategy, labeled as ‘LMMSE’ in the figures;
3. To evaluate the gain introduced by the AO strategy, we also include a baseline LMMSE estimator that directly uses $\mathbf{\Phi}_\tau^{(0)}$ without executing the AO algorithm. This estimator is labeled as ‘LMMSE-$\mathbf{\Phi}_0$’ in the figures.

Moreover, in the proposed AO approach, we set $\alpha = 0.5$ and $\epsilon = 10^{-3}$. To accelerate the simulation, the diagonally scaled steepest descent method is applied only once per PG loop.

Figure 6(a) and Figure 6(b) illustrate the NMSE of various estimators as a function of SNR in the SISO and Multiple-Input-Single-Output (MISO) scenarios, respectively. In these figures, the pilot length is set to $\tau = \text{rank}\{\mathbf{R}_{g_m} \odot \mathbf{R}_h\} = 15$, which corresponds to the minimum number of pilots required by the RS-LS estimator [38]. In the MISO case, the BS is equipped with a 2×1 antenna with $d_M = \lambda = 1$ mm. For comparison, we also include the result of the proposed estimator in the interference-free case (i.e., $\text{SIR} = \infty$), denoted with the suffix ‘-w/o-EMI’. It can be observed that the NMSE of all estimators decreases with increasing SNR. In the presence of EMI, it can be clearly observed that when the SNR is lower than the Signal-to-Interference Ratio (SIR), the NMSE of all estimators consistently improves with increasing SNR. However, when SNR $\geq$ SIR, the performance of the estimators tends to saturate, indicating that the SIR becomes the dominant factor limiting performance in this regime. Moreover, although both ‘RS-LS’ and ‘LMMSE-$\mathbf{\Phi}_0$’ use the same RIS configuration, ‘LMMSE-$\mathbf{\Phi}_0$’ achieves approximately a 5-dB gain over ‘RS-LS’ in Figure 6(a) due to the LMMSE criterion. More importantly, the proposed LMMSE estimator consistently outperforms ‘LMMSE-$\mathbf{\Phi}_0$’ by at least

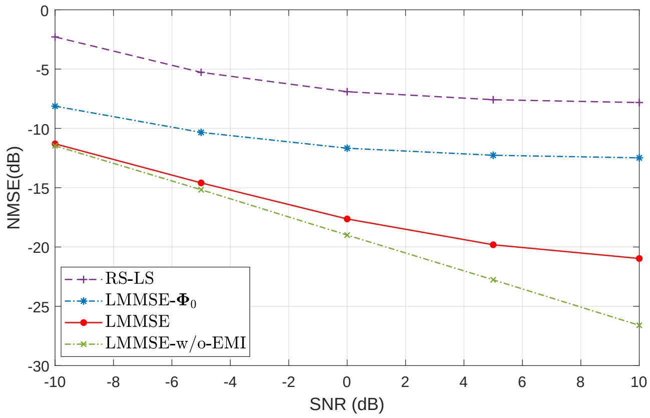

(a) The SISO case at $f_0 = 0.3\,$THz.

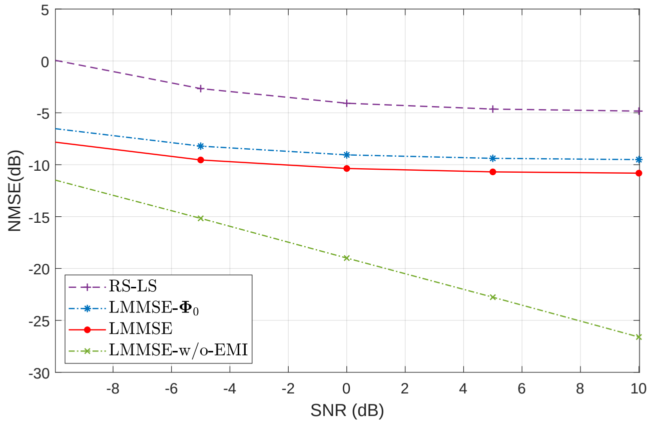

(b) The MISO case at $f_0 = 0.3\,$THz.

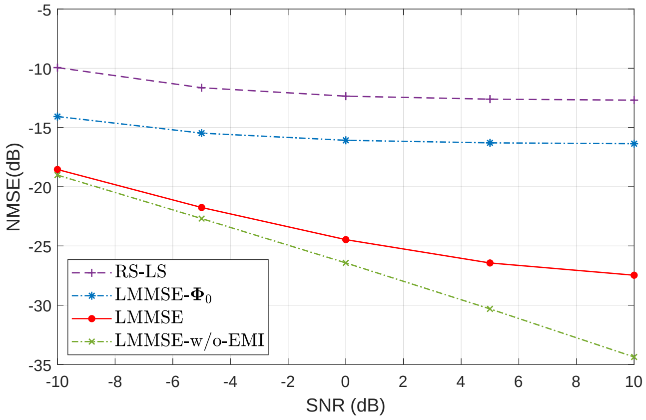

(c) The SISO case at $f_0 = 3\,$THz.

Figure 6: The NMSE vs. SNR for different estimators.

3.3 dB, demonstrating the effectiveness of the AO algorithm. Comparing Figure 6(a) and Figure 6(b), the performance of all estimators in the MISO case closely aligns with that observed in the SISO case. However, due to the increased dimensionality of the channel to be estimated, the overall performance of the estimators is degraded compared to the SISO scenario. Subsequently, Figure 6(c) illustrates the performance of various estimators in the SISO case at $f_0 = 3\,$THz. To ensure that the UE, BS, and EMI all lie within the NF region of the RIS, the number of RIS elements is increased to $N = 36 \times 4$. Under this setting, all estimators exhibit performance trends consistent with those observed in Figure6(a), thereby confirming the broad applicability of

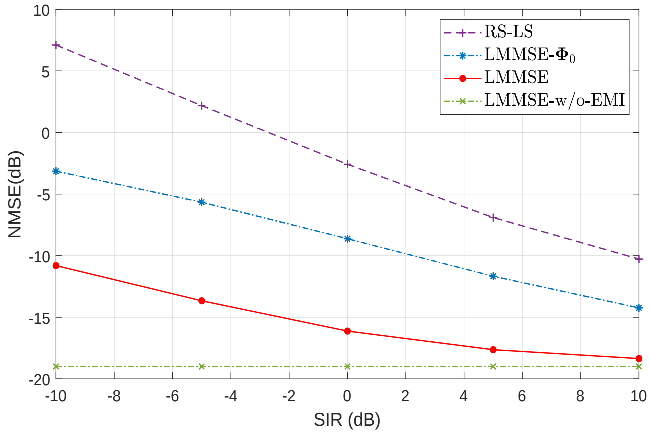

Figure 7: The NMSE vs. SIR for different estimators.

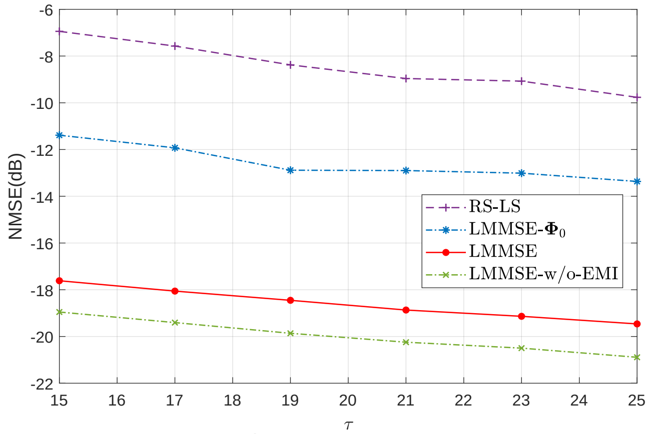

Figure 8: The NMSE vs. τ for different estimators.

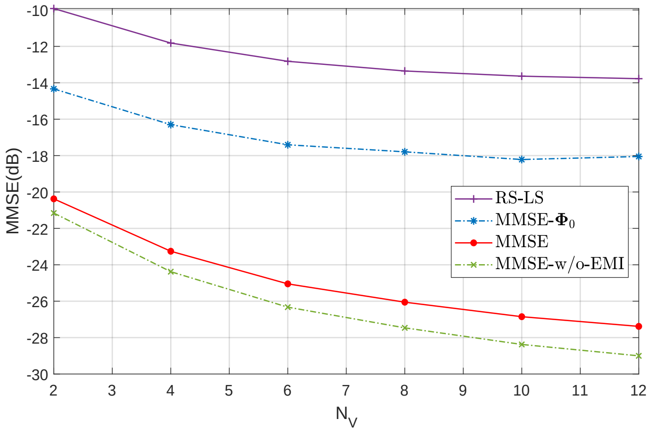

Figure 9: The NMSE vs. N_V for different estimators.

the proposed method in the THz band.

Figure 7 shows the NMSE as a function of SIR when $\text{SNR} = 0\,\text{dB}$. As expected, the NMSE of all estimators decreases as the SIR increases. Besides, consistent with earlier observations, the proposed LMMSE estimator achieves the best performance and gradually approaches the 'LMMSE-w/o-EMI' baseline as SIR increases.

Figure 8 presents the NMSE as a function of the pilot length τ at $\text{SNR} = 0\,\text{dB}$ and $\text{SIR} = 5\,\text{dB}$. The range of τ spans from $\tau = \text{rank}\{\mathbf{R}_{g_m} \odot \mathbf{R}_h\} = 15$ to N. As expected, the performance of all estimators naturally improves with increasing τ. Besides, the proposed LMMSE estimator consistently achieves superior performance, offering stable gains of about $10.5\,\text{dB}$ and $6.2\,\text{dB}$ over 'RS-LS' and 'LMMSE-$\mathbf{\Phi}_0$', respectively.

Figure 9 depicts the NMSE as a function of the number of RIS elements N at $\mathrm{SNR} = 0\,$dB and $\mathrm{SIR} = 5\,$dB, where N_H is fixed at 12, while N_V varies from 2 to 12. Note that increasing N also increases the required pilot length for channel estimation. However, the impact of increasing N on the cascaded channel rank, i.e., the effective channel dimension, is limited. Therefore, the dominant effect of a larger RIS is the enhancement of the RIS gain, which benefits near-field channel estimation. Accordingly, the NMSE of all estimators decreases as N increases. It is also worth noting that the performance gap between 'LMMSE' and both 'RS-LS' and 'LMMSE-$\mathbf{\Phi}_0$' widens with increasing N.

4.2 LoS-MIMO capacity evaluation in RIS-enhanced factory environments

When a MIMO system is operated in its near-field under LoS conditions, the spherical model of wave propagation enables the creation of independent channel paths, rendering the corresponding LoS-MIMO channel matrix to have a rank larger than one. This phenomenon can be effectively used for facilitating spatial multiplexing transmission under LoS to a single user, yielding an increase of its communication capacity. In RIS-enhanced environments, where the link between a BS and a UE is established through the reflection of a LoS path via a RIS, the size of the RIS effectively determines the degree of spatial multiplexing that can be supported for this link. Transmitter and receiver both need to be equipped with multiple antennas to allow for the separation of the spatially multiplexed beams and for precisely shaping the beam direction towards the RIS.

The solution for beam design and configuration of the RIS weights to enable high-rank LoS-MIMO transmission in RIS-enhanced communication links has been proposed in [39]. In TIMES project, the approach proposed therein has been reused to study the practically achievable capacity of a UE moving through a factory hall and being served by a BS through a RIS-enhanced communication link. Initially, the capacity for different RIS deployment setups has been evaluated assuming ideal conditions (i.e., continuous phase shifts for the RIS weights, ideal channel knowledge), where results indicate a high potential to increase the link capacity by using the strongest eigenchannels of the LoS-MIMO channel. However, analyzing the sensitivity with respect to practical impairments, such as displacement errors and cross-talk between adjacent RISs, reveals that these gains quickly deteriorate, in particular if eigenchannels of higher order are used.

4.2.1 Evaluation scenario and system model

The scenario under study is depicted in Figure 10. A cuboid factory hall is assumed of width $W = 20$ m, length $L = 30$ m and height $H = 10$ m. The (x, y, z) coordinate system is defined as shown in the figure with its origin in the (lower) top left corner of the hall. The BS is mounted at a height $h = 8$ m, i.e., 2 meters below the ceiling, in the center of the hall at coordinates $(15, 10, 8)$. Two RISs are mounted to the left wall at height $h = 4$ m with equal distance d to the x-coordinate of the BS, given by the coordinates $(15 \pm d, 0, 4)$. Finally, a UE with height $h = 1$ m centered in the hall moves along the x-axis from one end of the hall to the other, i.e., $(x, 10, 1)$ with $x \in [0, 30]$. BS, UE and RIS are assumed to use a UPA with K antennas (BS, UE) and N antennas (RIS), respectively, where $N \gg K$. The UPA is assumed to be centered at the respective location of the communication node, ant its orientation is assumed perpendicular to the y/z plane. While the UPAs at the RIS are planar with the wall they are mounted at, the UPAs at BS and UE are assumed to be tilted by 45° towards the RIS.

The UE is served by LoS links via RIS 1 (the right one from BS perspective) as long as it moves along the x-axis in the range $x \in [0, 15]$, while it is served via RIS 2 thereafter, $x \in [15, 30]$. In the initial investigations, possible cross-talk from the other RIS not used for serving the UE is ignored, whereas later the potential power from the interfering RIS will be analyzed. The near-field channels between RIS and BS and RIS and UE, respectively, are modelled according to the free-space propagation model, yielding for channel coefficient between the n-th RIS element and the k-th element at the BS (or UE, respectively):

$$h_{nk} = \frac{\lambda}{4\pi d_{nk}} \exp\left(j\frac{2\pi}{\lambda} d_{nk}\right) \tag{54}$$

with d_{nk} the distance between antenna element n at the RIS and k at the BS/UE. For calculating the capacity achievable by spatial multiplexing over the RIS-enhanced LoS channels, we reuse the solution proposed in [39]. Therein, the authors have shown that the complex unit-amplitude RIS weights w_n achieving the capacity of the overall LoS-MIMO channel are obtained for each RIS element n by the product of two phase factors, which represent the phase shift between that RIS element and the center of the BS UPA, and the phase shift between that RIS element and the center of the UE UPA, respectively:

$$w_n = \exp\left(j\frac{2\pi}{\lambda} d_n^{UE}\right) \cdot \exp\left(-j\frac{2\pi}{\lambda} d_n^{BS}\right) \tag{55}$$

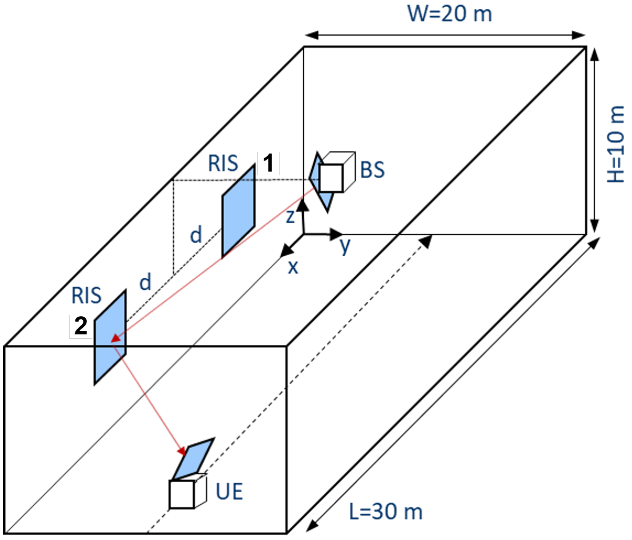

Figure 10: RIS-enhanced deployment scenario in factory hall

These RIS weights can be stacked into a diagonal (unitary) matrix $\mathbf{\Phi}$ of dimension $N \times N$. Similarly, by stacking the channel coefficients for the BS-RIS link into a $N \times K$ matrix $\mathbf{H}_B$ and those for the RIS-UE link into a $K \times N$ matrix $\mathbf{H}_U$, the effective channel for the RIS-enhanced link can be given as

$$\mathbf{H}_{eff} = \mathbf{H}_U \mathbf{\Phi} \mathbf{H}_B \tag{56}$$

Since this effective channel has a rank larger one when BS and UE are in the near-field of RIS, it is well known that the optimum transmission scheme is the singular value decomposition (SVD)-based precoding, which enables to use the eigenchannels for transmission. The authors in [39] have shown that the left and right singular vectors of $\mathbf{H}_{eff}$ to be used as the pre- and decoders for Eigenchannel transmission are identical to the left singular vector of $\mathbf{H}_U$ and the right singular vector of $\mathbf{H}_B$, respectively. Further, they have shown that each singular value of $\mathbf{H}_{eff}$ is identical to the product of the corresponding singular values of $\mathbf{H}_U$ and $\mathbf{H}_B$.

To achieve the maximum capacity C, optimal power distribution over the eigenchannels of $\mathbf{H}_{eff}$ according to the Waterfilling algorithm is mandated. However, the authors have further shown that, as long as the magnitudes of the eigenchannels σ_i of the L channels selected for transmission, $i \in \{1, \dots, L\}$, do not differ by a factor larger than 10, the difference to the much simpler equal power allocation scheme is insignificant. For the capacity evaluations carried out here, we have therefore selected the number of eigenchannels achieving the highest capacity under the equal power allocation scheme, where the above requirement is safely satisfied:

$$C = \max_L \sum_{i=1}^{L} \log\left(1 + \frac{P_0}{L N_0} \sigma_i^2\right), \tag{57}$$

where P_0 and N_0 is the transmit power and noise power, respectively.

4.2.2 Capacity gains under ideal assumptions

The system evaluation is carried out for a carrier frequency $f = 120$ GHz, yielding a wavelength $\lambda = 2.5$ mm. The assumed bandwidth is $B = 10^8$ Hz = 100 MHz, which leads to a noise power $N_0 = -174 + 80 = -94$ dBm. The transmit power is set to $P_0 = 20$ dBm. At the BS and the UE, an UPA with 8×8 antennas is assumed with an antenna spacing of 8λ, while the RIS is assumed to have an UPA with 64×64 RIS elements with a spacing of 2λ. Note that the choice of these parameters (including the carrier frequency) is mainly due to complexity issues. As a matter of fact, higher frequencies can lead to a reduced spacing between the elements, increasing the RIS size (in terms of the number of RIS elements) and its related complexity, but without achieving a significant benefit in terms of performance. Hence, the carrier frequency has been chosen at the lower end of the sub-THz range, as it allows for a general analysis of the near-field effects of LoS-MIMO in the given scenario while selecting parameters of reasonable magnitude for the RIS configuration, thus keeping the complexity in balance. For supporting higher frequencies, the RIS configuration needs to be scaled up accordingly to maintain the identical near-field range. Finally, the bandwidth has been selected to ensure achieving a SNR for the given signal power budget that is high enough for enabling the activation of additional eigenchannels when the UE moves closer to the RIS.

Capacity evaluation results for the UE moving along the x-axis in the center of the factory hall are shown in Figure 11 (a). As the baseline scenario, a direct communication between BS and UE (i.e, without RIS enhancements) is considered, where the tilting of the UPAs at BS and UE is set to 90° (i.e., parallel to the ceiling) to align their orientation with the direct LoS link between the two. In this case, the channel matrix for this LoS link has rank one, i.e., only one Eigenchannel, and the corresponding plot in the figure exhibits the lowest capacity. When the RIS-enhanced deployment is considered, a significant boost of the capacity can be observed, which reaches gains up to 50% compared to the baseline. This capacity boost is attributed to a substantially increased number of activated eigenchannels in the LoS-MIMO link, as shown in Figure 11 (b), where it is seen that up to 5 eigenchannels are activated. The different placements of the RIS considered in Figure 11, characterized by the distance d to the BS, reveal that by placing the RIS close to the BS, a high capacity – supported by a large number of activated eigenchannels – can be achieved close to the center of the hall, while it drops steeply when the UE moves towards the corners. Opposed to that, a placement of the RIS further away from the BS exhibits a more balanced capacity distribution in the factory hall, which is supported by a constant number of activated eigenchannels; the average and peak capacity is significantly lowered compared to the former case, though.

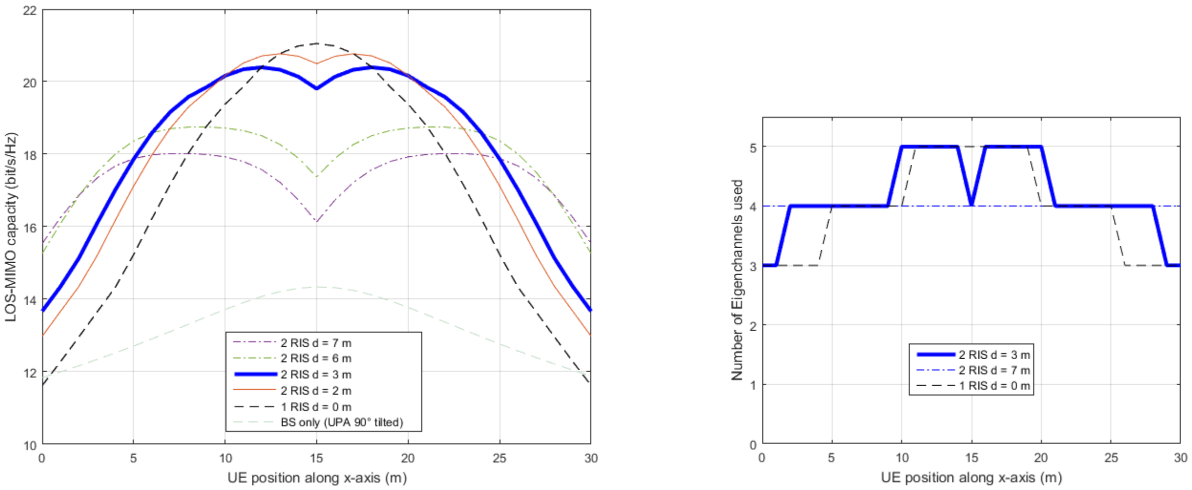

(a) LoS-MIMO capacity

(b) Number of activated eigenchannels

Figure 11: (a) LoS-MIMO capacity and (b) number of used eigenchannels vs. UE position for different RIS placements.

4.2.3 Sensitivity analysis

For the sensitivity analysis, the RIS placement is fixed to $d = 3$ m. The focus is first turned on the sensitivity of the Eigenvectors used for beamforming at UE and BS with respect to UE displacements in x, y and z direction. In particular, it is assumed that these beamformers are determined as the left singular vectors $\mathbf{u}_i$ of the channel $\mathbf{H}_U$ between RIS and the UE located at a pre-defined position, but the actual channel $\hat{\mathbf{H}}_U$ used for transmission is then calculated for the UE position with an offset in x, y or z direction. Since the beamformers at the BS are determined as the right singular vectors $\mathbf{v}_i$ of the channel $\mathbf{H}_B$ between RIS and BS, these are not affected by the UE displacement. The effective channel thus becomes

$$\hat{\mathbf{H}}_{eff} = \hat{\mathbf{H}}_U \mathbf{\Phi} \mathbf{H}_B \, . \tag{58}$$

The useful signal power P_i of the i-th Eigenchannel and the interference power Z_i of the other right singular vectors $\mathbf{v}_j$ distorting the i-th Eigenchannel then amount to

$$\begin{aligned} P_i &= P_0/L \, ||\mathbf{u}_i^H \hat{\mathbf{H}}_{eff} \mathbf{v}_i||^2 && (59) \\ Z_i &= P_0/L \sum_{j=1, j\neq i}^{L} ||\mathbf{u}_i^H \hat{\mathbf{H}}_{eff} \mathbf{v}_j||^2 \, . && (60) \end{aligned}$$

The SINR for the i-th Eigenchannel then yields $\frac{P_i}{Z_i+N_0}$, which allows to calculate the capacity achievable under UE displacement according to

$$C_{\text{displacement}} = \sum_{i=1}^{L} \log\left(1 + \frac{P_i}{Z_i + N_0}\right) \, . \tag{61}$$

The drop in capacity for UE displacements of different offsets is shown in Figure 12(a). We have investigated displacements in x, y and z direction by an offset of 2 cm $= 8\lambda$ and 5 cm $= 20\lambda$, respectively. It is observed that the lowest sensitivity is for offsets in the y-direction, which is the direction perpendicular to the RIS. The sensitivity for offsets in x and z direction, i.e. parallel to the RIS surface, is significantly more pronounced, though: It is observed that a displacement of 5 cm in x-direction yields a capacity drop by 20% close to the edges and by 50% in the worst case. The strongest capacity drop is observed at the x-positions where the RISes are located (x=12 m and x=18 m), i.e, where the LoS link between UE and RIS is closest to each RIS' perpendicular. Moreover, we observe from the plots that the most pronounced drops in capacity appear in the region with the highest number of activated eigenchannels; compare with Figure 11 (b). Hence, we also plotted for the maximum displacement considered the case where the total number of activated eigenchannels is limited to 3. From those plots, it can be observed that the capacity drops smoothen considerably, which suggests that reducing the number of eigenchannels in case of larger expected displacement is a useful approach to alleviate the induced capacity drop.

Up to now, it has been assumed that the UE is served via a single selected RIS, while the second RIS does not create any interference. However, this assumption holds true only if the second RIS can be switched off[2]. If it is not (which should be assumed as the common case now), then fractions of the power used to feed the first RIS by the BS precoders may be received at the second RIS, which could then create additional interference in the communication environment. This situation has been analyzed in Figure 12(b): While moving the UE in the range $x \in [0, 15]$ m, the UE is connected to RIS1 and will support up to 5 eigenchannels, as shown above. The figure depicts the power of each precoding vector used to create the Eigenchannel (subject to the equal transmit power allocation from above) that is received at RIS1 and the second RIS2. The received power is normalized to the one related to the strongest Eigenchannel received at RIS1, yielding a power level of 0 dB for the latter. It can be observed that for the precoding vector related to the strongest Eigenchannel, the power received by RIS2 lies more than 20 dB below the power level at RIS1, and hence it can be considered negligible. However, for the eigenchannels following, it is observed that while their power received at RIS1 decreases, their power received at RIS2 increases, diminishing this power gap, until both received powers achieve a similar level, as observed for the 5th eigenchannel. This finding underscores a severe challenge for the application of high-rank LoS-MIMO communication via RISes in practice: As long as the interference created by the power received at the second RIS cannot be mitigated otherwise, only few of the available eigenchannels may practically be used.

---

[2] RIS can be realized based on different technologies. There are indeed some technologies used in FR1 and FR2 that allow switching off the RIS elements completely; e.g., CMOS switches can be designed to disconnect the RIS elements or route signals to absorptive loads, effectively turning the RIS off.

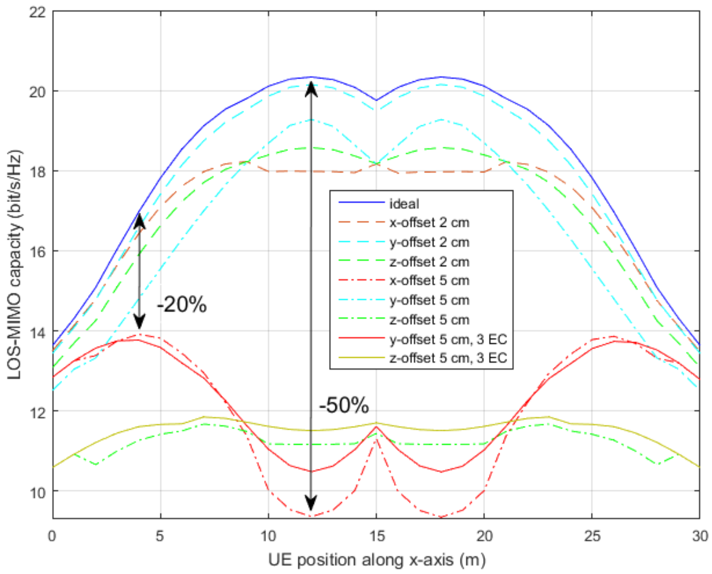

(a) Capacity drop due to UE displacement in x, y, z

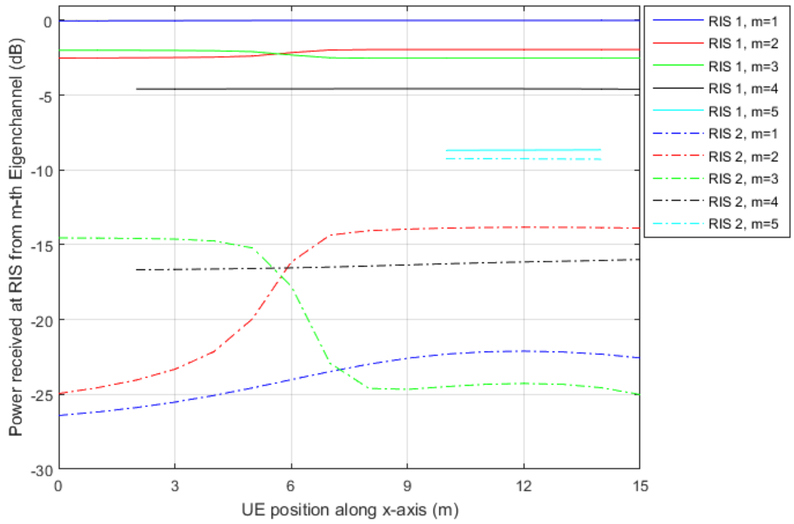

(b) Power received at RIS1 (desired) and RIS2 (cross-talk)

Figure 12: Sensitivity analysis: (a) Capacity drop due to UE displacements and (b) RIS cross-talk.

4.3 Solutions for mitigation of RF impairments on RIS-based communication links

This section introduces a transmit power adaptation approach to improve the performance of a RIS assisted communication link at THz frequencies, where the transmitter and the RIS are respectively impaired by PA non-linearities and phase errors at the RIS due to phase quantization at each of the RIS elements. Although one could, for instance, partially mitigate the considered impairments by e.g. applying a pre-distortion algorithm to linearize the output power at the transmitter, it was concluded that pre-distortion at such high frequencies is a challenging task in practice. Therefore, as we will show throughout this section, opting for a transmit power control approach without resorting to pre-distortion can provide considerable communication performance gains if transmit power values are wisely chosen.

4.3.1 System Model

The system model is comprised of a Node A that transmits information to Node B through a RIS due to the absence of a direct link. A single carrier waveform is assumed to be transmitted[3]. Given a QAM transmission system, the ideal transmitted $M \times 1$ data vector has the following form

$$\mathbf{x} = \left[a_1 e^{j\phi_1}, \dots, a_M e^{j\phi_M}\right]^T, \tag{62}$$

where $a_m,\ m \in \{1, \dots, Q\}$ belongs to a set of amplitudes specific to a certain modulation type, and ϕ_m represents the possible phase values. For instance, for a 4-QAM modulated signal, $Q = 1$ while for 16-QAM, $Q = 3$. It is also assumed that

$$\frac{\text{trace}\left(\mathbf{x}\mathbf{x}^H\right)}{M} = 1, \tag{63}$$

where $M = \text{card}(\mathbf{x})$. When PA non-linearity is present, the modified tranmsitted vector after the PA is modeled by applying

$$\begin{aligned} \tilde{\mathbf{x}} = \Big[& F_{\text{AMAM}}(\delta a_0) e^{j(F_{\text{AMPM}}(\delta a_0) + \phi_0)}, \dots, \\ & F_{\text{AMAM}}(\delta a_{M-1}) e^{j(F_{\text{AMPM}}(\delta a_{M-1}) + \phi_{M-1})}\Big]^T, \end{aligned} \tag{64}$$

where $\delta \in [0, \delta_{max}]$ is a control parameter that sets the desired transmission power to a certain value of interest, δ_{max} is a maximum power control limit parameter, $F_{\text{AMAM}}(x)$ is the amplitude-to-amplitude transfer function, and $F_{\text{AMPM}}(x)$ is the amplitude-to-phase modulation transfer function, which are , e.g., defined as in the equation set defined in (5).

After passing through the RIS assisted channel, it can be shown that the received signal maybe then written as

$$\mathbf{r}(m) = \sqrt{\alpha} \sum_{n=0}^{Q-1} e^{j(\theta_{t,n} + \theta_{r,n} + \omega_n)} \tilde{\mathbf{x}}(m) + \mathbf{v}(m), \tag{65}$$

where α is the path gain between Node A and Node B which was derived in [41], adopted in Deliverable 4.3 [3], and may be written as

$$\alpha = \frac{G_A G_B}{(4\pi)^2} \left(\frac{lw}{d_{AR} d_{RB}}\right)^2 \cos(\psi), \tag{66}$$

---

[3] In this deliverable, we adopt a single-carrier waveform where we align with the THz single carrier mode (THz-SC PHY) in IEEE Std 802.15.3-2023 intended for high data rate communications [40]. 4, 16 and 64-QAM modulation orders are proposed.

where l, w are the length and width of the reflecting element at the RIS, d_{AR} and d_{RB} are the separation distances between Node A and the RIS, and the RIS and Node B, respectively, G_A and G_B represent the gains of the transmitter and receiver antennas, respectively, and ψ is the angle of incidence of the reflecting wave. In (65), $\theta_{t,n}$ is the phase of the wave from Node A to the n^{th} RIS reflecting element, and $\theta_{r,n}$ is the phase of the wave from the n^{th} RIS reflecting element to Node B, and ω_n is the weight applied at the RIS to compensate for the phases at the n^{th} RIS reflecting element.

At Node B, the receiver attempts to detect the transmitted symbol. Since the RIS accounts for aligning the phases at the receiver (Node B), it only applies a gain control to rectify the concatenated propagation loss. Thus, Node B applies

$$\hat{\mathbf{r}}(m) = \frac{1}{\beta}\mathbf{r}(m), \tag{67}$$

$$= \frac{1}{\beta}\left(\sqrt{\alpha}\sum_{i=0}^{Q-1} e^{j\epsilon_i}\tilde{\mathbf{x}}(m) + \mathbf{v}(m)\right), \tag{68}$$

where $\beta = Q\sqrt{P\alpha}$, $P = \frac{\tilde{\mathbf{x}}^H\tilde{\mathbf{x}}}{M}$, ϵ_i represents the phase error resulting from the limited phase quantization at the RIS elements.

4.3.2 Performance Analysis

In this section, we carry out an analysis to quantify the performance of a RIS assisted communication link impaired by phase errors resulting from phase quantization and PA nonlinearities. The analysis is carried out by deriving the MSE at Node B.

The MSE may be calculated by applying

$$\varepsilon \triangleq \frac{1}{M}\mathbb{E}\left\{(\hat{\mathbf{r}} - \mathbf{x})^H(\hat{\mathbf{r}} - \mathbf{x})\right\}. \tag{69}$$

Towards deriving (69), the following calculations are needed

$$\mathbb{E}\left\{\text{trace}(\hat{\mathbf{r}}\hat{\mathbf{r}}^H)\right\} = \frac{M}{Q^2}\mathbb{E}_\epsilon\left\{\left|\sum_{i=0}^{Q-1} e^{j\epsilon_i}\right|^2\right\} + \frac{M}{\beta^2}\sigma_{\mathbf{v}}^2, \tag{70}$$

where this result was found by adopting the Additive White Gaussian Noise (AWGN) properties of the vector $\mathbf{v}$. Moreover, we need to calculate

$$\mathbb{E}\left\{\text{trace}(\mathbf{x}\hat{\mathbf{r}}^H)\right\} = \frac{M}{Q\sqrt{P}}\mathbb{E}_\epsilon\left\{\sum_{i=0}^{Q-1} e^{j\epsilon_i}\right\}P_{\tilde{\mathbf{x}}\mathbf{x}}, \tag{71}$$

where $P_{\tilde{\mathbf{x}}\mathbf{x}} \triangleq \tilde{\mathbf{x}}^H\mathbf{x}$. To perform the above calculations, we use the circular moment result of a continuous uniform random variable of the phase quantization error as specified in Section 3.1

$$\mathbb{E}_\epsilon\{e^{j\epsilon}\} = \text{sinc}\left(\frac{\pi}{L}\right). \tag{72}$$

Thus, exploiting the linearity of the expectation operator, we can write

$$\mathbb{E}\left\{\text{trace}(\mathbf{x}\hat{\mathbf{r}}^H)\right\} = \frac{M}{\sqrt{P}}\text{sinc}\left(\frac{\pi}{L}\right)P_{\tilde{\mathbf{x}}\mathbf{x}}. \tag{73}$$

Moreover, the expectation in (70) can be calculated by applying

$$\mathbb{E}\left\{\left|\sum_{i=0}^{Q-1} e^{j\epsilon_i}\right|^2\right\} = Q + \sum_{n=0}^{Q-1}\sum_{\substack{m=0\\ n\neq m}}^{Q-1}\mathbb{E}\left\{e^{j\epsilon_n}\right\}\mathbb{E}\left\{e^{j\epsilon_m}\right\} = Q + Q(Q-1)\text{sinc}^2\left(\frac{\pi}{L}\right), \tag{74}$$

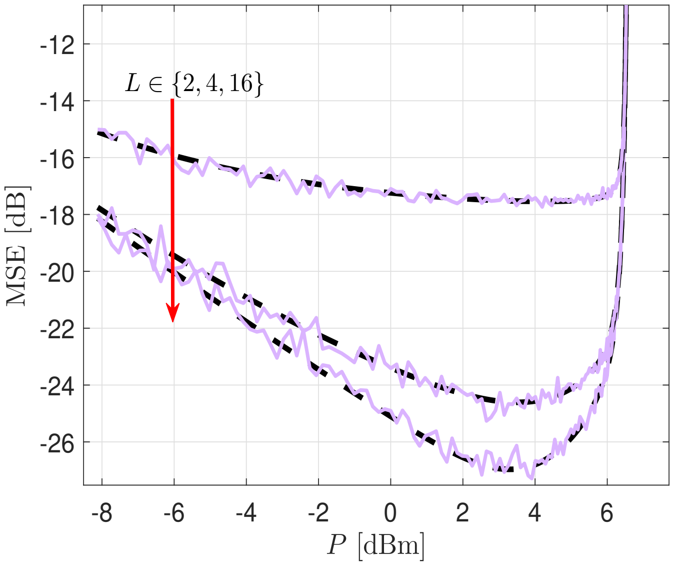

Figure 13: MSE vs. P for different quantization levels L.

where this result was achievable due to the assumption that the phase quantization errors at each of the RIS elements are independent and identically distributed. We then can re-write (73) as

$$\mathbb{E}\left\{\text{trace}(\hat{\mathbf{r}}\hat{\mathbf{r}}^H)\right\} = \frac{M}{Q}\left(1 + (Q-1)\text{sinc}^2\left(\frac{\pi}{L}\right)\right) + \frac{M}{\beta^2}\sigma_{\mathbf{v}}^2. \tag{75}$$

Using the results in (75) and (74), and substituting them in (69), it can be shown that

$$\varepsilon = \left(\frac{1}{Q} + \frac{Q-1}{Q}\mathfrak{sinc}^2\left(\frac{\pi}{2^L}\right)\right) + \frac{1}{\beta^2}\sigma_{\mathbf{v}}^2 - \frac{1}{\sqrt{P}}\text{sinc}\left(\frac{\pi}{L}\right)2\Re\left\{P_{\mathbf{x}\tilde{\mathbf{x}}}\right\} + 1. \tag{76}$$

The above expression will be tested using Monte Carlo simulations in the next Section.

4.3.3 Adaptive Transmit Power Control

As the transmitter (Node A) is impaired with PA non-linearity and the RIS is limited by its phase quantization capability, then herein, we can show how Node A, Node B and the RIS can jointly coordinate to minimize a certain performance metric, ε in this specific case, that reflects the quantization limitation at the RIS and the PA non-linearity. To justify the aim behind using adaptive power control and to verify the accuracy of the derived expression in (76), we first start by depicting Figure 13, where it presents the analytical vs. simulated results of the MSE as a function of P, and for different phase quantization levels L. Unless stated otherwise, the adopted simulation parameters are listed in Table 3.

The first observation is that as L is increased, the MSE is lower due to the more accurate tuning of the phases at each of the reflecting RIS elements. Moreover, for each quantization level L, and for small P, the MSE is relatively higher due to the overall effect of the channel, i.e. not meeting the link budget requirements. As P is increased, the MSE keeps improving (decreasing) until it starts degenerating again, and this time due to distortions introduced by PA nonlinearities. One can also notice the higher dependency of ε on P as L is increased. This can be explained by the improved effective channel when the RIS approaches almost perfect phase tuning, hence, increasing P becomes beneficial. However, when quantization is coarser, i.e. $L = 2$, increasing P brings little improvement to the MSE.

Table 3: Simulation parameters. The PA nonlinearity parameters are taken from [3]

Symbol	**Description**	**Value**
d_{AR}	Distance between node A and RIS	50m
d_{RB}	Distance between RIS and node B	50m
f_c	Carrier frequency	280 GHz
BW	Signal bandwidth	2 GHz
G_A	Antenna gain of node A	20 dB
G_B	Antenna gain of node B	20 dB
l	Length of the RIS element	$\lambda/2$
σ_v^2	Noise power	-80 dBm
ψ	The wave's angle of incidence	40°
Q	Number of RIS elements	2^{20}
w	Width of the RIS elements	$\lambda/2$
G	Small signal gain (AMAM model)	7.913
p	Smoothness factor (AMAM model)	1.132
V_{sat}	Saturation voltage (AMAM model)	0.0674 V
A	AMPM model parameter	-20075
B	AMPM model parameter	80.98
q_1	AMPM model parameter	2.109
q_2	AMPM model parameter	31.45

Given this discussion, we aim at selecting the output power P such that ε is minimized, i.e.

$$\arg\min_{P} \varepsilon, \tag{77}$$

$$\text{s. t. } v_{in} \in \mathcal{R} \tag{78}$$

where $\mathcal{R}$ is the range of the PA input voltage (v_{in}) that the PA is characterized within. One way to solve the above problem is to find the partial derivative

$$\frac{\partial \varepsilon}{\partial P} = 0, \tag{79}$$

and then solve for P. Since this is a complex task and the function can be seen to be convex for the adopted range of output powers, we can simply use the analytical result and find the unique minimum of the function to get P that minimized the MSE. The parameters used in Eq. (76), can be forwarded to the transmitter to perform the above optimization. For instance, the receiver noise level σ_v^2 and β (can be estimated in the channel estimation process at Node B) can be shared between the Nodes B and A, while Q and L can be shared with Node A through a control communication channel between Node A and the RIS using a control communication module at the RIS [42]. Figure 14 shows the effect of changing the dimensions of the RIS, in this case the parameter Q and for $L = 16$. As can be seen form the results, it is clear that the MSE overall increases when Q is decreased. Further, the optimal power transmission power increases as the Q is decreased. This is obviously due to the decrease of the number of reflecting elements, and the system

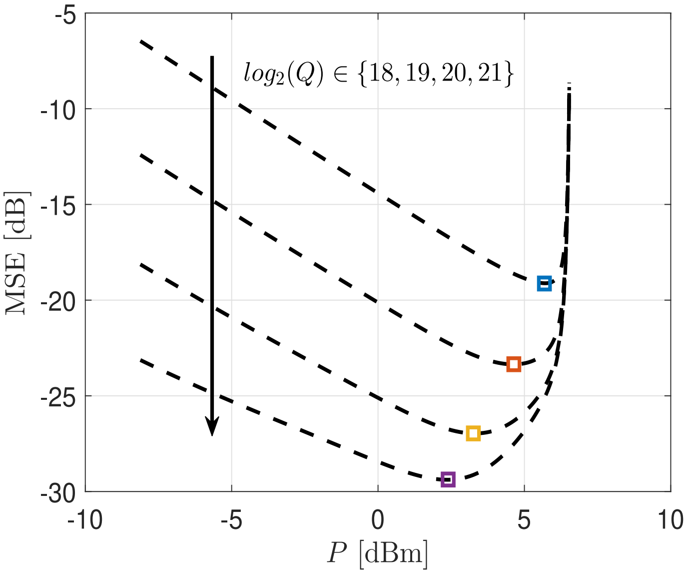

Figure 14: MSE vs. P *for different sizes of the RIS* Q*.* $L = 16$*. Square labels indicate minimum MSE.*

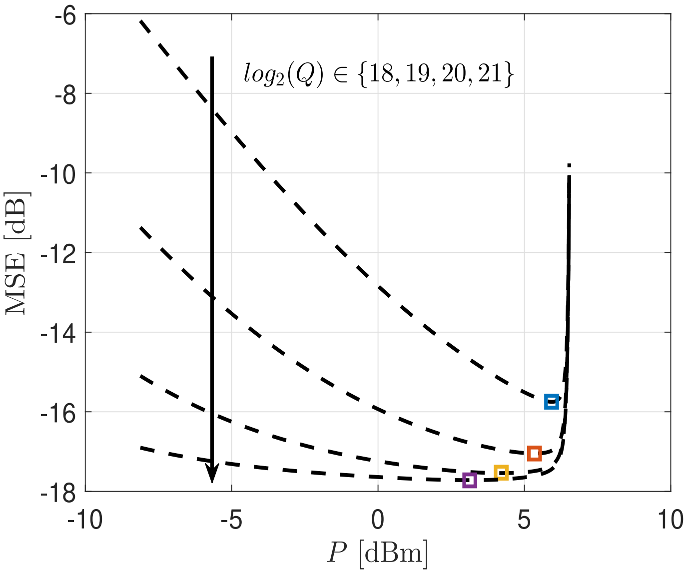

Figure 15: MSE vs. P *for different sizes of the RIS* Q*.* $L = 2$*. Square labels indicate minimum MSE.*

affords relatively more errors caused by the transmitter PA to compensate for the lack of enough reflections that improves the performance at the receiver. Moreover, we depict in Figure 15 the same MSE vs. P, but this time for a lower L, namely $L = 2$. In this case, the same trends observed in Figure 14 can be deduced, but with an increased overall MSE due to the lower quantization capability at the RIS. An important observation is that the optimal transmit power does not significantly improve the MSE for different values of Q. The reason for this behavior is that the system is overwhelmed by the quantization errors, reducing the gains obtained by varying the transmit power.

4.4 Near field sensing for theoretical estimation of target position and speed

Unlike traditional far-field systems, where the planar wavefront assumptions dominate, near-field systems operate in a regime where spherical wavefronts and spatial signal variations become prominent. These effects are especially pronounced with the emergence of Electrically Large Aperture Arrays (ELAAs) , which promise to push the boundaries of localization and sensing precision [43]. In this Section, we investigate the impact of the adoption of ELAAs on the estimation of some parameters of interest for localization and sensing, in particular target position (range and Angle of Arrival (AoA)) and speed. Specifically, a theoretical analysis is performed, deriving the corresponding Cramér-Rao Lower Bounds (CRLBs) to highlight how system parameters, the propagation regime, and array geometry influence the estimation accuracy. In the following, two cases are examined: the asymptotic CRLB derivation for target position estimation differentiating the analysis between localization of an active UE and sensing of a passive target, in a condition for which the distance to the UE/target is less than the array aperture; the performance analysis in terms of CRLB for velocity estimation in the near field of the ELAA, that allows the estimation not only of the radial component, as in far field, but also of the transverse one.

4.4.1 Asymptotic Behavior of Near-Field Localization and Sensing

The adoption of ELAAs introduces new challenges to be tackled. Specifically, a reevaluation of traditional models and performance limits is needed. Indeed, the well-known path-loss scaling laws and the number of available communication modes fail when the transmitter-receiver distance approaches the array size. Although this phenomenon has been widely investigated for communication, it remains less understood in the field of localization and sensing. In this section, we derive and evaluate the performance limits for range and AoA estimation in terms of CRLB, to assess how the estimation limits change as traditional scaling laws are no longer fulfilled.

System model Two cases are evaluated: (a) localization of a single-antenna UE using a BS equipped with an ELAA of K elements; (b) sensing of a passive target using a Single Input Multiple Output (SIMO) monostatic configuration at the BS, namely with a single transmitter and K receiving elements. The ELAA is arranged in a uniform linear deployment and the antenna aperture is defined as $D_{\text{elaa}} = (K-1)\delta$, where δ is the antenna spacing. The reference system is placed with the origin corresponding to the central element of the ELAA at the BS, which is oriented along the x axis. Denote with $\tilde{\mathbf{p}}_k = [x_k, y_k]^\top$ the position of the k-th receiving antenna at the BS , for $k = -(K-1)/2, \dots, (K-1)/2$, so that $y_k = 0$, $\forall k$, and $x_0 = 0$. The position of the UE/target, instead, is expressed in a polar coordinate system for convenience, as $\mathbf{p} = [\theta, d]^\top$ where θ denotes the angle with respect to the broadside direction of the receiving ELAA (i.e., the AoA of the signal received at the BS), and d denotes the distance between the UE/target and the central element of the receiving ELAA at the BS (i.e., the reference antenna). We denote by $d_k = ||\tilde{\mathbf{p}} - \tilde{\mathbf{p}}_k||$ the distance between the k-th receiving antenna and the UE/target

$$d_k = \sqrt{d^2 \sin^2\theta + x_k^2 - 2x_k d \sin\theta + d^2\cos^2\theta} = d\sqrt{1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}}. \qquad (80)$$

When considering case (a) of localization, the UE transmits an uplink data packet spanning N subcarriers of an Orthogonal Frequency Division Multiplexing (OFDM) signal. On the other hand, when considering case (b) of sensing, the central element of the BS transmits the OFDM signal. This signal is then reflected by the target in position $\mathbf{p}$, and all antenna elements of the BS receive the signal, as for the active localization case (SIMO radar). The BS aims at estimating the distance $d = ||\tilde{\mathbf{p}}||$ (i.e., the *range*) and the angle θ (i.e., the AoA) processing the signal received at its ELAA. After standard cyclic prefix removal and Fast Fourier Transform (FFT) processing, the received signal for a given OFDM symbol, in complex baseband, can be written as [44]

$$r_{n,k} = y_{n,k} + z_{n,k} = \sqrt{P}\, x_n \beta_{n,k}\, e^{-\mathrm{j}2\pi f_n \tau_k} e^{\mathrm{j}\varphi} + z_{n,k} \qquad (81)$$

where $P = P_{\mathrm{T}}/N$ is the power allocated to each subcarrier, P_{T} is the total transmit power, and x_n is the data symbol (e.g., a pilot used for channel estimation and localization) transmitted in the n-th subcarrier with $\mathbb{E}\left\{|x_n|^2\right\} = 1$, for $n = -(N-1)/2, \dots, (N-1)/2$. Moreover, $f_n \triangleq f_{\mathrm{c}} + n\Delta f$ is the frequency associated with the n-th subcarrier, Δf is the subcarrier spacing (SCS), φ accounts for the phase synchronization mismatch between the transmitter and the receiver, and $\beta_{n,k}$ are the channel scaling coefficients. The term τ_k denotes the time taken by the signal to travel between the UE and the k-th BS antenna (for localization) or between the transmitting antenna, the target and the k-th BS antenna (for sensing). The term $z_{n,k}$ denotes the AWGN, with $z_{n,k} \sim \mathcal{CN}(0, \sigma^2)$. Assuming a small bandwidth, we have $f_n \approx f_{\mathrm{c}}$, and the channel scaling coefficients can be considered constant on the different subcarriers, i.e., $\beta_{n,k} = \beta_k \ \forall n$. As a consequence, the SNR at the k-th receiving antenna can be expressed as $\mathrm{SNR}_k = P\beta_k^2/\sigma^2 = \alpha_k^2/\sigma^2$. The channel scaling coefficients for localization (a) and sensing (b) are given by the Friis laws in the active and passive case, that are, respectively

$$\beta_k^{2(\mathrm{loc})} = \frac{G_{\mathrm{T}} G_{\mathrm{R}} \lambda^2}{(4\pi d_k)^2}, \qquad \beta_k^{2(\mathrm{sens})} = \frac{G_{\mathrm{T}} G_{\mathrm{R}} \lambda^2 \rho}{(4\pi)^3 \, d^2 d_k^2} \tag{82}$$

where G_{T} and G_{R} are the transmitting and receiving antenna gains, $\lambda = c/f_{\mathrm{c}}$ is the wavelength, c is the speed of light, and ρ is the RCS of the target. When $d > D_{\mathrm{elaa}}$, the channel coefficients can be considered almost equal across the array, leading to $\beta_{n,k} = \beta, \forall n, k$. For $d < D_{\mathrm{elaa}}$, instead, this approximation is no longer valid. Indeed, it is necessary to compute the path loss at each antenna element, because d_k can change significantly across the array aperture. Starting from (82) and using (80), the corresponding SNR at the k-th antenna can be written for the two cases as

$$\mathrm{SNR}_k^{(\mathrm{loc})} = \frac{P\, G_{\mathrm{T}} G_{\mathrm{R}} \lambda^2}{(4\pi)^2 \, d^2 \sigma^2} \frac{1}{\left(1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}\right)} = \frac{\mathrm{SNR}_0}{f_k} \tag{83}$$

$$\mathrm{SNR}_k^{(\mathrm{sens})} = \frac{P\, G_{\mathrm{T}} G_{\mathrm{R}} \lambda^2 \rho}{(4\pi)^3 \, d^4 \sigma^2} \frac{1}{\left(1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}\right)} = \frac{\mathrm{SNR}_0}{f_k} \tag{84}$$

where $\mathrm{SNR}_0 = P\beta_0^2/\sigma^2$ is the SNR at the reference antenna in the localization/sensing case and $f_k = 1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}$. Thus, $\beta_k^2 = \beta_0^2/f_k$ and from (83) and (84) the resulting SNR is the scaled version of that at the reference antenna. Moreover, once defined the SNR at the reference antenna, no differences are experienced between localization and sensing for the SNR scaling law across the array aperture. For what concerns the time displacement parameter across the array τ_k, for localization (a) we have $\tau_k = d_k/c$, while for sensing (b) $\tau_k = (d + d_k)/c$, considering the transmitter synchronized in time with the receiver. Notice that, for sensing (b), we can incorporate the phase shift due to the forward path between the transmitting antenna and the target within the common phase term φ. Therefore, hereafter we no longer make a distinction between localization and sensing by considering always $\tau_k = d_k/c = d\sqrt{f_k}/c$. Since the transmitted symbols are known by the receiver and used as pilots, we set $x_n = 1$. Thus, for the sake of performance limits derivation, we consider the received signal model

$$r_{n,k} = y_{n,k} + z_{n,k} = \frac{\alpha}{\sqrt{f_k}} \, e^{-\frac{\mathrm{j}2\pi f_{\mathrm{c}} d \sqrt{f_k(\mathbf{p})}}{c}} \, e^{\mathrm{j}\varphi} + z_{n,k} \tag{85}$$

where $\alpha = \sqrt{P}\beta_0$ and we have highlighted the dependence on $\mathbf{p} = [\theta, d]$ for the phase at each antenna.

Performance limits We now evaluate the performance limits for the estimation of the range d and AoA θ by exploiting the information coming from the phase shifts across the antenna elements in near-field conditions. To this end, we compute the CRLB for the set of parameters $\mathbf{\Theta} = \{d, \theta, \varphi, \alpha\}$ in (85). The (i, j) element of the Fisher Information Matrix (FIM) can be obtained as [45]

$$[\mathbf{J}]_{i,j} = \frac{2}{\sigma^2} \Re \left\{ \sum_n \sum_k \left[\frac{\partial y_{n,k}}{\partial \mathbf{\Theta}_i} \right]^* \left[\frac{\partial y_{n,k}}{\partial \mathbf{\Theta}_j} \right] \right\} \tag{86}$$

where $y_{n,k}$ is the noise-free version of $r_{n,k}$ in (85), leading to a 4×4 FIM in the form

$$\mathbf{J} = \begin{bmatrix} J_{dd} & J_{d\theta} & J_{d\varphi} & J_{d\alpha} \\ J_{\theta d} & J_{\theta\theta} & J_{\theta\varphi} & J_{\theta\alpha} \\ J_{\varphi d} & J_{\varphi\theta} & J_{\varphi\varphi} & J_{\varphi\alpha} \\ J_{\alpha d} & J_{\alpha\theta} & J_{\alpha\varphi} & J_{\alpha\alpha} \end{bmatrix} = \begin{bmatrix} \mathbf{J}_A & \mathbf{J}_B \\ \mathbf{J}_B^{\mathrm{T}} & \mathbf{J}_C \end{bmatrix} \tag{87}$$

since $J_{d\theta} = J_{\theta d}$, $J_{d\varphi} = J_{\varphi d}$, $J_{d\alpha} = J_{\alpha d}$, $J_{\varphi\theta} = J_{\theta\varphi}$, $J_{\alpha\theta} = J_{\theta\alpha}$, $J_{\alpha\varphi} = J_{\varphi\alpha}$. In particular, we have

$$\mathbf{J}_A = \begin{bmatrix} J_{dd} & J_{d\theta} \\ J_{\theta d} & J_{\theta\theta} \end{bmatrix}, \qquad \mathbf{J}_B = \begin{bmatrix} J_{d\varphi} & J_{d\alpha} \\ J_{\theta\varphi} & J_{\theta\alpha} \end{bmatrix}, \qquad \mathbf{J}_C = \begin{bmatrix} J_{\varphi\varphi} & J_{\varphi\alpha} \\ J_{\alpha\varphi} & J_{\alpha\alpha} \end{bmatrix}. \tag{88}$$

To obtain the inverse $\mathbf{J}^{-1}$ of the FIM and compute the CRLB for the range d and AoA θ, thanks to its block form, we exploit the matrix inversion lemma. In this manner, the inverse sub-matrix related to the parameters of interest (i.e., d and θ) is obtained, reducing the computation complexity. Specifically, we have

$$\mathbf{J}^{-1}_{1:2,1:2} = \mathbf{H}^{-1} \tag{89}$$

where $\mathbf{H} \triangleq \mathbf{J}_A - \mathbf{J}_B \mathbf{J}_C^{-1} \mathbf{J}_B^{\mathrm{T}}$, that allows to take into account matrices with lower dimensions. Thus, we have the CRLB for near-field range and AoA estimation given by, respectively:

$$\mathrm{CRLB}^{(\hat{d})} = \mathbf{J}^{-1}_{1,1}\,, \qquad \mathrm{CRLB}^{(\hat{\theta})} = \mathbf{J}^{-1}_{2,2}\,. \tag{90}$$

The received signal in (85) presents a phase term dependent on d_k (thus, d and θ) through τ_k, where $f_k(\mathbf{p}) = f_k(\theta, d)$ represents the wavefront curvature information [46]. To derive the performance bounds in the near-field region, two main assumptions are usually carried out [47]: (*i*) to simplify the strong nonlinearity of the phase profile in the near-field regime according to (85), the Fresnel approximation is exploited; (*ii*) the amplitude variations across the array are neglected, by considering a constant SNR for all the K receiving antennas, i.e., $\alpha_k = \alpha\ \forall k$. Thanks to this approximations, closed-form expressions can be derived for the CRLB for range and AoA estimation [48]:

$$\mathrm{CRLB}^{(\hat{d})} = \frac{6c^2 d^2 \left(\delta^2 \left(K^2-4\right)\sin^2\theta + 15d^2\right)}{\pi^2 f_c^2 NK\,\mathrm{SNR}_0 \cos^4\theta\,\delta^4 \left(K^2-4\right)\left(K^2-1\right)}, \tag{91}$$

$$\mathrm{CRLB}^{(\hat{\theta})} = \frac{3c^2}{2\pi^2 f_c^2 NK\,\mathrm{SNR}_0 \cos^2\theta\,\delta^2 (K^2-1)}\,. \tag{92}$$

However, when $d < D_{\mathrm{elaa}}$, which is denoted as *geometric near-field* region, the above conditions are no longer valid. In fact, such a boundary depends only on geometrical arguments and it is not related to the carrier frequency as it is for the definition of the Fraunhofer distance $r_{\mathrm{ff}} = 2D^2_{\mathrm{elaa}}/\lambda$. In this case, it is impossible to apply the Fresnel approximation, since its requirements are no longer fulfilled. In addition, the path loss at each antenna element must be considered separately. Therefore, in the following, we consider the original non-linear phase profile according to (85). In this case, we can compute the derivatives in (86) by obtaining[4]

$$J_{dd} = \frac{8\pi^2 f_c^2 N\,\mathrm{SNR}_0}{c^2} \sum_k \frac{\left(1 - \frac{x_k \sin\theta}{d}\right)^2}{\left(1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}\right)^2}\,, \tag{93}$$

$$J_{\theta\theta} = \frac{8\pi^2 f_c^2 N\,\mathrm{SNR}_0}{c^2} \sum_k \frac{(x_k \cos\theta)^2}{\left(1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}\right)^2}\,. \tag{94}$$

---

[4] The other FIM terms are omitted for brevity.

The previous derivatives are calculated considering the dependence on $\mathbf{p}$ of the phase profile only, according to (85). Thus, a performance bound is obtained for an estimator extracting location-dependent information solely from the phase term. For what concerns range estimation, when the array size becomes extremely large, we can compute the limit for the FIM term J_{dd} in (93) as $K \to \infty$. For simplicity, by assuming a UE/target on the broadside direction of the array (i.e., $\theta = 0$), we obtain for the right-hand term in (93) the symmetric series:

$$\sum_{k=-\infty}^{\infty} \frac{1}{\left(1 + a^2 k^2\right)^2} \tag{95}$$

with $a^2 = \frac{\delta^2}{d^2}$ which converges to $\xi = \frac{\pi a \coth(\pi/a) + \pi^2 \operatorname{csch}^2(\pi/a)}{2a^2}$. Therefore, the information on the range saturates as the array becomes extremely large when considering the actual dependence on the non-linear phase term and path loss. Consequently, the related CRLB for range estimation necessarily presents a horizontal asymptote for large K, and is lower bounded as

$$\mathrm{CRLB}^{(\hat{d})} \geq \frac{\lambda^2}{8\pi^2 N\, \mathrm{SNR}_0\, \xi} \approx \frac{\lambda^2 \delta}{4\pi^3 N\, \mathrm{SNR}_0\, d} \tag{96}$$

where the approximation holds for small a (e.g., a standard $\lambda/2$ spaced array and practical UE/target distance). The lower bound is obtained considering that the inverse FIM on range equals the CRLB when all the other parameters are already known, or they are uncorrelated. Thus, any asymptote derived for $1/J_{dd}$ is a lower bound for the actual CRLB, as the other parameters are unknown and the performance accuracy cannot be improved with respect to this case. As can be seen, saturation of the performance occurs for ELAAs, and the saturation value depends on the UE/target distance d.

For AoA estimation, instead, considering again $K \to \infty$ and $\theta = 0$, we obtain for the right-hand term in (94) the symmetric series:

$$\delta^2 \sum_{k=-\infty}^{\infty} \frac{k^2}{\left(1 + a^2 k^2\right)^2} \tag{97}$$

which converges to $\xi' = \frac{\pi \delta^2 \left(a \coth(\pi/a) - \pi \operatorname{csch}^2(\pi/a)\right)}{2a^4}$. Therefore, even the information on the AoA saturates as the array becomes extremely large when considering the actual dependence on the non-linear phase term and path loss. Consequently, the related CRLB for AoA estimation necessarily presents a horizontal asymptote for large K, and, by following the same reasoning as before, is lower bounded as

$$\mathrm{CRLB}^{(\hat{\theta})} \geq \frac{\lambda^2}{8\pi^2 N\, \mathrm{SNR}_0\, \xi'} \approx \frac{\lambda^2 \delta}{2\pi^3 N\, \mathrm{SNR}_0\, d^3} \tag{98}$$

where the approximation holds for small a. It is possible to notice that, differently from what predicted by the traditional expression, the accuracy limit does depend on the UE/target distance d.

If the same analysis is performed again considering the true phase profile (thus still avoiding the Fresnel approximation), but with a hypothetical constant path loss, we obtain that the performance for range estimation saturates, while the performance for AoA estimation continues to improve unboundly. This result highlights that saturation for range estimation mainly arises due to the fact that the true phase profile is considered, while for AoA estimation is caused by the non-uniform signal strength across the array. In fact, if the Fresnel approximation is employed, a non-linear phase profile is assumed for all the elements composing the ELAA. This is not true when $d < D_{\mathrm{elaa}}$, as the wavefront is almost planar for the array elements far from the array center. On the contrary, AoA estimation does not require a non-linear phase profile: indeed, it is feasible also in far field. In this case, saturation occurs since the array elements far from the array center do not collect significant energy and thus do not contribute to the estimation quality as far as the array becomes extremely large.

Numerical results Simulations are performed considering a half-wavelength spaced array ($\delta = \lambda/2$), with $f_c = 300$ GHz, an overall SNR at the reference antenna over the N subcarriers $N\,\mathrm{SNR}_0 = -20$ dB, $\theta = 0$, and two different UE/target distances $d = 10$ cm and $d = 1$ m. In Figure 16(a) the root-CRLB for range estimation

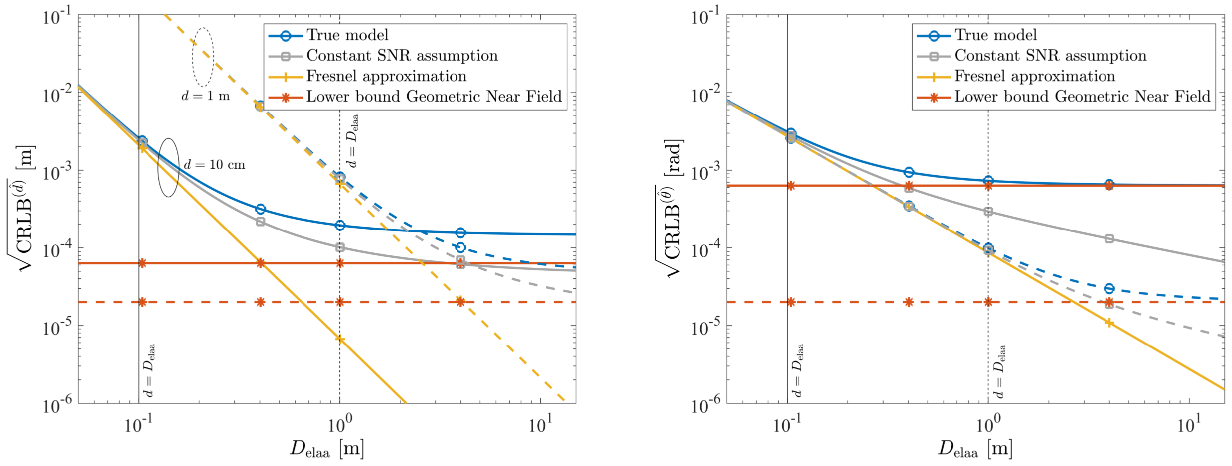

(a) Range estimation (b) AoA estimation

Figure 16: Root-CRLB for range (a) and AoA (b) estimation as a function of the array aperture D_{elaa}. UE/target distance of 10 cm (continuous lines) and 1 m (dashed lines).

is depicted as a function of the array aperture D_{elaa} (i.e., for increasing values of K). Continuous lines (–) are for $d = 10$ cm and dashed lines (- -) are for $d = 1$ m. The blue curves ($\circ$) refer to the actual CRLBs derived with the FIM terms obtained with the true phase profile and path loss model in the geometric near-field region; the gray curves ($\square$) refer to the CRLBs obtained with a hypothetical constant path loss, namely considering the original phase profile, but $\beta_k \approx \beta \ \forall k$; lastly, the yellow curves ($+$) show the traditional CRLBs obtained according to the Fresnel and constant path loss approximations. For reference, the red lines ($*$) are the lower bounds (asymptotes in (96)). It is possible to notice that traditional results based on the Fresnel approximation start to lose their validity when the array aperture becomes comparable to d, i.e., $d \approx D_{\mathrm{elaa}}$ (see the vertical black lines). Indeed, the performance predicted according to the Fresnel approximation improves unbounded as the aperture increases, since all the antennas consider a non-linear phase profile for the received signal. This is not true when an ELAA is employed, as only the signal received by the central portion of the array exhibits a strong nonlinearity for its phase distribution. In contrast, when the true phase profile is adopted, we observe a saturation of the performance. The gap between the actual CRLB (blue lines) and the asymptotes (red lines) comes from the other unknown parameters. When a constant path loss is considered, together with the true phase profile, saturation is still experienced, as the main effect on the range estimation performance is due to the use of the Fresnel approximation, rather than to the constant path loss assumption. Indeed, the non-linearity of the phase profile provides a significant improvement of the range estimation performance and it is not considered in this case, that is why saturation still occurs.

In Figure 16(b) results are reported for AoA estimation. In this case, when considering the classical CRLB obtained with the Fresnel approximation, the bound is independent of the distance d (yellow curves), and the accuracy improves unbounded as the array size increases. Differently, the bounds derived with the FIM terms obtained with the true phase profile and path loss model in the geometric near-field region, show a saturation as predicted by the asymptotes in (98). Moreover, it is shown that the CRLB for AoA estimation becomes dependent on the UE/target distance, differently from what would be expected in far-field conditions or even in near-field conditions when $d > D_{\mathrm{elaa}}$. Notice that for the case $\theta = 0$ we are considering, the CRLB and the inverse of the information coincide, so that the lower bound in (98) predicts exactly the saturation value of the CRLB.

4.4.2 Velocity Estimation using Extremely Large Antenna Arrays

Thanks to the adoption of ELAAs, the ability to observe targets from different directions permits the estimation of the different velocity components, as for multi-static radars, but with a monostatic configuration. In particular, it allows to estimate not only the radial component of the velocity, as is traditionally the case in the far field, but also the transverse one. This is the concept of *velocity sensing* recently introduced in [49] when employing ELAAs working in near-field conditions, enabling to get a full picture of the target velocity and trajectory using a monostatic radar instead of multi-static radars that require challenging synchronization and expensive hardware [50]. Such a capability makes ELAAs attractive for velocity estimation, in addition to the already-known benefits of near-field operations, such as improved spatial multiplexing even in LOS conditions [51, 52], the possibility of single-anchor localization [46] and high-resolution sensing [43, 52].

System model We consider sensing of a passive point target in a 2D scenario using a SIMO monostatic radar at the BS, arranged according to a linear ELAA deployment. The parameters describing the considered setup are the same as above in the previous subsection for the sensing case (b). We assume a point target placed in position $\tilde{\mathbf{p}} = [d \sin\theta, d\cos\theta]^\top$ moving with velocity $\mathbf{v} = [v_\mathrm{r}, v_\mathrm{t}]^\top$ tangent to its trajectory on the plane, where v_r and v_t are the radial and transverse velocity components, respectively, with respect to the center of the ELAA. For convenience, let us consider a polar coordinate system, so that the distance d_k between the k-th receiving antenna and the target can be written as in (80) . The central antenna element of the ELAAs, which is selected as reference for distance, angle and velocity estimation, transmits an OFDM signal spanning N subcarriers and M OFDM symbols. At the receiver side, after classical cyclic prefix removal and FFT processing, the received signal is

$$r_{m,n,k} = y_{m,n,k} + z_{m,n,k} = \sqrt{P}\beta_{m,n,k} x_{m,n} e^{-\mathrm{j}2\pi f_n \tau_k} e^{\mathrm{j}2\pi \nu_{n,k} m T_\mathrm{sym}} + z_{m,n,k} \tag{99}$$

where $T_\mathrm{sym} = T + T_\mathrm{cp}$ is the symbol time with T_cp is the cyclic prefix duration and $T = 1/\Delta f$, and the round-trip Doppler shift is represented as $\nu_{n,k} = \frac{f_n}{c}\left(v_\mathrm{r} + \mathbf{v}\cdot\mathbf{e}_k\right)$, with $\mathbf{a}\cdot\mathbf{b}$ indicating the scalar product between vectors $\mathbf{a}$ and $\mathbf{b}$. Here, the term $\mathbf{e}_k = \mathbf{d}_k/||\mathbf{d}_k||$ is a unit norm vector denoting the direction between the k-th receiving antenna element and the target, so that $\nu_{n,k}$ includes a component which is the projection of the target's velocity along this direction for each receiving antenna. In this case, we consider the channel gain equal for all the antennas, which is reasonable for the practical size of the receiving ELAA, in particular when $d > D_\mathrm{elaa}$ in LOS-dominated near-field scenarios, and small relative bandwidth $B = N\Delta f \ll f_\mathrm{c}$, so that we can write $\beta_{m,n,k} = \beta$, $\forall m, n, k$. Starting from the received signal (99), the BS aims at estimating the position $\mathbf{p}$ and velocity $\mathbf{v}$ through the estimation of their components $[d, \theta]$ and $[v_\mathrm{r}, v_\mathrm{t}]$, respectively, thanks to the use of the ELAA.

Performance limits The optimal way of processing the receiving signal would require the joint estimation of all the unknown parameters $\boldsymbol{\Theta} = \{d, \theta, v_\mathrm{r}, v_\mathrm{t}\}$ from the $M \times N \times K$ observations. This would imply a very high complexity, especially when operating in the near-field region where the planar wavefront approximation does not hold. To focus here on the estimation of the radial and transverse velocity components when adopting the ELAA, we assume that the distance d and the AoA θ have already been estimated, as done also in [49]. This can be realized by exploiting the information coming from the different subcarriers (e.g., for distance estimation), and/or from the phase profile along the array caused by the spherical wavefront (e.g., for distance and angle estimation). For this reason, τ_k is known and does not need to be estimated. Moreover, we consider $x_{m,n} = 1$, as the transmitted symbols are known at the receiver and used as pilots. Starting from this assumption, we compute the CRLB considering $\boldsymbol{\Theta} = \{v_\mathrm{r}, v_\mathrm{t}\}$. The (i, j)-th element of the FIM can be obtained as [45]

$$[\mathbf{J}]_{i,j} = \frac{2}{\sigma^2}\Re\left\{\sum_{m,n,k}\left[\frac{\partial y_{m,n,k}}{\partial \boldsymbol{\Theta}_i}\right]^*\left[\frac{\partial y_{m,n,k}}{\partial \boldsymbol{\Theta}_j}\right]\right\} \tag{100}$$

Thus we have a 2×2 FIM in the form

$$\mathbf{J} = \begin{bmatrix} J_{v_\mathrm{r} v_\mathrm{r}} & J_{v_\mathrm{r} v_\mathrm{t}} \\ J_{v_\mathrm{t} v_\mathrm{r}} & J_{v_\mathrm{t} v_\mathrm{t}} \end{bmatrix} \tag{101}$$

with $J_{v_r v_t} = J_{v_t v_r}$, and the corresponding CRLBs on the radial and transverse velocities are given by, respectively,

$$\mathrm{CRLB}^{(v_r)} = \frac{1}{\det \mathbf{J}} J_{v_t v_t}\,, \qquad \mathrm{CRLB}^{(v_t)} = \frac{1}{\det \mathbf{J}} J_{v_r v_r} \tag{102}$$

where $\det \mathbf{J} = J_{v_r v_r} J_{v_t v_t} - J^2_{v_t v_r}$. According to the geometry of the scenario, we can project the radial and transverse velocity components along the direction $\mathbf{e}_k$ by obtaining the projections $v_{\mathrm{r}k}$ and $v_{\mathrm{t}k}$, respectively. It holds

$$v_k = v_{\mathrm{r}k} + v_{\mathrm{t}k} \tag{103}$$

where v_k is the projection of the velocity $\mathbf{v}$ along the direction $\mathbf{e}_k$. Then, we can write

$$\nu_{n,k} = \frac{f_n}{c}(v_{\mathrm{r}} + v_k). \tag{104}$$

The projections $v_{\mathrm{r}k}$ and $v_{\mathrm{t}k}$ can be obtained as

$$v_{\mathrm{r}k} = q_k v_{\mathrm{r}} \tag{105}$$
$$v_{\mathrm{t}k} = p_k v_{\mathrm{t}} \tag{106}$$

where

$$q_k = \frac{d - x_k \sin\theta}{d_k} = \frac{1 - \frac{x_k \sin\theta}{d}}{\sqrt{1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}}} \tag{107}$$
$$p_k = \frac{x_k \cos\theta}{d_k} = \frac{x_k \cos\theta}{d\sqrt{1 + \frac{x_k^2}{d^2} - \frac{2x_k \sin\theta}{d}}}\,. \tag{108}$$

By substituting (105) and (106) in (103), and considering (104) in the model (99), we can compute the derivatives in (100) and obtain the components of the FIM

$$J_{v_r v_r} = \sum_n I_n \sum_k (1 + q_k)^2 \tag{109}$$

$$J_{v_t v_t} = \sum_n I_n \sum_k p_k^2 \tag{110}$$

$$J_{v_t v_r} = \sum_n I_n \sum_k p_k\,(1 + q_k) \tag{111}$$

where

$$I_n = \frac{2\pi^2 f_n^2\, M\, \mathrm{SNR}\,(M^2 - 1)\, T^2_{\mathrm{sym}}}{3c^2} \tag{112}$$

and we defined the SNR as $\mathrm{SNR} = P\beta^2/\sigma^2$. Considering the practical case $d > D_{\mathrm{elaa}}$, it holds $J_{v_r v_r} J_{v_t v_t} \gg J^2_{v_t v_r}$ so that $\det \mathbf{J} \approx J_{v_r v_r} J_{v_t v_t}$ and we have $\mathrm{CRLB}^{(v_r)} \approx 1/J_{v_r v_r}$ and $\mathrm{CRLB}^{(v_t)} \approx 1/J_{v_t v_t}$. This assumption is verified exploiting the 2nd order Taylor expansion for q_k in (107) and p_k in (108). In this way, considering $x_k^2/d^2 \to 0$ when $d > D_{\mathrm{elaa}}$ and $\cos^2\theta$ spans between 0 and 1, an upper bound is found such that the ratio between $J^2_{v_t v_r}$ and $J_{v_r v_r} J_{v_t v_t}$ results less than 0.08. This leads to the conclusion that the estimation of the radial and the transverse velocity components are practically decoupled in this configuration. Since $B \ll f_{\mathrm{c}}$, $f_n \approx f_{\mathrm{c}}$ and, as a consequence, $I_n \approx I$. When $d \to \infty$, the right-hand summation in (109) tends to $4K$ and, consequently, the inverse of the radial velocity information becomes the traditional CRLB in far field

$$\mathrm{CRLB}^{(v_r)} = \frac{3c^2}{8\pi^2 f_{\mathrm{c}}^2\, MNK\mathrm{SNR}\,(M^2 - 1)\, T^2_{\mathrm{sym}}} \tag{113}$$

where $(M^2-1)T_{\mathrm{sym}}^2$ is approximately the squared signal duration T_{obs}^2, and a MNK SNR gain is experienced due to the $M\times N\times K$ observations. When $\theta=0$, (109) becomes

$$J_{v_r v_r}=\frac{2\pi^2 f_c^2 MN\mathrm{SNR}(M^2-1)T_{\mathrm{sym}}^2}{3c^2}\sum_k\left(1+\frac{1}{\sqrt{1+k^2\frac{\delta^2}{d^2}}}\right)^2. \tag{114}$$

Interestingly, the information on the target's radial velocity reduces as the distance d decreases, since $k^2\delta^2/d^2>0$, resulting in an SNR gain lower than K for a small distance. This is reasonable since when the target approaches the array, the antennas farther from the array center sense the target in a direction different from that corresponding to the radial velocity, which is defined according to the reference central element. However, for practical array dimensions and operating distances (in particular, when the distance becomes higher than the largest x_k, i.e., than the array aperture D_{elaa}), the term $k^2\delta^2/d^2$ is small. Therefore, as far as the distance d exceeds D_{elaa}, we have $1/J_{v_r v_r}\to\mathrm{CRLB}^{(v_r)}$ in (113) which does not depend on the distance d and on the array aperture. Thus, no differences are experienced between far-field and near-field radial velocity estimation quality.

For what concerns the transverse velocity, its estimation at a large distance is not possible; specifically, we have $p_k\to 0$ for $d\to\infty$ (in practice, when $d\gg D_{\mathrm{elaa}}$) so that $J_{v_t v_t}\to 0$ and no information can be retrieved. In fact, when the target is far, the whole array sees it under a single direction, which is that of the radial velocity. Differently, when the array aperture increases and/or the distance decreases, we have $p_k\neq 0$, and transverse velocity estimation becomes feasible. Therefore, the possibility of gathering information on the transverse velocity can be attributed to the possibility of projecting the target velocity along the set of directions corresponding to the different antennas of the array. The information on transverse velocity is maximum on the boresight direction of the ELAA according to (110) and (108) (i.e., $\theta=0$); differently from the radial velocity, it is not possible to estimate the transverse velocity when $\theta=\pm\frac{\pi}{2}$. In fact, in this case, the direction $\mathbf{e}_k$ is the same for all the antennas and corresponds to the radial direction, so no further information becomes available rather than the radial velocity. For the case of transverse velocity estimation, we can write for $\theta=0$

$$J_{v_t v_t}=\frac{2\pi^2 f_c^2 MN\mathrm{SNR}\,(M^2-1)T_{\mathrm{sym}}^2\delta^2}{3c^2 d^2}\sum_k\frac{k^2}{1+k^2\frac{\delta^2}{d^2}} \tag{115}$$

which decreases with the square of the distance d for fixed SNR. Considering the practical case $d>D_{\mathrm{elaa}}$, which also allows to define the SNR according to (112), the term $k^2\delta^2/d^2$ is generally small, and the following approximate expression is obtained

$$\begin{aligned}J_{v_t v_t}&\approx\frac{\pi^2 f_c^2 MNK\mathrm{SNR}\,(M^2-1)T_{\mathrm{sym}}^2(K^2-1)\delta^2}{18c^2 d^2}\\&\approx\frac{\pi^2 f_c^2 MNK\mathrm{SNR}\,T_{\mathrm{obs}}^2 D_{\mathrm{elaa}}^2}{18c^2d^2}.\end{aligned} \tag{116}$$

The information on the transverse velocity is proportional to the squared signal duration, as for the radial velocity in (114). Larger carrier frequency f_c is beneficial for both radial and transverse velocity estimation. The adoption of a proper array aperture D_{elaa} can be traded to achieve the required transverse velocity accuracy depending on the operating distance d, while the SNR and the symbol time have the same impact on both the components. Expression (116) shows a deep difference between what happens in the near field for distance estimation and transverse velocity estimation. In fact, in near-field distance estimation, the CRLB is proportional to $d^2/f_c^2 D_{\mathrm{elaa}}^4$; as a result, in that case, the root-CRLB is inversely proportional to the Fraunhofer distance $d_{\mathrm{ff}}=2D_{\mathrm{elaa}}^2/\lambda$, with $\lambda=c/f_c$. Therefore, increasing the array aperture D_{elaa} (i.e., number of antennas, antenna spacing) and/or the carrier frequency is beneficial for the estimation quality, with a larger impact of the array aperture. Instead, in transverse velocity estimation using ELAAs, the parameters affecting the quality are the same according to (116), but they influence the bound all with the second power. Differences are evident when considering a half-wavelength spaced array ($\delta=\lambda/2$). In this case, the information on the transverse velocity is

$$J_{v_t v_t}\approx\frac{\pi^2 MNK\mathrm{SNR}\,(M^2-1)T_{\mathrm{sym}}^2(K^2-1)}{72d^2} \tag{117}$$

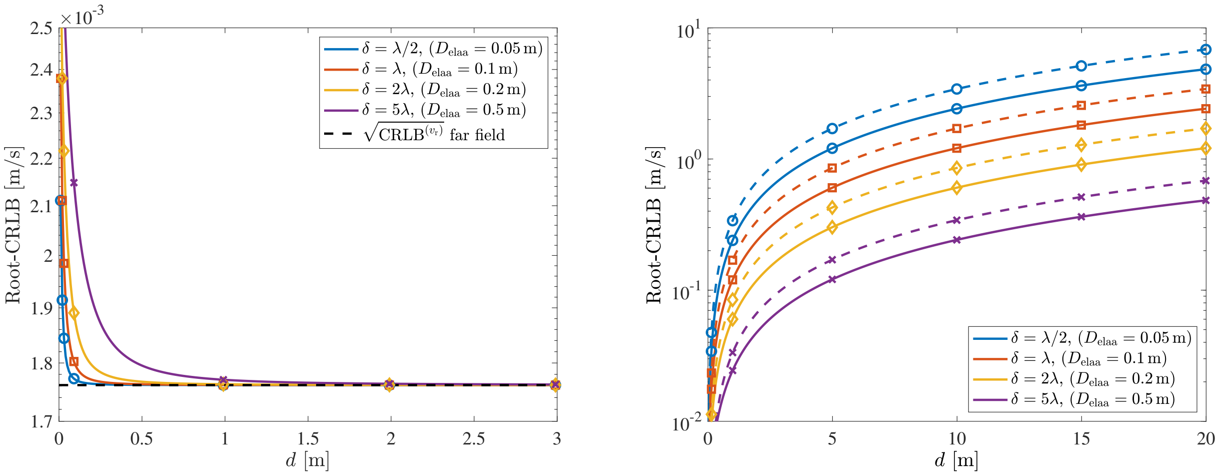

(a) Radial velocity estimation

(b) Transverse velocity estimation

Figure 17: Root-CRLB for radial velocity (a) and transverse velocity (b) estimation as a function of the distance d for different apertures D_{elaa}. Markers stands for the root-inverse information in the same setting.

which does not depend on the carrier frequency, differently from the radial velocity estimation or what happens for near-field distance estimation.

Numerical Results We present some results concerning the quality of velocity estimation using a linear ELAA. If not differently specified, we consider $f_c = 300$ GHz, $K = 101$ antennas, and $N\,\mathsf{SNR} = 10$ dB. A packet with $M = 14$ OFDM symbols is considered, and symbol time $T_{\mathsf{sym}} = 66.6\ \mu$s.

Figure 17(a) shows the root-CRLB for the radial velocity estimation in (102)-left as a function of the distance d, for a small distance from the array. Markers indicate the approximation obtained as $\sqrt{1/J_{\mathsf{v_r v_r}}}$, with $J_{\mathsf{v_r v_r}}$ given by (114), showing that it is very tight for the condition of interest. The traditional far-field CRLB is reported for comparison. Results are given for different apertures D_{elaa}, but with the same number of antennas K; this choice allows to consider a fixed overall received power, thus characterizing the impact of the array aperture under a constant number of observations MNK and thus SNR gain. As it is possible to notice, a small deviation from the far-field CRLB is experienced only at very small distance from the array, in particular, for a distance below the array aperture D_{elaa}, region also known as *geometric near field*. In Figure 17(b) the root-CRLB for the transverse velocity in (102)-right is reported as a function of the distance d, considering $\theta = 0$ and $\theta = 45°$. Again, the comparison between the exact root-CRLB and its approximation $\sqrt{1/J_{\mathsf{v_t v_t}}}$ (markers) shows a very good agreement. It can be noticed that the accuracy decreases as the distance d increases, as well as when the angle increases. The array aperture D_{elaa} has a fundamental role in achieving a good estimation quality.

Figure 18 compares the radial and transverse velocity accuracy in terms of root-CRLB as a function of the distance d for a fixed number of antennas K, $\delta = \lambda/2$, $\theta = 0$, and two different carrier frequencies, i.e., $f_c = 28$ GHz and $f_c = 300$ GHz. Thus, the two conditions translate in two different array apertures D_{elaa}. It is possible to notice that the radial accuracy is close to the traditional far-field CRLB, and it is always higher than the transverse one. In fact, by comparing (113) and (116) it can be seen that when $d = D_{\mathsf{elaa}}/(4\sqrt{3})$, the estimation quality for the radial velocity equals the estimation quality for the transverse one, at $\theta = 0$. Thus, an operating distance smaller than the array aperture D_{elaa} is required to achieve a balance in terms of estimation quality among the different components of the velocity. The transverse velocity accuracy does not depend on the carrier frequency, since we are comparing the results considering half-wavelength antenna spacing; differently, the estimation quality for the radial velocity improves as the carrier frequency increases.

To summarize the impact of the carrier frequency on the performance bounds for near-field distance, AoA,

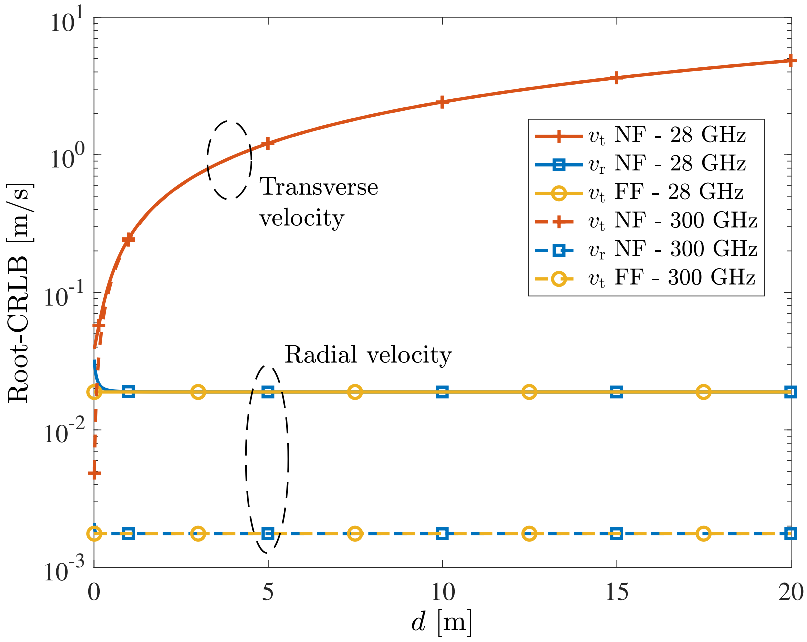

Figure 18: Root-CRLB for the radial and transverse velocities in near-field (NF) conditions. Comparison with the far-field (FF) CRLB in (113)*.*

and velocity estimation, Tab. 4 is shown. It is straightforward that, in general, increasing the carrier frequency is beneficial for the achievable performance in all the considered cases. However, when an half-wavelength array is considered, the CRLB for angle estimation and the CRLB for transverse component estimation are independent on the carrier frequency. In fact, the increase of the accuracy due to the carrier frequency is compensated by the smaller aperture.

Table 4: Comparison of the near-field CRLB for distance, AoA, and velocity estimation at $\theta = 0$*.*

	Distance	AoA	Radial Velocity	Transverse Velocity
Fixed geometry (K, δ, d, β)	$\propto \frac{1}{f_c^2}$	$\propto \frac{1}{f_c^2}$	$\propto \frac{1}{f_c^2}$	$\propto \frac{1}{f_c^2}$
Fixed signal (f_c, M, N, SNR)	$\propto \frac{d^4}{KD_{\text{elaa}}^4}$. More antennas in a fixed aperture: linear SNR gain only.	$\propto \frac{1}{KD_{\text{elaa}}^2}$. More antennas in a fixed aperture: linear SNR gain only.	$\propto \frac{1}{K}$. Linear SNR gain.	$\propto \frac{d^2}{KD_{\text{elaa}}^2}$. More antennas in a fixed aperture: linear SNR gain only.
Half-wavelength array ($\delta = \frac{\lambda}{2}$)	$\propto f_c^2$. The increase of the accuracy due to the carrier frequency through the term λ^2 cannot compensate the loss due to the smaller aperture.	Independent on f_c. The increase of the accuracy due to the carrier frequency through the term λ^2 is compensated by the smaller aperture (no effect).	$\propto \frac{1}{f_c^2}$. Accuracy improves with the frequency as in far-field Doppler-based velocity estimation.	Independent on f_c. The increase of the accuracy due to the carrier frequency through the term λ^2 is compensated by the smaller aperture (no effect).

4.5 Beamfocusing algorithms via trasmittive RIS for active multi-antenna feeder systems

In this section, we address the extremely large-scale multiple-input multiple-output (XL-MIMO) concept, which has emerged as a key enabler for THz communications, leveraging massive antenna deployments in the radiative near-field. However, fully active XL-MIMO arrays pose significant challenges in terms of hardware cost, complexity, and power consumption. An alternative architecture combining an active multi-antenna feeder (AMAF) with a RIS has been introduced to address these issues. This design exploits near-field electromagnetic (EM) space-feeding and programmable wavefront shaping to achieve scalable and energy-efficient wireless connectivity. The AMAF-RIS system reduces reliance on active RF chains, supports reconfigurable beamforming, and improves energy efficiency compared to conventional hybrid architectures [53]. Several works have investigated the propagation characteristics and beamforming capabilities of this architecture. In particular, near-field interactions between a RIS and uniform linear arrays (ULAs) or UPAs have been characterized in [54–58], with singular value decomposition (SVD)-based precoding schemes showing effective beam control over wide bandwidths. Nevertheless, scalability remains a concern, especially for supporting multi-UE transmission. Related studies [59, 60] have analyzed different RIS feeding strategies and low-complexity precoding schemes, highlighting trade-offs between performance and computational burden. Despite promising results, efficient multi-UE communications with this novel antenna architecture remains an open problem, particularly in designing beam patterns with low inter-UE interference. The work [61] reported in this section addresses these challenges by developing and evaluating effective near-field beamforming strategies for RIS-enabled AMAF architectures.

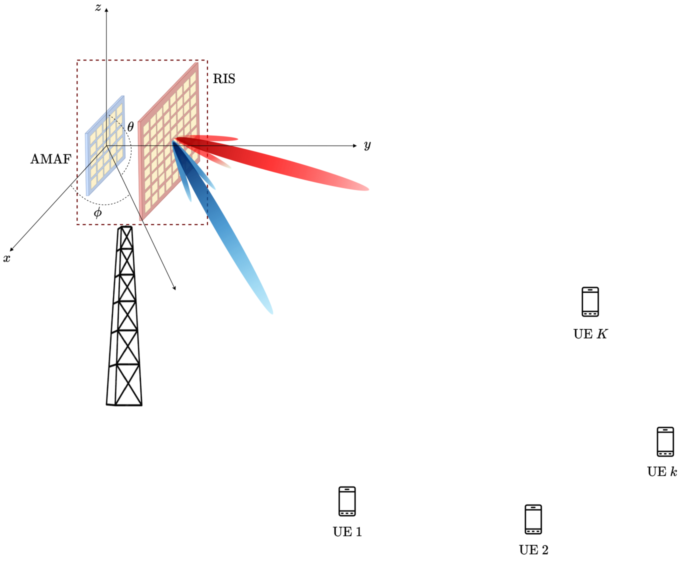

Figure 19: AMAF-RIS based antenna architecture for multi-UE MIMO communications in LOS conditions.

System Model We consider a transmitting BS equipped with an AMAF-RIS antenna array serving K UEs, as illustrated in Figure 19. The AMAF array is modeled as a square UPA comprising N_T active elements positioned at $\mathbf{p}_{\mathrm{T},i} = [x_{\mathrm{T},i},\, y_{\mathrm{T},i},\, z_{\mathrm{T},i}]^T$, with $i \in \mathcal{N}_\mathrm{T} = \{1, 2, \dots, N_\mathrm{T}\}$, and inter-element spacing of $\lambda/2$. This array illuminates a square RIS, made of N unit cells located at $\mathbf{p}_{\mathrm{RIS},j} = \left[x_{\mathrm{RIS},j},\, y_{\mathrm{RIS},j},\, z_{\mathrm{RIS},j}\right]^T$, for $j \in \mathcal{N} = \{1, 2, \dots, N\}$, also spaced by $\lambda/2$. Each UE is equipped with a single antenna at position $\mathbf{p}_\mathrm{R}^{(k)} = \left[x_\mathrm{R}^{(k)},\, y_\mathrm{R}^{(k)},\, z_\mathrm{R}^{(k)}\right]^T$, where $k \in \mathcal{K} = \{1, 2, \dots, K\}$.

The RIS is assumed to operate in the radiative near-field of the AMAF array, which is satisfied if

$$2D\sqrt{2N} \leq d \leq \frac{2\left(D\sqrt{2N}\right)^2}{\lambda} = d_{\mathrm{FF}} \,, \tag{118}$$

where d denotes the distance between the array and RIS centers, and D is the diagonal of the square RIS. This condition defines the Fresnel region of the transmit array, where both amplitude and phase variations across the spherical EM wavefronts must be accounted for in inter-element phase evaluations [62]. Conversely, the RIS-to-UE link may fall either in the near-field or far-field, depending on their relative distance d', which may be smaller or greater than d_{FF} [63].

We focus on the single-UE transmission scenario, in which the AMAF sends a symbol $x_k \in \mathbb{C}$ to the kth UE, with unit average power, i.e., $\mathbb{E}[|x_k|^2] = 1$. The received signal $y_k \in \mathbb{C}$ is expressed as

$$y_k = \mathbf{h}_k^T \mathbf{\Phi} \mathbf{G} \mathbf{b}_k x_k + w_k = \bar{\mathbf{h}}_k^T \mathbf{b}_k x_k + w_k \,, \; k \in \mathcal{K} \,, \tag{119}$$

where $\mathbf{b}_k = [b_1, b_2, \ldots, b_{N_{\mathrm{T}}}]^T \in \mathbb{C}^{N_{\mathrm{T}} \times 1}$ denotes the beamforming vector at the AMAF, and $\mathbf{G} = \{g_{j,i}\} \in \mathbb{C}^{N \times N_{\mathrm{T}}}$ represents the AMAF-to-RIS channel matrix. The matrix $\mathbf{\Phi} = \mathrm{diag}\left(\boldsymbol{\varphi}\right) = \mathrm{diag}\left(\varphi_1, \varphi_2, \ldots, \varphi_N\right) \in \mathbb{C}^{N \times N}$ entails the phase shifts applied by the RIS. We assume a transmissive, passive, and lossless RIS, meaning it receives energy from the AMAF side and re-radiates it towards the UEs with transformations dictated by $\mathbf{\Phi}$, and satisfies $\mathbf{\Phi}^H \mathbf{\Phi} = \mathbf{I}_N$. Hence, each coefficient $\varphi_n = e^{\jmath \alpha_n}$ has unit amplitude and phase $\alpha_n \in [0, 2\pi)$. The vector $\mathbf{h}_k^T = \left[h_1^{(k)}, \ldots, h_j^{(k)}, \ldots, h_N^{(k)}\right] \in \mathbb{C}^{1 \times N}$ models the RIS-to-kth UE channel, and $w_k \sim \mathcal{CN}\left(0, \sigma^2\right)$ is the additive white Gaussian noise. For compactness, we define the kth AMAF-UE effective channel as $\bar{\mathbf{h}}_k^T = \mathbf{h}_k^T \mathbf{\Phi} \mathbf{G} \in \mathbb{C}^{1 \times N_{\mathrm{T}}}$.

Under the assumption of free-space LOS propagation, the entries of $\mathbf{G}$ are given by

$$g_{j,i} = \frac{\lambda}{4\pi d_{i,j}} \sqrt{G_{\mathrm{T}}\left(\mathbf{\Theta}_{i,j}\right)} \, e^{-\jmath \frac{2\pi}{\lambda} d_{i,j}} \,, \tag{120}$$

where $[d_{i,j}, \mathbf{\Theta}_{i,j}] = [d_{i,j}, \phi_{i,j}, \theta_{i,j}]$ denote the distance and angular parameters between the ith transmit antenna and the jth RIS element, and $G_{\mathrm{T}}\left(\mathbf{\Theta}_{i,j}\right)$ is the AMAF transmit gain in the direction of departure $\mathbf{\Theta}_{i,j}$. Similarly, the RIS-to-kth UE channel coefficients are expressed as

$$h_j^{(k)} = \frac{\lambda}{4\pi d_j^{(k)}} \sqrt{G_{\mathrm{R}}\left(\mathbf{\Theta}_j^{(k)}\right)} \, e^{-\jmath \frac{2\pi}{\lambda} d_j^{(k)}} \,, \; k \in \mathcal{K} \,, \tag{121}$$

with $\left[d_j^{(k)}, \mathbf{\Theta}_j^{(k)}\right] = \left[d_j^{(k)}, \phi_j^{(k)}, \theta_j^{(k)}\right]$ indicating the link geometry from the jth RIS element to the kth UE, and $G_{\mathrm{R}}\left(\mathbf{\Theta}_j^{(k)}\right)$ representing the receive gain in the direction of arrival $\mathbf{\Theta}_j^{(k)}$. This signal model remains valid for both near-field and far-field regimes, as it incorporates exact distances and angles among all antenna pairs [64].

In addition, the SVD of the AMAF-RIS channel $\mathbf{G}$ is expressed as

$$\mathbf{G} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^H = \sum_{i=1}^{M} \xi_i \mathbf{u}_i \mathbf{v}_i^H , \tag{122}$$

where $\mathbf{\Lambda} = \mathrm{diag}(\xi_1, \ldots, \xi_M) \in \mathbb{C}^{N \times N_{\mathrm{T}}}$ contains the singular values of $\mathbf{G}$, with $M = \min(N_{\mathrm{T}}, N)$ and $\xi_i = \sigma_i(\mathbf{G}) \in \mathbb{R}_0^+$. The matrices $\mathbf{U} \in \mathbb{C}^{N \times N}$ and $\mathbf{V} \in \mathbb{C}^{N_{\mathrm{T}} \times N_{\mathrm{T}}}$ are unitary, with columns $\mathbf{u}_i$ and $\mathbf{v}_i$ representing the eigenvectors of $\mathbf{G}\mathbf{G}^H$ and $\mathbf{G}^H\mathbf{G}$, respectively.

In parallel, we define the matrix $\mathbf{H} = [\mathbf{h}_1, \ldots, \mathbf{h}_K]^T \in \mathbb{C}^{K \times N}$, aggregating the RIS –UE channels across the K receivers. Its SVD is given by

$$\mathbf{H} = \mathbf{P} \mathbf{\Sigma} \mathbf{Q}^H = \sum_{i=1}^{M'} \rho_i \mathbf{p}_i \mathbf{q}_i^H , \tag{123}$$

with $\mathbf{\Sigma} = \text{diag}(\rho_1, \dots, \rho_{M'}) \in \mathbb{C}^{K \times N}$, where $\rho_i = \sigma_i(\mathbf{H}) \in \mathbb{R}_0^+$ and $M' = \min(K, N)$. The matrices $\mathbf{P} \in \mathbb{C}^{K \times K}$ and $\mathbf{Q} \in \mathbb{C}^{N \times N}$ are unitary, and contain the eigenvectors of $\mathbf{H}\mathbf{H}^H$ and $\mathbf{H}^H\mathbf{H}$, respectively, with $\mathbf{p}_i$ and $\mathbf{q}_i$ indicating their ith columns.

We now consider the case of simultaneous transmission towards K independent UEs. Let $\mathbf{x} = [x_1, x_2, \dots, x_K]^T \in \mathbb{C}^{K \times 1}$ denote the transmit symbol vector, where each entry represents an independent data stream with unit power, i.e., $\mathbb{E}[\mathbf{x}\mathbf{x}^H] = \mathbf{I}_K$. The precoding matrix at the AMAF is denoted as $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_K] \in \mathbb{C}^{N_\text{T} \times K}$, subject to a total power constraint $\text{tr}(\mathbf{B}\mathbf{B}^H) = P_\text{T}$.

Under this configuration, the sum-rate in bits/s/Hz achievable by the system is given by [65]

$$\Gamma = \sum_{k=1}^{K} \gamma_k = \sum_{k=1}^{K} \log\left(1 + \frac{|\bar{\mathbf{h}}_k^T \mathbf{b}_k|^2}{\sigma^2 + \sum_{\substack{i=1 \\ i \neq k}}^{K} |\bar{\mathbf{h}}_k^T \mathbf{b}_i|^2}\right), \tag{124}$$

where γ_k denotes the rate for the kth user, while the denominator on the right hand side of the expression models the interference from the other streams plus the noise at that receiver. The analysis assumes no user cooperation at the receiver side.

Consequently, we aim to maximize the overall system throughput by jointly designing the precoding matrix at the AMAF and the RIS phase profile. This leads to the following constrained optimization problem

$$\underset{\mathbf{B}, \boldsymbol{\varphi}}{\text{maximize}} \quad \Gamma \tag{125}$$

$$\text{subject to} \quad \mathbf{\Phi} = \text{diag}\left(\boldsymbol{\varphi}\right), \tag{126}$$

$$|\varphi_i|^2 = 1, \quad i = 1, \dots, N, \tag{127}$$

$$\text{tr}\left(\mathbf{B}\mathbf{B}^H\right) = P_\text{T}. \tag{128}$$

This problem aims to jointly optimize the transmit beamformer $\mathbf{B}$ at the AMAF and the transmission coefficients $\boldsymbol{\varphi}$ of the RIS to maximize the overall sum-rate. Given the non-convex nature of (125)—stemming from the interdependence of Γ on both $\mathbf{B}$ and $\boldsymbol{\varphi}$—we refrain from pursuing an exact global solution. Instead, we propose a low-complexity heuristic approach yielding closed-form expressions for both variables. This method, leveraging peculiar system characteristics, can also serve as a suitable initialization for more advanced iterative procedures.

RIS Configuration Strategies for Sum-Rate Optimization We develop a tailored configuration strategy that exploits the near-field nature of the AMAF-RIS link while maintaining a practical diagonal structure for the RIS. Additionally, we consider the alternating optimization scheme in [66] using the obtained solution as initial configuration for benchmarking to evaluate the quality of the obtained solution.

For beamforming at the AMAF, [67] demonstrates that, in single-UE settings with near-field propagation, optimal mutual information is achieved by aligning the beamformer with the right singular vectors of the AMAF-RIS channel. Extending this principle, we define $\mathbf{B} = \mathbf{V}$, assigning each UE an orthogonal eigenvector, i.e., $\mathbf{b}_k = \mathbf{v}_k$ for $k \in \mathcal{K}$.

For the RIS design, we propose a diagonal approximation combining mode selection and focusing. In the first step, similar as in [57], we define the RIS transmission coefficients vector for a single user k according to

$$\boldsymbol{\psi}_k = \mathbf{u}_k^H \odot \mathbf{f}_k^{(\text{NF})}, \tag{129}$$

where $\odot$ is the element-wise multiplication operator and

$$\mathbf{f}_k^{(\text{NF})} = \left[e^{-j\kappa d_1^{(k)}}, \dots, e^{-j\kappa d_i^{(k)}}, \dots, e^{-j\kappa d_N^{(k)}}\right]^T \in \mathbb{C}^{N \times 1} \tag{130}$$

represents the near-field focusing vector towards the kth UE, where $\kappa = 2\pi/\lambda$ is the wavenumber, and with $d_i^{(k)} = \left|\mathbf{p}_{\text{RIS},i} - \mathbf{p}_\text{R}^{(k)}\right|$ denoting the distance between the ith RIS element and the kth UE. Then, these K vectors are used to construct the diagonal RIS transmission matrix as

$$\mathbf{\Phi} = \text{diag}\left(e^{j\cdot\angle\left(\sum_{k=1}^{K}\psi_k\right)}\right), \tag{131}$$

thus enforcing the phase-only constraint for the RIS coefficients.

As an alternative to near-field beam focusing, we consider an minimum mean-square error (MMSE) precoding approach, widely adopted in MIMO systems [68]. The precoding matrix $\mathbf{L} = [\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_K] \in \mathbb{C}^{N\times K}$ is computed as

$$\mathbf{L} = \eta\left(\sigma_{\text{TX}}^2\mathbf{I}_N + \mathbf{H}^H\mathbf{H}\right)^{-1}\mathbf{H}^H, \tag{132}$$

where $\sigma_{\text{TX}}^2 = \delta\left(||\mathbf{H}||_{\text{F}}^2/N\right)$ introduces a regularization term modulated by a design parameter $\delta > 0$, and $\eta = 1/\left|\left|\mathbf{L}^{(\text{MMSE})}\right|\right|_{\text{F}}$ensures power normalization.

Building upon this precoding structure, the RIS phase profile is constructed similarly to the focusing-based method. For each user k, the transmission coefficients vector is defined as

$$\boldsymbol{\psi}_k = \mathbf{u}_k^H \odot \boldsymbol{\ell}_k, \quad \forall k \in \mathcal{K}, \tag{133}$$

which incorporates the MMSE precoding direction $\boldsymbol{\ell}_k$for calculating the per-user weights instead of the focusing vector $\mathbf{f}_k^{(\text{NF})}$ from above. The final RIS configuration is then computed using the same procedure as in (131), by extracting the phase of the aggregate signal over all users.

Numerical Results The system operates at a carrier frequency $f_{\text{c}} = 120\,\text{GHz}$ ($\lambda \simeq 0.25\,\text{cm}$) with a bandwidth of $B = 120\,\text{MHz}$. The carrier frequency has been set at the lower end of the sub-THz range, as this choice enables a general analysis of near-field propagation effects in the considered scenario while ensuring parameter values of practical magnitude for the RIS configuration, thus keeping the complexity in balance. At higher frequencies, the RIS dimensions would need to be scaled accordingly to preserve the same near-field region. The system bandwidth has been selected such that, given the assumed transmit power budget, the resulting SNR guarantees a sufficient link margin to ensure that the orthogonal near-field channels exhibit adequate coupling strength. The transmit power is fixed at $P_{\text{T}} = 10\,\text{dBm}$, with a noise power spectral density of $\sigma^2 = -170\,\text{dBm/Hz}$ and a noise figure of $3\,\text{dB}$. The AMAF comprises $N_{\text{T}} = 16$ elements arranged as a $(\lambda/2)$-spaced UPA, while the RIS is positioned 8λ from the AMAF along the y-axis. The RIS features $N = 2500$ elements ($25\lambda \times 25\lambda$, with $\lambda/2$ spacing) and is deployed at a height of $z = 10\,\text{m}$ without tilt. Two single-antenna UEs ($K = 2$) located in the near field are considered: one fixed at $\mathbf{\Theta}^{(1)} = [30°, 135°]$ at a distance $d^{(1)} = 14\,\text{m}$, relative to the RIS center, and another moving along a $10\,\text{m}$-radius circular trajectory in the xy plane. In addition, the directional antenna gain functions $G_{\text{T}}(\mathbf{\Theta})$ and $G_{\text{R}}(\mathbf{\Theta})$, respectively associated with the AMAF and UEs antenna elements, are assumed to be identical for all array elements and are given by $G_{\text{T}}(\mathbf{\Theta}) = G_{\text{R}}(\mathbf{\Theta}) = 2\sin(\theta)\sin(\phi)$. The AMAF beamforming matrix is initialized as $\mathbf{B} = \mathbf{V}$ for all considered configurations. We then compare two different RIS design strategies: the focusing-based (FOC) approach and the MMSE-based (MMSE) approach (with $\delta/N = 0.2$). Specifically, each UE is assigned a RIS transmission coefficient vector computed according to (129) in the FOC case, and according to (133) in the MMSE case.

For benchmarking and evaluating our engineered solution, an alternating optimization procedure is applied using our solution as initial configuration, where the AMAF beamforming matrix $\mathbf{B}$ is first optimized via the weighted minimum mean-square error (WMMSE) approach to yield $\mathbf{B}^*$, followed by gradient-based optimization of the RIS configuration matrix $\mathbf{\Phi}$ to obtain $\mathbf{\Phi}^*$. This iterative optimization process, described in detail in [66], gradually enhances the system sum-rate by alternately improving the active and passive beamforming components until convergence to a local optimum is achieved.

Figure 20 compares the per-user rate γ_k, $k \in \mathcal{K} = 1, 2$ (shown in Figure 20(a)), and the overall system sum-rate Γ (shown in Figure 20(b)) as functions of the azimuthal separation $\Delta\varphi$ between the two UEs, which move along a circular trajectory. Solid lines refer to the optimized configuration derived via the algorithm in [66], while dashed lines represents our engineered phase configuration from (131).

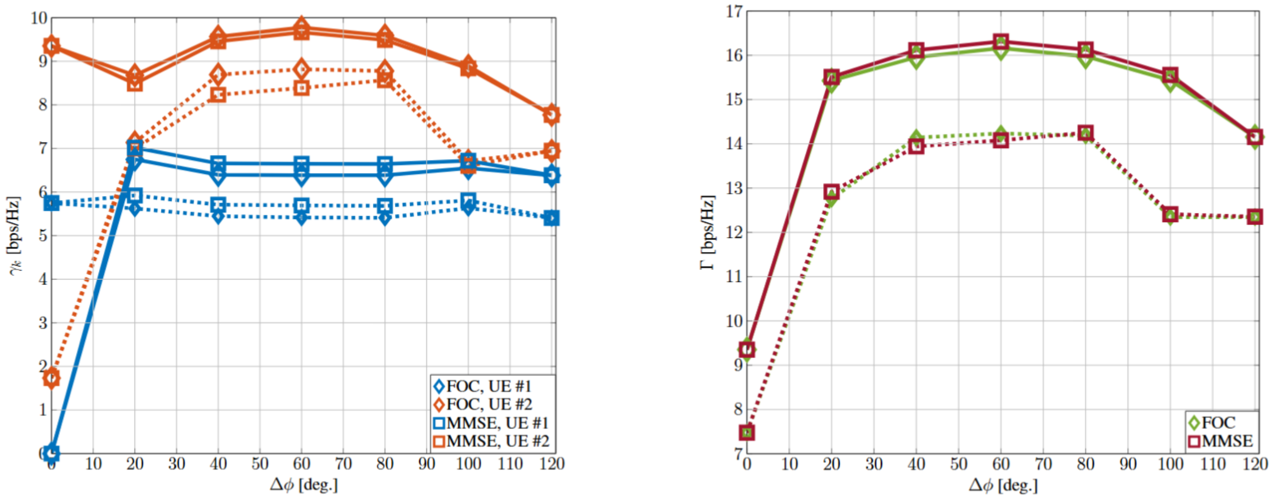

(a) Per-UE rate vs. azimuthal spacing

(b) Sum rate vs. azimuthal spacing

Figure 20: Comparison between (a) the per-UE rate γ_k, $k = 1, 2$, and (b) the system sum rate Γ as functions of the azimuthal spacing between UEs along a circular trajectory, under the assumption that the RIS applies only phase shifts. Solid lines $(-)$ represent the optimized configuration, while dashed lines $(- -)$ correspond to the engineered strategy for $\boldsymbol{\Phi}$ matrix configuration without optimization.

In Figure 20(a), the individual per-UE rates exhibit a clear dependence on the angular separation. When $\Delta\varphi$ is small, i.e., the UEs are closely spaced, the inter-user interference cannot be fully mitigated, resulting in asymmetric rate performance: typically, only one of the two users benefits from favorable channel conditions. As $\Delta\varphi$ increases, both users experience improved performance, particularly when one UE aligns with the boresight of the RIS. In this configuration, the transmit beam pattern of the RIS provides optimal power focusing, which is especially evident under the optimized scheme. However, the performance gap between the optimized and the heuristic configurations remains limited. This indicates that despite its simplicity, the proposed engineered design captures the optimal solution's essential structure and performs comparably well. Moreover, the difference between the FOC and MMSE-based combining schemes is negligible, implying that our heuristic design enables close-to-optimal performance without requiring complex precoding schemes. Notably, MMSE precoding demands accurate Channel State Information (CSI), including knowledge of the channel between each UE and the RIS, which introduces significant overhead in both estimation and feedback. In contrast, our engineered scheme eliminates the need for such procedures, rendering it more practical for real-time or large-scale implementations.

Figure 20(b) further supports these observations by illustrating the system sum-rate Γ. The trends mirror those seen in the per-user case: the heuristic configuration approaches the performance of the optimized one across the entire range of $\Delta\varphi$, with a performance gap of less than 20% . Additionally, both FOC and MMSE combining strategies yield almost overlapping curves, reinforcing the effectiveness of the engineered configuration regardless of the precoding technique employed. This result is particularly relevant when considering large-scale RIS deployments or scenarios with many users, where the computational burden of optimization and channel estimation becomes prohibitive. Interestingly, these observations hold across both near-field and far-field propagation regimes. As such, the proposed scheme demonstrates strong scalability, making it suitable for diverse deployment scenarios and frequency bands, including those envisioned for future sixth-generation (6G) systems.

In conclusion, the results highlight a fundamental trade-off between achievable performance and implementation complexity. The proposed engineered RIS configuration emerges as a compelling solution due to its low computational cost and negligible performance loss. Despite its reduced complexity, the proposed focusing-based strategy achieves performance levels close to those attained by computationally demanding optimization algorithms. These findings validate the practical relevance of the proposed approach, which enables efficient and scalable multi-user support in multi-UE MIMO systems without the need for full channel knowledge or real-time optimization.

4.6 Simulation-based evaluation of RIS placement and achievable data rates

In this section, we present a comprehensive simulation-based analysis of the impact of RIS placement on the wireless propagation environment and the resulting achievable data rates in an industrial setting. The evaluation is carried out at 285 GHz, leveraging the ray tracing capabilities of the SiMoNe to capture the complex electromagnetic interactions within the Robopac factory hall, where POC1 will be showcased. Indeed, WP2 has defined the configuration of POC1 within Deliverable D2.3 [69], and therefore this activity of WP4 will consider the same POC scenario, that is, the Robopac industry plant. High-gain antennas with narrow beam patterns are deployed at the TX and RX ends to establish a LoS link enhanced by the RIS, which reflects the incident signal to the receiver.

The pathloss predictions derived from the ray tracing simulations serve as a basis to determine possible RX placement and to assess the influence of the RIS RCS on system performance. These results will be used as a starting point for the definition of the layout that will be used in the TIMES POC, and to obtain an estimate of the achievable performance. The underlying calculation of the pathloss in the SiMoNe ray tracer, including antenna gains and RIS response, is detailed in Section 3.2.4. Subsequently, the achievable data rates are computed by mapping the pathloss values to link-level metrics under realistic assumptions of modulation and coding schemes. This approach allows us to quantify the potential throughput gains enabled by RIS-assisted communication in industrial scenarios.

Figure 21 illustrates the layout of the scenario within the Robopac factory, where the transmitter is located in the top right, the RIS is placed centrally at the bottom, and the receiver is positioned in the top left corner behind the machinery. This configuration serves as the basis to investigate the feasibility of wireless communication between industrial machines and central servers in realistic factory environments.

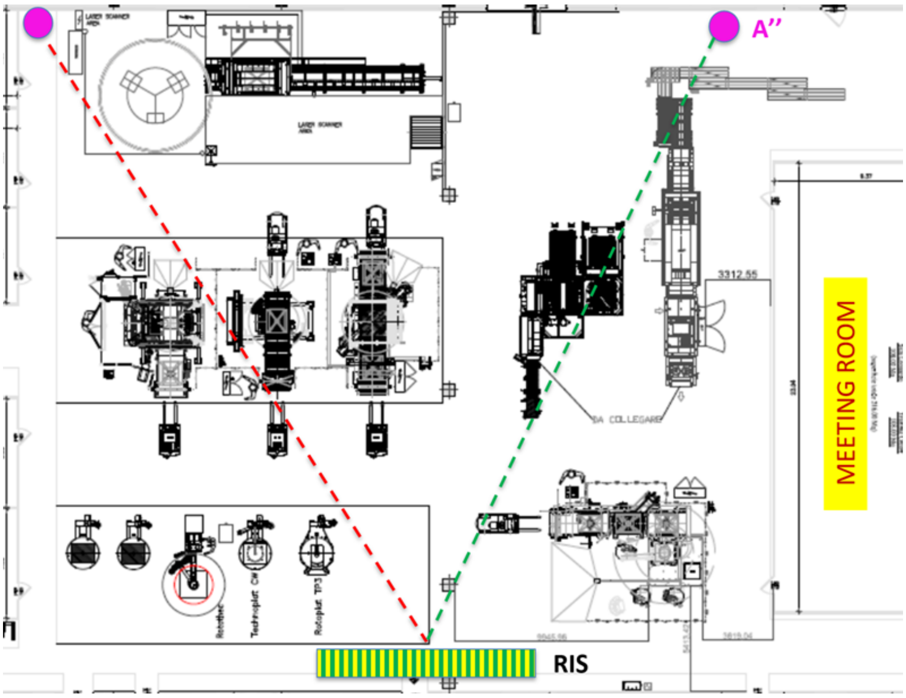

Figure 21: Layout of the Scenario in the Robopac Factory

4.6.1 Radio map prediction in the factory

The digital model of the Robopac factory hall described in [18] is used for conducting ray tracing simulations. The simulation utilizes high-gain directive TX and RX antennas as well as a reflective RIS which emulate the custom devices developed within the TIMES project that will be used in the POC activities. Their behavior is described by electromagnetic patterns provided by Anteral, as shown in Figures 22 and 23. The employed antennas exhibit a narrow Half Power Beam Width (HPBW) of only $1°$, ensuring a highly focused radiation in the main lobe direction. The side lobes are located at approximately $\phi = \pm 60°$ and are suppressed by roughly $40\,\mathrm{dB}$ relative to the main lobe, significantly reducing unwanted signal dispersion.

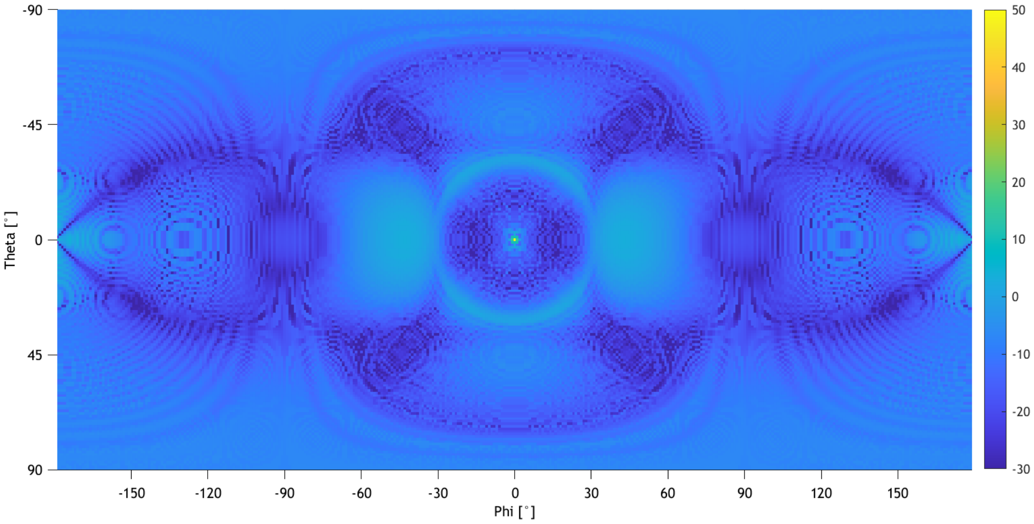

Figure 22: Antenna Pattern of the High Gain Antennas. Main Lobe at $\theta = 0°$ and $\phi = 0°$

The RIS pattern is defined for an incident wave arriving at an azimuth angle of $\phi = 0°$ and an elevation angle of $\theta = 0°$. Under this illumination condition, the maximum RCS of the RIS is measured to be $-77.34\,\mathrm{dBsm}$, with the reflection peak observed at $\phi = 32°$ and $\theta = 0°$. In contrast, the specular reflection at $\phi = 0°$ and $\theta = 0°$ forms a side lobe approximately $18\,\mathrm{dBsm}$ below the peak. This pattern underlines the highly directional response of the RIS and its suitability for spatially controlled wave redirection in indoor environments.

To evaluate the feasibility and effectiveness of RIS-assisted communication in the factory hall, map-based predictions were carried out using the SiMoNe ray-tracing framework. These simulations aim to determine viable placement configurations for TX, RX, and RIS units within the structural and operational constraints of the Robopac facility. For this Proof-of-Concept (PoC), all antennas and the RIS are mounted at a uniform height of $5.05\,\mathrm{m}$, which is above the height of the industrial machinery yet below the ceiling-mounted lighting infrastructure. As a result, all angular alignments can be considered in the azimuth plane only, simplifying the simulation setup.

Since the RIS pattern is only valid for an incident angle of $\Phi = 0°$ and features a maximum reflection toward $\phi = 32°$, the geometric configuration formed by the TX, RIS, and RX must result in a deflection angle of $32°$. Given that the TX is to be installed adjacent to the right-side wall and the RIS on the opposite wall, both components must be positioned such that the $32°$ reflection angle allows coverage of the RX region located within the machinery area. Additionally, structural obstacles such as the vertical support columns—marked as crosses in the scenario layout in Figure 21—must be considered when determining viable placements.

Two candidate positions were defined for the TX. The first one is located closer to the wall at ($x = 32.95\,\mathrm{m}$, $y = 31.14\,\mathrm{m}$), while the second is further away at ($x = 38.00\,\mathrm{m}$, $y = 31.14\,\mathrm{m}$). For the first TX position, two possible RIS placements were tested to direct the beam between structural columns: one with the TX aligned between columns 2 and 3, resulting in a RIS position at ($x = 16.53\,\mathrm{m}$, $y = 0.358\,\mathrm{m}$), and another between columns 3 and 4, placing the RIS at ($x = 22.41\,\mathrm{m}$, $y = 0.358\,\mathrm{m}$). These are referred to as Position 1 and Position 2, respectively. For the second TX placement, alignment between columns 3 and 4 was necessary to maintain visibility of the

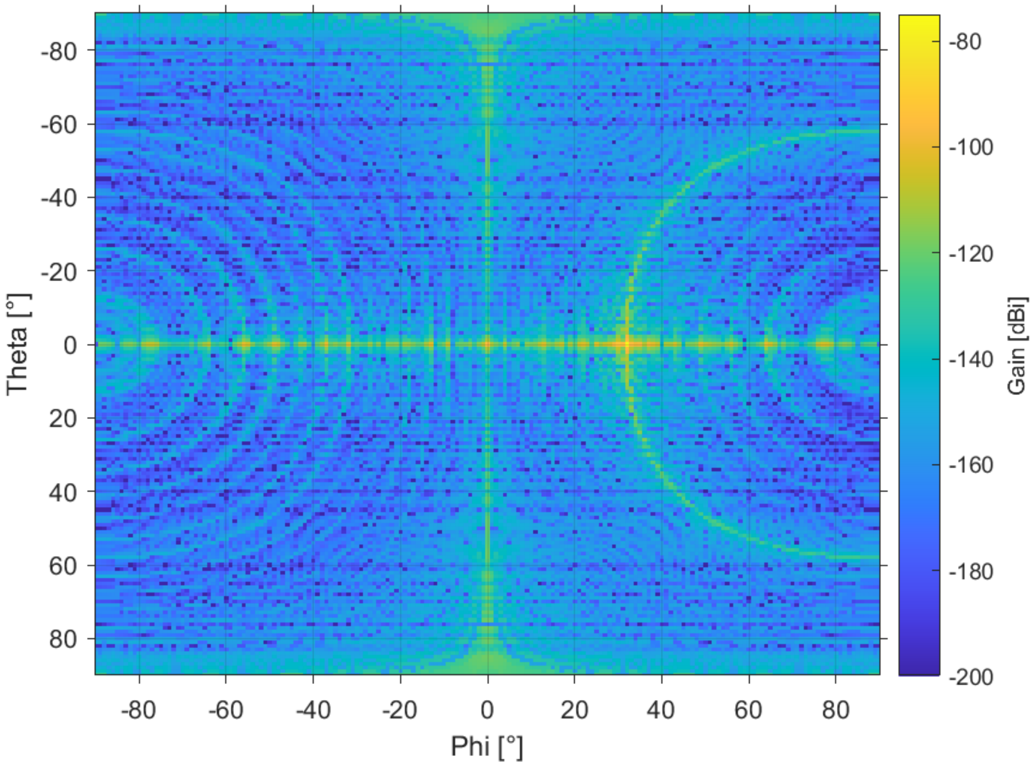

Figure 23: Pattern of the RIS. Main Lobe at $\theta = 0^\circ$ and $\phi = 32^\circ$

RIS toward the machine area. The corresponding RIS location was set to ($x = 19.28\,\text{m}$, $y = 0.358\,\text{m}$), denoted as Position 3. Note that the y-coordinate of all RIS placements is constrained by the fixed distance from the wall due to the mounting system.

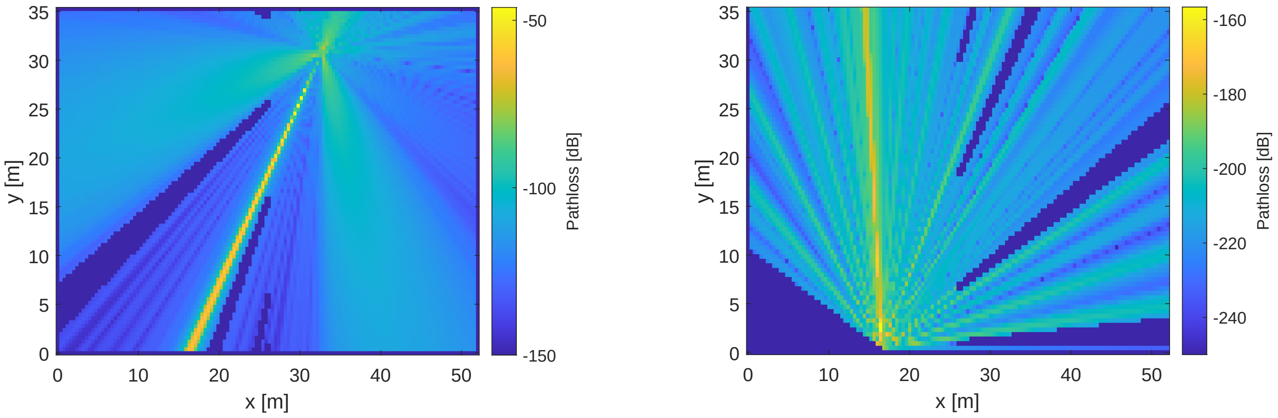

(a) Combined received power from TX and RIS

(b) Received power via RIS only

Figure 24: Map-based prediction of received power across the factory floor for Position 1

Due to the high directivity of both the high-gain TX antenna and the RIS, the ray-tracing simulation for the map-based prediction was conducted without considering multipath reflections. Figure 24 shows the resulting received power across the entire Robopac factory for Position 1. In Subfigure (a), the total received power is depicted, including contributions from both the direct path via the TX and the reflected path via the RIS. However, the simulated radar cross section (RCS) of the RIS is significantly underestimated, as only a subsection of the surface was modeled. While the overall shape of the RIS radiation pattern is expected to remain consistent, the absolute gain would be higher if the full surface were included. As a result of this limited RIS gain, no discernible influence of the RIS can be observed in Subfigure (a).

To isolate the contribution of the RIS, Subfigure (b) shows the received power exclusively from rays reflected by the RIS. In both subfigures, a clearly defined main lobe is visible, demonstrating that specular propagation dominates and reflections have minimal impact on the received power in this scenario. Additionally, shadowing effects from the structural columns are evident in both visualizations. Due to the narrow HPBW of the antenna and the RIS, the immediate vicinity around each element is not fully illuminated, which causes the main lobe to appear at a certain distance from each source.

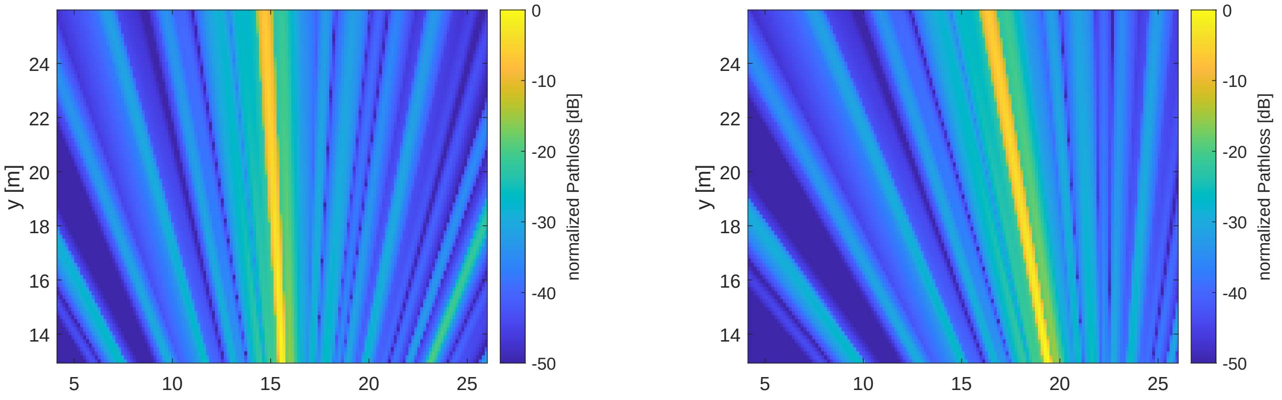

(a) Position 1 (b) Position 2

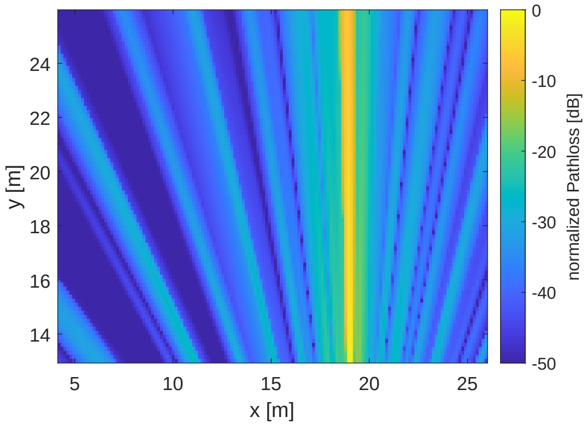

(c) Position 3

Figure 25: Map-based prediction of received power via the RIS in the machine area for all three positions

Since the absolute RIS gain in the simulations is underestimated due to the limited simulated surface area, the predicted path losses are significantly overestimated. To enable a meaningful comparison between the three investigated configurations, the received power values in all map-based predictions within the machine area were normalized to the highest received power value observed across the scenarios.

The results shown in Figure 25 indicate that all three TX-RIS configurations provide comparable performance within the machine area, with only minor differences of a few decibels. Consequently, the following section will evaluate the achievable data rates for each configuration in order to assess their relative effectiveness more quantitatively.

4.6.2 Evaluation of RIS-Assisted Throughput in an Industrial Setting

In order to evaluate the influence of different RIS placements on the achievable throughput, three configurations of TX and RIS are considered, as defined in the previous subsection. For each configuration, three distinct RX positions are selected to represent different regions of the factory floor: in front of the machines ($y = 14.5\,\mathrm{m}$), on top of the machines ($y = 19\,\mathrm{m}$), and behind the machines ($y = 24.5\,\mathrm{m}$). Due to the constraint that the angle between the incoming and outgoing ray at the RIS must be 32°, the corresponding x-coordinates of the RX for each TX-RIS configuration are uniquely determined. The resulting RX coordinates for each setup are summarized in Table 5.

Configuration	**RX Position 1 (front)**	**RX Position 2 (top)**	**RX Position 3 (rear)**
Position 1	$(15.56,\ 14.5)$	$(15.25,\ 19)$	$(14.88,\ 24.5)$
Position 2	$(19.12,\ 14.5)$	$(18.07,\ 19)$	$(16.79,\ 24.5)$
Position 3	$(19.11,\ 14.5)$	$(19.06,\ 19)$	$(18.99,\ 24.5)$

Table 5: RX coordinates for each TX-RIS configuration based on the 32° *angular constraint.*

To estimate the achievable data rates, the simulated pathloss values are compared to the reference receiver sensitivity thresholds defined in the IEEE 802.15.3-2023 standard [40]. Table 6 lists the sensitivity levels for selected Modulation and Coding Schemes (MCSs), considering only BPSK, QPSK, 16-QAM, and 64-QAM as relevant for this evaluation. The values depend on the employed bandwidth and define the minimum required receive power for successful decoding.

MCS ID	**Mod.**	**FEC Rate**	**2.16 GHz**	**4.32 GHz**	**8.64 GHz**	**12.96 GHz**	**17.28 GHz**	**25.92 GHz**	**34.56 GHz**	**69.12 GHz**
0	BPSK	11/15	-67	-63	-60	-57	-55	-54	-52	-49
1	BPSK	14/15	-64	-62	-59	-56	-53	-52	-50	-47
2	QPSK	11/15	-64	-60	-57	-55	-51	-50	-48	-45
3	QPSK	14/15	-61	-59	-56	-53	-50	-48	-46	-43
8	16-QAM	11/15	-56	-52	-48	-46	-43	-41	-39	-36
9	16-QAM	14/15	-52	-48	-45	-42	-39	-37	-35	-32
10	64-QAM	11/15	-51	-45	-42	-40	-36	-35	-33	-30
11	64-QAM	14/15	-46	-43	-40	-37	-34	-33	-31	-29

Table 6: Reference sensitivity levels (in dBm) for selected MCS defined in the IEEE 802.15.3-2023 [40].

For the evaluation of achievable data rates, only a bandwidth of 2.16 GHz was considered, since the RIS exhibits an approximately constant RCS only over a bandwidth slightly exceeding 2 GHz. If a higher bandwidth is to be employed, an RCS corresponding to the respective sensitivity levels in Table 6 must be achieved.

Furthermore, a transmit power of 0 dBm was assumed in the evaluation. If a different transmit power is used, the required RCS of the RIS must be adjusted accordingly—higher transmit power allows for lower RCS, and vice versa. The exact influence of the RCS on the received signal strength is discussed in Section 3.2.4. A 1 dB increase in RCS leads to a 1 dB increase in path gain.

The graphs in Fig. 26 show the achievable data rates as a function of the RCS for three different receiver positions. The simulations were carried out for RCS values between -30 dBsm and 6 dBsm. Values below

MCS ID	Modulation	FEC rate	Data rate in [GBps]
0	BPSK	11/15	1.29
1	BPSK	14/15	1.64
2	QPSK	11/15	2.58
3	QPSK	14/15	3.29
8	16-QAM	11/15	5.16
9	16-QAM	14/15	6.45
10	64-QAM	11/15	7.74
11	64-QAM	14/15	9.86

Table 7: Achievable data rates in Gbps for selected MCS at 2.16 GHz bandwidth, based on IEEE 802.15.3-2023 [40].

-30 dBsm were excluded because no data transmission was possible in this regime. Similarly, values above 6 dBsm were not simulated, as the maximum achievable data rate was already reached and no further gains were expected beyond this point.

The three subfigures represent different receiver positions, while the color of each curve indicates the receiver's relative placement with respect to the machine. The blue curve corresponds to receiver positions in front of the machine, the red curve represents positions directly on the machine, and the yellow curve indicates positions behind the machine.

As expected, the highest data rates are achieved for receiver positions in front of the machines (blue curves), due to the shorter propagation distance and consequently lower path loss. Starting from an RCS of -18 dBsm, all three receiver positions enable data transmission, regardless of the specific configuration of the TX and the RIS. A significant increase in data rate can be observed between -12 dBsm and -6 dBsm, which corresponds to the transition point at which the use of 16-QAM becomes feasible. Once the RCS exceeds 0 dBsm, all combinations of TX, RIS, and receiver allow for transmission with 64-QAM and a Forward Error Correction (FEC) code rate of $14/15$, resulting in the maximum achievable data rate for the simulated bandwidth of 2.16 GHz.

These results demonstrate that, assuming sufficiently high RCS values at the RIS, reliable communication links over distances exceeding 50 m are feasible when employing highly directive antennas and reflective surfaces. However, it is important to note that certain impairments—such as beamsquint—were only implicitly considered in the analysis through the choice of a relatively narrow bandwidth. If comparable RCS values can be maintained over larger bandwidths, even higher data rates may be achievable.

These considerations will directly inform the design of the POC1 to be implemented in the Robopac factory hall. In particular, if an RCS of 0 dBsm can be achieved, the KPIs defined in [69] will be fulfilled, enabling a theoretical throughput of up to 9.86 Gbps.

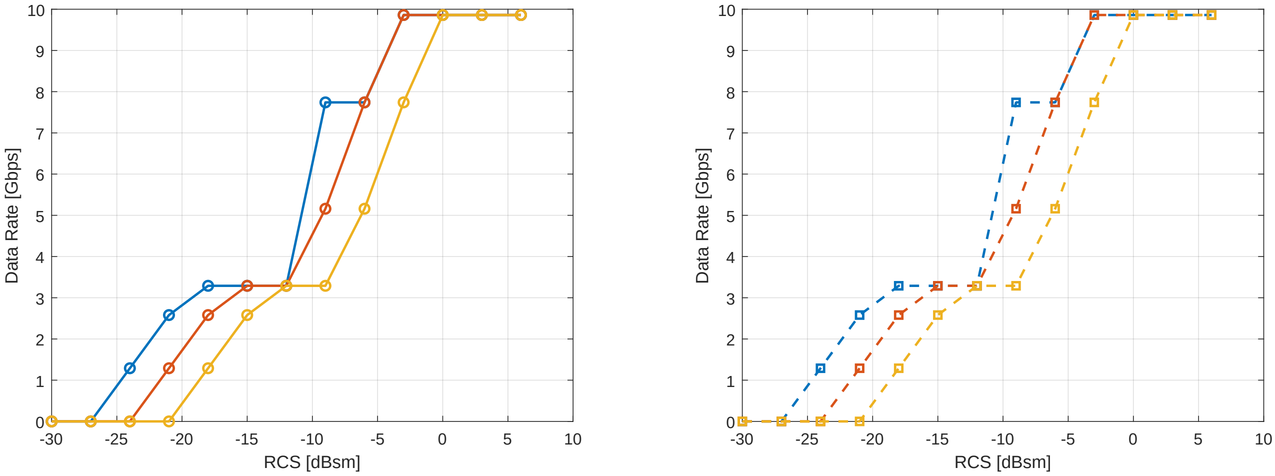

(a) Position 1

(b) Position 2

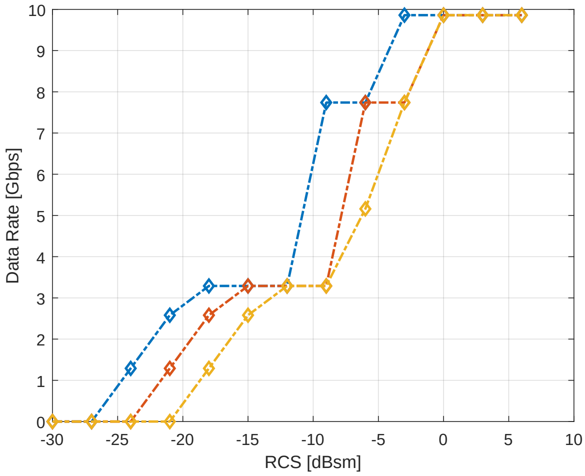

(c) Position 3

Figure 26: Achievable data rates for different RCS values and receiver positions. The blue line represents positions in front of the machine, the red line corresponds to positions on the machine, and the yellow line shows positions behind the machine.

5 Network optimizations for THz-based systems in industrial environments

In this section, we present upper-layer solutions for intelligent THz networks, including both model-based and data-driven protocol designs and network optimizations for multi-goal industrial mesh networks, as well as predictive algorithms for blockage detection and proactive resource allocation.

5.1 MAC and routing protocols design in multi-hop THz networks: a model-based and reinforcement learning approach

Despite their potential, THz networks face critical limitations in terms of coverage, primarily due to high path loss and sensitivity to blockage [70, 71]. Multi-hop communication emerges as a promising solution to address these challenges by leveraging intermediate nodes to extend coverage and maintain robust connectivity, particularly in complex environments such as industrial facilities.

While multi-hop routing has been extensively explored in traditional mesh networks [72, 73], its application to THz systems remains constrained by the limitations of existing protocols, which often rely on complex routing logic, significant control overhead, or unrealistic assumptions such as always-available relay nodes [74–77]. Furthermore, the rapid variability of the THz channel—exacerbated by its millimeter-scale wavelength—poses additional challenges for real-time route adaptation, especially when relying on control-plane signaling [78].

To address these gaps, this work introduces two fully distributed multi-hop schemes tailored for THz mobile radio networks, with a focus on simplicity and adaptability to varying network conditions. Unlike conventional solutions, our approaches operate solely on user-plane data transmissions, avoiding the use of centralized coordination or control signaling. Moreover, we consider realistic constraints such as the inability of relays (which are not chosen a priori) to forward packets when not in reception mode.

The two proposed schemes, Table-Less (TL) and Table-Based (TB), leverage reception phases to forward data but adopt different relay selection strategies: TL sends a duplicate of its data to all neighbors without maintaining routing tables, while TB relies on route discovery and routing tables to determine the next hop towards the destination.

In addition to the two protocols introduced above, a third variant of the TB protocol based on Multi-Agent Deep Reinforcement Learning (MADRL) is proposed. In this scenario, each UE operates as an independent agent, learning its own policy to optimally choose between unicast and broadcast transmission modes. In highly dynamic or congested environments, broadcast transmissions are often preferred, as they increase the likelihood of reaching the final receiver (i.e., the BS in our case) by avoiding reliance on a single next-hop, which may be congested or no longer available. This motivates the design of a variant of the TB protocol that can exploit broadcast transmissions—typically associated with TL—to proactively update routing tables and enhance overall network performance. By leveraging the adaptability and decision-making capabilities of Deep Reinforcement Learning (DRL) in dynamic environments [79], this TB variant is implemented using Deep Q-Learning (DQL) techniques, incorporating advanced mechanisms such as experience replay, fixed Q-value targets, and Double Deep Q Network (DDQN) to ensure both stability and efficiency.

The evaluation uses an Unslotted-Aloha MAC protocol [80], well-suited for THz frequencies due to the limited coverage area of each device, which reduces spatial collisions, and the long, variable propagation delays between devices that lower the likelihood of simultaneous data reception at the receiver.

5.1.1 System Model

Scenario To explore the capabilities of THz frequencies, we consider an IIoT scenario modeled as a parallelepiped with length l, width w, and height h. Within this environment, N industrial assets (e.g., robotic arms,

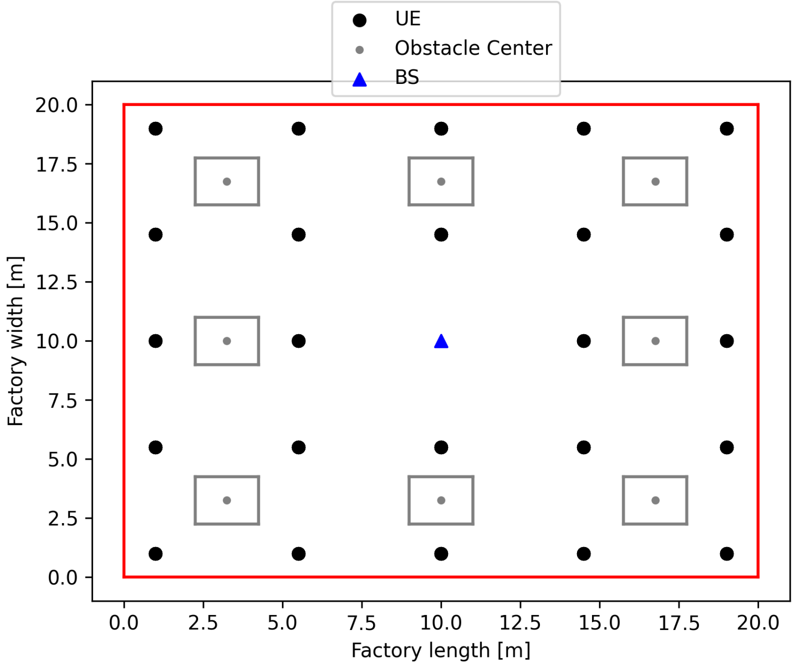

Figure 27: The reference Industrial Internet of Things (IIoT) scenario with UEs organized in a grid layout and the BS at the center.

pumps, valves, and pistons) are arranged in a grid layout [81], with each asset positioned at uniform intervals within a defined framework. This grid-based arrangement ensures consistent spacing between assets, providing a controlled environment for analyzing propagation characteristics such as attenuation and interference. Furthermore, UEs are placed so that each of them can reach at least another UE or the BS (i.e., it has sufficient link budget), ensuring network connectivity throughout the environment. Additionally, this layout reflects typical industrial setups, where assets are evenly distributed to optimize workspace organization and accessibility. The industrial assets, referred to as UEs, require real-time communication with a remote Programmable Logic Controller (PLC). To facilitate this communication, a BS is positioned at the center of the grid. However, communication between the UEs and the BS may experience significant attenuation due to obstructing elements, denoted as 'obstacles', such as walls or metal slabs. These obstacles are modeled as cubes with side s and placed at fixed locations to simulate typical industrial obstructions, as shown in Figure 27.

Traffic Model The achievable network performance is analyzed in a worst-case scenario in terms of offered traffic. This scenario arises when the buffers of the UEs are consistently full, meaning that UEs always have a new DATA packet ready for transmission to the BS immediately after receiving the Acknowledgment (ACK) for the previous transmission. This analysis is crucial for evaluating the maximum network throughput of a THz network (with queue always full), as this metric depends on the type of traffic. Therefore, in the following sections, we will assume that all UEs generate traffic with a fixed DATA packet size, P.

5.1.2 MAC Layer

MAC protocol As a trade-off between simplicity and performance [82], and given the unique characteristics of THz communications that help mitigate collisions (e.g., due to the long propagation delays), we consider a simple MAC protocol based on Unslotted Aloha [80], incorporating a random Back-Off (BO) even in the first transmission to reduce collisions.

In particular, the MAC protocol works as follows. Each UE is initially in IDLE mode. When new DATA has to be transmitted, the UE initiates a BO for $T_{\mathrm{B}} = T_{\mathrm{BO}}\,\xi$, where T_{BO} is the minimum BO duration; ξ is a integer random

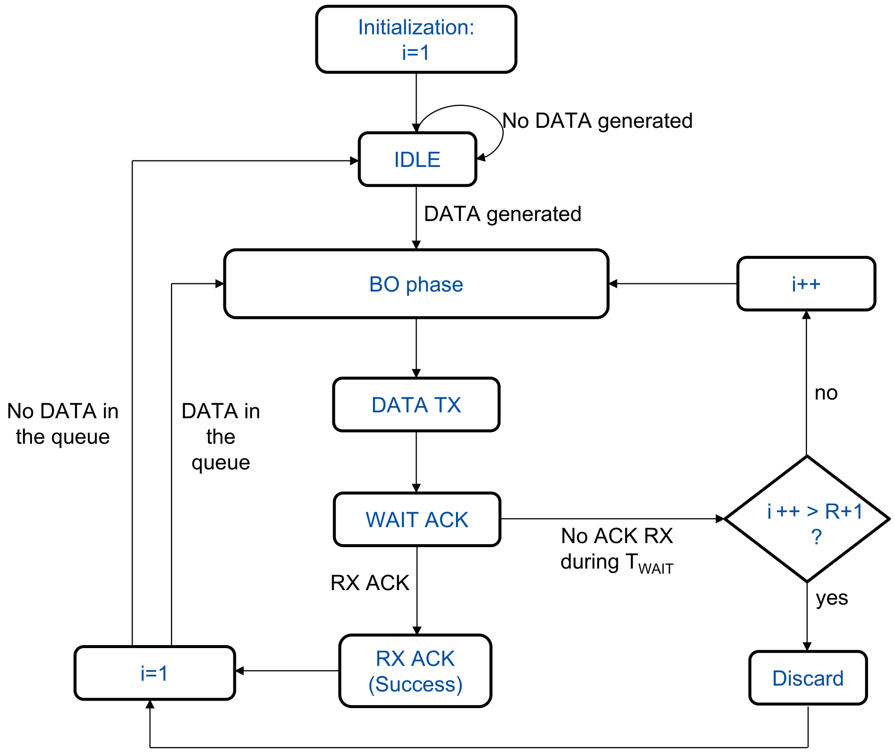

Figure 28: Flowchart of the MAC layer protocol from the UE side.

variable, uniformly distributed in the range $[1; 2^i C]$, with C being an integer defining the maximum duration of the contention window; and i being an integer counting the number of transmission attempts (starting at 1 for the first attempt). At the end of the BO period, the UE sends the DATA and enters reception mode for a maximum period, T_{WAIT}. Specifically, we set $T_{\mathrm{WAIT}} = T_{\mathrm{ACK}} + 2\,\tau_{\mathrm{p_{max}}}$, since it is the time needed to transmit the DATA and receive the corresponding ACK, where $T_{\mathrm{ACK}} = \frac{8\,P_{\mathrm{A}}}{R_{\mathrm{b}}}$, P_{A} is the number of bytes forming an ACK, $R_{\mathrm{b}} = B\log_2(M)$ is the bit rate, M is the modulation order, and $\tau_{\mathrm{p_{max}}}$ is the maximum possible propagation delay in the considered IIoT scenario.

If an ACK is received during T_{WAIT}, the UE either goes back to BO, if a new DATA packet is in the MAC layer queue (as in the case of full queue traffic), or returns to IDLE if no DATA is queued. Conversely, if no ACK is received during T_{WAIT}, the UE retries the transmission up to a maximum number of attempts, R. The flowchart of the MAC layer protocol from the UE side is shown in Figure 28.

Interference modeling Collisions are computed at the receiver side, taking into account the actual propagation delays from the transmitter(s). Specifically, for each useful data transmission received from UE_i, we evaluate whether it is subject to interference from other transmissions, $DATA_j$ originating from user equipment UE_j, where $j \neq i$. The level of interference is computed by first determining the temporal overlap between each interfering transmission $DATA_j$ and the useful transmission $DATA_i$. The overlap duration is normalized by the duration of $DATA_i$ to obtain a weighted overlap factor $t_j = \frac{t_o}{t_i}$, defined as the ratio of the overlapping time, t_o, to the duration of the useful transmission, t_i.

This factor $t_j \in [0, 1]$ represents the fraction of $DATA_i$ that is subject to interference from $DATA_j$. A value of $t_j = 1$ indicates full overlap, while $t_j = 0$ means no overlap.

Each overlap factor t_j is then used to weight the interfering power $I_{j,i}$, which denotes the received interference power at the receiver of UE_i due to the transmission from UE_j. This weighted contribution reflects the effective impact of UE_j's transmission on $DATA_i$. The average interference power observed during the reception of $DATA_i$, denoted by I_{avg}, is then computed as the normalized sum of these weighted interference terms:

$$I_{avg} = \sum_j t_j \sum_i I_{j,i} \tag{134}$$

This expression ensures that only the overlapping portions of interfering transmissions contribute to the interference level and appropriately normalizes their combined impact.

The Signal-to-Interference-plus-Noise Ratio (SINR) at the receiver is subsequently calculated as:

$$\mathrm{SINR}_i = \frac{P_{\mathrm{RX}}}{P_{\mathrm{N}} + I_{\mathrm{avg}}} \tag{135}$$

where $P_{\mathrm{RX}} = 10^{\frac{P_{\mathrm{RX}}[dBm]}{10}}$ is the useful signal power received from UE_i, and $P_{\mathrm{RX}}[dBm]$ is computed as in (9), following the far-field channel model derived experimentally from WP3 (see 3.2.1). $P_{\mathrm{N}} = kT_0F_{\mathrm{RX}}B$ is the noise power assuming the antenna temperature equal to the reference temperature T_0, with k being the Boltzmann constant, F_{RX} is the receiver noise figure, and B the bandwidth.

To assess whether the received transmission can be successfully decoded, the SINR is mapped to the Bit Error Rate (BER), assuming interference behaves as additive Gaussian noise. This is done through the following steps:

1. The SINR is first converted to the ratio between the energy per bit of the received signal divided by the noise power spectral density, $\frac{E_b}{N_0}$ using the relationship:

$$\frac{E_b}{N_0} = \frac{\mathrm{SINR}}{\log_2(M)} \tag{136}$$

   where M is the modulation order of the digital modulation scheme.

2. The BER is then computed assuming an M-QAM modulation:

$$\mathrm{BER} = \frac{(L-1)}{L\log_2(L)} \cdot \mathrm{erfc}\left(\sqrt{\frac{E_b}{N_0} \cdot \frac{3\log_2(L)}{L^2-1}}\right) \tag{137}$$

   where $L = \sqrt{M}$ and $\mathrm{erfc}(\cdot)$ is the complementary error function.

3. Finally, the probability of correctly receiving a data packet of m bits is approximated assuming no forward error correction as:

$$p_s = (1 - \mathrm{BER})^m \tag{138}$$

A transmission is considered successfully received if the computed SINR exceeds a pre-defined threshold $\mathrm{SINR}_{\mathrm{th}}$, which depends on the modulation and coding scheme.

Additionally, we assume a half-duplex mode of operation, meaning that if a device is busy transmitting an ACK and receives DATA simultaneously, the reception is discarded.

Note that we account for propagation delays, which vary for each BS-UE and UE-UE link, as they can be longer than transmission times at THz frequencies and thus impact the overall performance. Indeed, this peculiarity may actually reduce collision probability, since simultaneous transmissions (either DATA or ACK) may not lead to collisions if they are received at different times by the receiver.

5.1.3 Multi-hop Approach

In this section, we describe our proposal, that is, a TL and TB multi-hop algorithm at the network layer. The key innovation of both approaches is that UEs switch to reception mode during the BO period, ensuring they can potentially receive DATA. Consequently, after BO, UEs transmit both their own DATA and any DATA received from other UEs. To realize the above approach, we set the BO duration equal to $T_{\mathrm{B}} = T_{\mathrm{DATA}} + \tau_{\mathrm{p_{max}}} + T_{\mathrm{BO}}\,\xi$, where $T_{\mathrm{DATA}} = \frac{8\,P}{R_{\mathrm{b}}}$ is the time needed to transmit a DATA, and P is the number of bytes forming the DATA. If during BO the UEs correctly receive DATA, and the number of DATA in the MAC layer queue is below its capacity Q, they will enqueue it. In this case, UEs will immediately acknowledge successful reception with an ACK. Note that, during the WAIT phase, each UE remains in reception mode to receive the ACK for all transmitted DATA. This also allows a UE to receive DATA from other UEs during T_{WAIT} and forward them accordingly.

Table-less multi-hopping In the TL solution, after the BO period, each UE always sends its own DATA with a broadcast address, thereby producing the well-known broadcast storm problem [72] but, at the same time, increasing the likelihood that at least one UE within its transmission range is listening to the channel (i.e., is in reception mode during the BO or WAIT phase) and can subsequently forward the DATA to the BS. Then, UEs enter the reception mode either (i) for a maximum duration of T_{WAIT} or (ii) until they correctly receive the ACK before T_{WAIT} expires. It is important to note that the non-negligible propagation delays, which contribute to reducing collisions and packet loss, have been taken into account. This approach contrasts with conventional flooding methods in the literature, where such delays are often overlooked.

Table-based multi-hopping In the TB approach, UEs create a neighbor table containing for each other UE: (i) an indication of whether it has the BS in its own table, (ii) the number of ACKs received from it; and (iii) the corresponding received powers P_{RX}. Hence, after the BO phase, each UE transmits its own DATA with a unicast address to the optimal receiver selected based on the neighbor table. Specifically, the optimal receiver is always the BS if it is present in the neighbor table; otherwise, it is the UE that has the BS in its table and/or from which the highest number of ACKs has been received at the highest P_{RX}.

When the neighbor table is empty, after BO, the UEs initiate a *network discovery* phase without adding any control plane overhead. Indeed, UEs simply transmit its own DATA in broadcast and then enter reception mode for T_{WAIT}. During this time, they populate the neighbor table based on the ACKs received from all receivers that have correctly collected the broadcast transmission. Once the neighbor table contains at least one entry, the network discovery phase concludes. This means that, in the next BO phase, DATA are sent in unicast to the optimal receiver and the WAIT phase ends before T_{WAIT} if the correct ACK is received (as in the TL case)[5].

Note that each UE removes an entry from the table after a Time To Live, TTL, expires, which is defined as the number of consecutive missed receptions of an ACK from a specific neighbor.

Final remarks For both TB and TL approaches, it is important to note that:

1. During T_{WAIT}, UEs can also receive DATA from other UEs. In this case, they will transmit the corresponding ACK(s) at the end of the WAIT phase;
2. The BS discards any duplicate DATA generated by the same UE but forwarded by different UEs;
3. UEs discard any received DATA that (i) is already present in their queue, (ii) originates from themselves (to prevent loops), or (iii) exceeds the predefined hop count limit H (i.e., the maximum number of relays to be traversed before reaching the BS);

---

[5] Note that the broadcast transmissions of DATA during the network discovery phase in TB differ from those in TL. In TB, UEs broadcast DATA packets during network discovery to build a neighbor table, which is then used to select the next hop toward the BS. Once the neighbor table is established, subsequent transmissions are performed via unicast toward the chosen next-hop. In contrast, TL does not maintain a neighbor table: each UE broadcasts its DATA using a broadcast address, and no unicast transmissions are employed.

4. The UEs discard any ACK received during BO or WAIT but not intended for them.

5.1.4 Multi-Agent Deep Reinforcement Learning

In the MADRL algorithm proposed both the training and execution processes are fully decentralized [83]. Multiple agents (i.e., the UEs) interact with a common and unknown environment, that is the industrial scenario described in Figure 27. Each UE independently learns a behavioral policy, denoted as π, to optimally select the most suitable transmission type (either unicast or broadcast) for the given task. The UEs store their interactions with the environment in a replay buffer, which is used to train and update their respective learning models that represent the policy π.

Specifically, the learning problem is formalized as a Markov Decision Process (MDP) by means of the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R})$:

- $\mathcal{S}$ is the *state space*, which represents the set of all possible *states*. A state is the limited observation of the environment available to the UE. At each time step t, the agent observes the state s_t. In particular, a state $s_{t,i} \in \mathcal{S}$ is represented by the ratio of the ACKs received by the neighbor i and the number of actions performed that involve that neighbor, $\alpha_{t,i} \in \mathcal{S}$. In order to have a state space with fixed dimension, we consider only the N_{BN} best neighbors, that is, those with the highest ratio.

$$s_{t,i} = ACK_{t,i}/\alpha_{t,i} \tag{139}$$

- $\mathcal{A}$ is the *action space*, which defines the set of all possible *actions* that the UE can take to interact with the environment. In this work, the action $a_t \in \mathcal{A}$ at time step t is binary, where 0 corresponds to broadcast and 1 to unicast;
- $\mathcal{T}$ is the transition probability function, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S} \times (0, 1)$. Given the current state s_t and action a_t, the function $\mathcal{T}$ describes the probability of transitioning to the next state s'_{t+1};
- $\mathcal{R}$ is the reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$. It assigns a real-valued score, r_t, based on the current state s_t, the action a_t performed, and the resulting next state s'_{t+1}. In particular, we define r_t as follows:

$$r_t = \begin{cases} \beta - \delta N_{\text{RTX}} & \text{if } a_t = 1, \\ \eta\,(\eta - \delta N_{\text{RTX}}) & \text{if } a_t = 0, \end{cases} \tag{140}$$

  where β is the unicast factor, δ is the retransmission factor, $N_{\text{RTX}} \leq R$ is the number of retransmissions before receiving the ACK, and η is the broadcast factor.

The learning agent interacts with the environment at discrete time steps $t \in \{0, 1, \dots, t_{\text{ep}} - 1\}$, where t_{ep} represents the number of steps in a single episode, which has a duration of T_{S}. At the start of each episode, every agent performs a forced broadcast action to discover its neighbors. This action is repeated until at least one neighbor is added to the neighbor table. Following this initialization phase, when new DATA needs to be transmitted, the UE observes the current state s_t and selects an action a_t according to the policy $\pi(a_t|s_t)$, or simply π for brevity. After executing action a_t, the environment provides a reward r_t to the UE, and it transitions to a new state s'_{t+1}. In our case, the reward is computed when the UE has received the ACK for the transmitted DATA or the UE has discarded the DATA after exceeding the maximum number of retransmissions, R.

The objective of the learning model is then to identify the optimal policy π^* that maximizes the expected discounted return $G_t = \sum_t^{t_{\text{ep}}-1} r_t \gamma^t$, where $\gamma^t \in [0, 1), \forall t \in \{0, 1, \dots, t_{\text{ep}} - 1\}$ is a hyperparameter called the discount factor, which quantifies the importance of future rewards when executing an action a_t at time step t. By resorting to the Bellman model [84], we define the optimal Q-values as the expectation of G_t, given that the agent executes action a_t in state s_t and subsequently acts optimally until the end of the episode, that is:

$$Q^*(s_t, a_t) = \mathbb{E}[G_t | s_t, a_t]. \tag{141}$$

Following this, it can thus be stated that the optimal policy π^* to be followed in state s_t corresponds to selecting the action a_t that yields the highest Q-value, expressed as:

$$\pi^*(s_\mathrm{t}) = \operatorname*{argmax}_{a_\mathrm{t}} Q^*(s_\mathrm{t}, a_\mathrm{t}), \tag{142}$$

In this work, owing to the potentially large number of states (i.e., neighbors), we adopt DQL [85], where each agent utilizes a Deep Neural Network (DNN) to approximate $Q^*(s_\mathrm{t}, a_\mathrm{t})$ for any state-action pair using a parameter vector $\boldsymbol{\Theta}$. A DNN used to estimate Q-values is called Deep Q Network (DQN), and is designed to take the state s_t as input and output the parametrized estimates of the Q-values, i.e., $Q_{\boldsymbol{\Theta}}(s_\mathrm{t}, a_\mathrm{t})$, for each possible action $a_\mathrm{t} \in \mathcal{A}$.

To train a DQN, at each time step t, each agent fills a replay buffer of dimension D with experiences e_t, which are tuples containing s_t, a_t, r_t, and $s'_{\mathrm{t}+1}$. In this way, it is possible to define the target Q-values as:

$$Q_\mathrm{target}(s_\mathrm{t}, a_\mathrm{t}) = r_\mathrm{t} + \gamma \cdot \max_{a'_{\mathrm{t}+1}} Q_{\boldsymbol{\Theta}}(s'_{\mathrm{t}+1}, a'_{\mathrm{t}+1}), \tag{143}$$

and therefore we can perform a training step applying any Gradient Descent algorithm. To balance between exploration of the MDP and exploitation of the learned policy, we adopt an ϵ-greedy approach where, at each time step t, the agent selects the action randomly with probability ϵ, or greedily (i.e., choosing the action associated with the highest Q-value) with probability $1-\epsilon$. During the training phase the ϵ starts from ϵ_max, then it decreases of ϵ_step until it reaches ϵ_min in 15 episodes.

To stabilize training [86], we also resort to the Fixed Q-Value Targets variant [85], where the agent utilizes two DQNs. One is called *online DQN* and is used to provide the Q-value estimates at each time step t by means of the parameter vector $\boldsymbol{\Theta}$, whereas the second is named *target DQN*, and is used to define the target Q-values with parameter vector $\boldsymbol{\Theta}^-$. The online DQN is trained every episode, and the target DQN is a clone of the online DQN, where $\boldsymbol{\Theta}^-$ is overwritten with $\boldsymbol{\Theta}$ every t_tr time steps. Consequently, eq. (143) is modified as follows:

$$Q_\mathrm{target}(s_\mathrm{t}, a_\mathrm{t}) = r_\mathrm{t} + \gamma \cdot \max_{a'_{\mathrm{t}+1}} Q_{\boldsymbol{\Theta}^-}(s'_{\mathrm{t}+1}, a'_{\mathrm{t}+1}), \tag{144}$$

To reduce the overestimation bias seen in traditional DQN [85], we consider the DDQN paradigm, where the online DQN is used to select the best actions for the next states, and the target DQN to estimate the Q-values for these best actions. Therefore, eq. (144) becomes:

$$Q_\mathrm{target}(s_\mathrm{t}, a_\mathrm{t}) = r_\mathrm{t} + \gamma \cdot Q_{\boldsymbol{\Theta}^-}(s'_{\mathrm{t}+1}, \operatorname*{argmax}_{a'_{\mathrm{t}+1}} Q_{\boldsymbol{\Theta}}(s'_{\mathrm{t}+1}, a'_{\mathrm{t}+1})). \tag{145}$$

It is worth mentioning that, however, we do not rely on the common Dueling Double Deep Q Network (3DQN) [85], as it is primarily designed for environments with a high number of actions that do not significantly alter the environment itself. In contrast, in our task, depending on the action chosen, the environment can vary significantly. For example, a broadcast action produces much more congestion than a unicast on, with a clear impact on the future observations of the agents, making the 3DQN unsuitable for our needs.

5.1.5 Key Performance Indicators

Network Throughput The network throughput, S, is defined as the number of information bits per second successfully received at the MAC layer of the BS, that is,

$$S = \frac{PN_\mathrm{R}}{T_\mathrm{S}}, \tag{146}$$

Table 8: Simulation parameters.

Symbol	Description	Value
f_c	Carrier frequency	300 GHz
B	Bandwidth	15 GHz
$P_{\mathrm{TX,UE}}$	Power transmitted from a UE	30 dBm
$P_{\mathrm{TX,BS}}$	Power transmitted from the BS	35 dBm
G_{UE}	Antenna gain of a UE	5 dB
G_{BS}	Antenna gain of the BS	10 dB
F_{UE}	Noise figure of a UE	5 dB
F_{BS}	Noise figure of the BS	4 dB
η_{UE}	Antenna efficiency of a UE	0 dB
η_{BS}	Antenna efficiency of the BS	0 dB
T_0	Reference temperature	290 K
M	Modulation order	4
$SINR_{\mathrm{th}}$	SINR threshold	10 dB
C	Integer value defining the BO period	5
R	Maximum number of retransmissions	3
P_{A}	Size of an ACK	10 byte
P	Size of a DATA	20, 50, 100 byte
O	Overhead of DATA	10 byte [7]
Q	Length of the UEs' queue	8
TTL	TTL for an entry in the neighbor table	3, 12
T_{BO}	BO minimum time slot duration	1.6 ns
N_{S}	Number of simulations	50
T_{S}	Simulation time	0.5 ms

where P is the DATA size, N_{R} is the number of DATA successfully received at the BS, and T_{S} is the simulation time.

Latency The average latency, $\overline{L}$, is defined as the average time needed by UEs to transmit DATA with success, that is,

$$\overline{L} = \frac{1}{N}\sum_{j=1}^{N}\frac{1}{N_{\mathrm{P}_j}}\sum_{i=1}^{N_{\mathrm{P}_j}} L_{j,i} \tag{147}$$

where N_{P_j} is the number of DATA generated by the j-th UE and not discarded, and $L_{j,i}$ is the time interval from the generation of the i-th DATA by the j-th UE to the reception of the corresponding ACK[6]. When the BS successfully receives the DATA from the first relay, we calculate the latency for that packet, denoted as $L_{j,i}$, by summing T_{ACK} and the propagation delay of the ACK. It is important to note that the ACK from the BS may not be received by the originating UE, as it is not in reception mode after already receiving confirmation from the first relay.

5.1.6 Numerical Results

Simulations parameters, if not otherwise specified, are reported in Table 8, where some of them were set based on discussions with WP5 partners. In particular, all results have been obtained by averaging over N_{S} simulations of duration T_{S}. We set the $SINR_{\mathrm{th}} = 10$ dB, to ensure a link success probability $p_s = 0.88$, when $M = 4$.

---

[6] Note that UEs can receive ACKs from other UEs acting as relays to the BS.

Table 9: MADRL parameters.

Symbol	Description	Value
α	Learning rate	0.01
γ	Discount factor	0.95
δ	Retransmission factor	0.1
D	Replay buffer size	15000
ϕ	Batch size	256
t_{tr}	Training periodicity of target DQN	15
ϵ_{max}	Maximum exploration probability	1.0
ϵ_{min}	Minimum exploration probability	0.01
ϵ_{step}	Decreasing step of the exploration probability	0.067
N_{BN}	Number of the best neighbors for input state	4

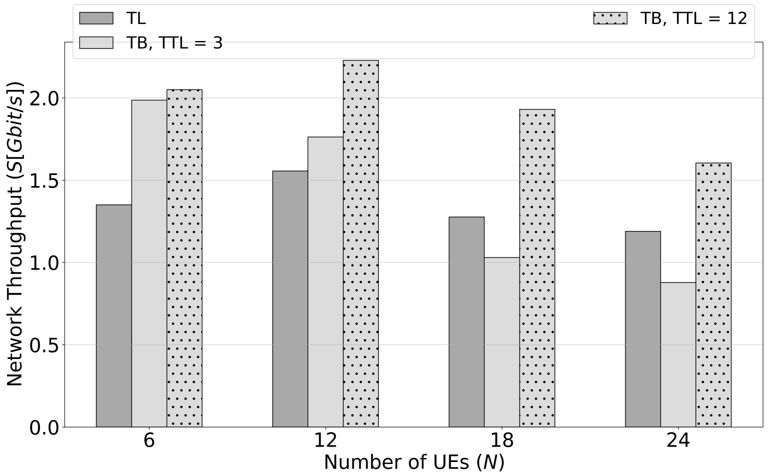

Figure 29: Network Throughput, S, as a function of the number of UEs, N, for TL and TB with $TTL = 3$ and 12.

Static scenario We start our analysis by comparing the performance of TB and TL in a static scenario where both UEs and obstacles occupy fixed positions.

Figure 29 presents the network throughput, S, as a function of the number of UEs, N, for TL and TB with TTL = 3 and 12. It can be observed that in all the three configurations, S reaches its maximum with $N = 12$ but then decreaes as the number of collisions increases. When $N = \{6; 12\}$ TB outperforms TL due to efficient unicast forwarding, which reduces channel overhead. TB with TTL = 12 achieves the highest overall throughput due to infrequent neighbor rediscovery and stable next-hop selection. As the number of UEs increases, $N = \{18; 24\}$, the throughput decreases for all schemes due to increased collisions. However, TB with TTL = 3 suffers a sharper decrease. The frequent rediscovery of neighbors leads to suboptimal next-hop choices and more transmission steps, resulting in a deterioration of performance. In contrast, TL becomes relatively more effective in dense scenarios: its broadcast approach increases the chance of reaching a UE that forwards the packet, despite the higher overhead. These results show that in high-frequency channels where collisions with multiple packets lead to losses (with our setup if the number of interferers is greater than 2, the DATA is lost) and in turn this can lead to the removal of a neighbor from the neighbor table, the benefits of lower overhead given by TB decrease with network density. A longer TTL = 12 mitigates this problem by preserving the state of neighbors for longer and reducing the instability induced by rediscovery.

Then, Figure 30, shows the Average Latency, $\overline{L}$, as a function of the number of UEs, N, for TL and TB with

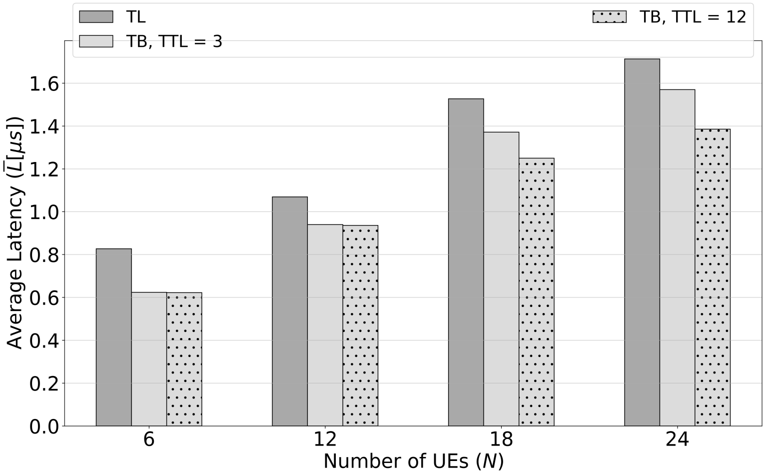

Figure 30: Average latency, $\bar{L}$, as a function of the number of UEs, N, for TL and TB with $TTL = 3$ and 12.

TTL = 3 and 12. As expected, $\bar{L}$ increases with N due to the higher number of collisions and retransmissions caused by the greater number of UEs competing for the channel, while still remaining below 1.75 μs. These results further confirm the better performance of TB under static conditions because of the higher reliability given by the choice of the next hop. Moreover, when $N = \{6; 12\}$, TB with TTL = 3 and 12 show comparable latency performance, as rediscovery of the next hop rarely occurs under low network load. However, when $N = \{18; 24\}$, the higher collision rate triggers more frequent rediscovery in the TTL = 3 configuration. This increases the likelihood of selecting suboptimal relays, resulting in higher latency, as also reflected in the throughput degradation shown in Figure 29.

Based on these findings, it can be concluded that, in static scenarios, there is a need for a proper tuning of the TTL parameter based on network congestion level. Specifically, in our studies, TB configuration with TTL = 12 consistently achieves higher performance compared to TTL = 3 case. Consequently, the subsequent analysis will consider only the TB configuration with TTL = 12.

Figure 31 shows the network throughput S, as a function of the payload size P for different contention window sizes W. We set $N = 12$, since it is the number of UEs for which the highest S is achieved in Figure 29 and we compare the performance of TB with TTL = 12 and TL. Increasing W improves performance for both TB and TL schemes by reducing collisions: a larger contention window introduces greater transmission time diversity among UEs, lowering the likelihood of packet overlap. This effect shifts the throughput peak from $P =$ 40 B with $W =$ 5, to $P =$ 60 B with $W =$ 10, reflecting improved robustness to channel contention. Across most payload sizes, TB outperforms TL thanks to its more efficient unicast forwarding and lower channel overhead compared to the broadcast transmissions used in TL. However, as payload size increases ($P = \{80; 100\}$), the collision rate rises and TB becomes more susceptible to neighbor rediscovery events, particularly when collisions disrupt ACK receptions. This results in a performance drop that eventually causes TL to surpass TB, as the broadcast-based forwarding in TL ensures a more consistent delivery path under heavy traffic conditions.

Figure 32 shows the Average Latency, $\bar{L}$, as a function of the payload size P for different contention window sizes W. As in Figure 31 we set $N = 12$ and we compare the performance of TB with TTL = 12 and TL. As shown in the figure, increasing the contention window size to $W = 10$ results in higher average latency across all schemes. This is attributed to the longer average BO duration before a UE initiates a transmission attempt. Additionally, the TL scheme provides higher latency compared to TB, particularly as P increases. This is due to the inherent inefficiency of broadcast-based forwarding in TL, which does not always ensure the selection of the most efficient path toward the BS. In contrast, TB leverages explicit next-hop selection, which more frequently results in shorter and more stable routing paths, thereby reducing overall latency.

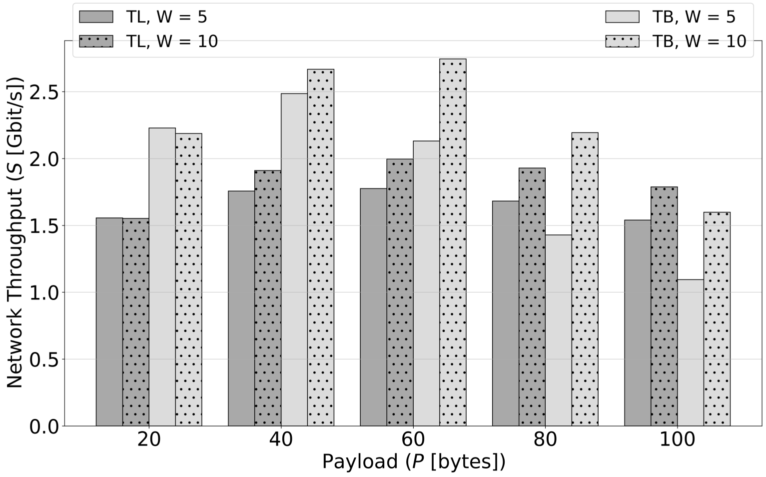

Figure 31: Network Throughput, S, as a function of the payload, P, and the contention window W. We set $N = 12$, for TL and TB with TTL = 12.

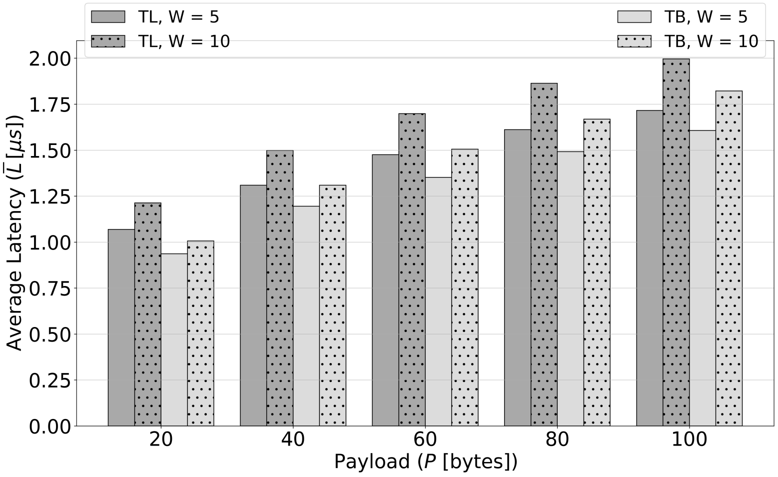

Figure 32: Average latency, $\bar{L}$, as a function of the payload, P, and the contention window W. We set $N = 12$, for TL and TB with TTL = 12.

Dynamic Scenario We extended our analysis by focusing on the TB and its MADRL extension under dynamic conditions, where UEs change their position during a simulation run. Specifically, UEs are moved to modify their potential neighbor(s). As a result, some UEs may lose their connection to the BS either because their next-hop moves to a new location or because they move outside the communication range of their next-hop.

Once the TTL expires, the UEs must start a new network discovery process to find a new path to reach the BS. As in the static case, we evaluate two TTL values: 3 and 12. To study the effect of mobility, each UE performs a fixed number of movements during a simulation run, and this number will be denoted as m in the following.

The higher the value of m, the greater the impact of mobility, since UEs must perform network discovery more frequently as they change position (to emulate potential environmental or channel variations). However, this type of dynamicity has a stronger effect in configurations with a large TTL, as UEs require more time to detect the need to reselect the next hop, leading to packet losses. Therefore, to ensure a fair comparison (i.e.,

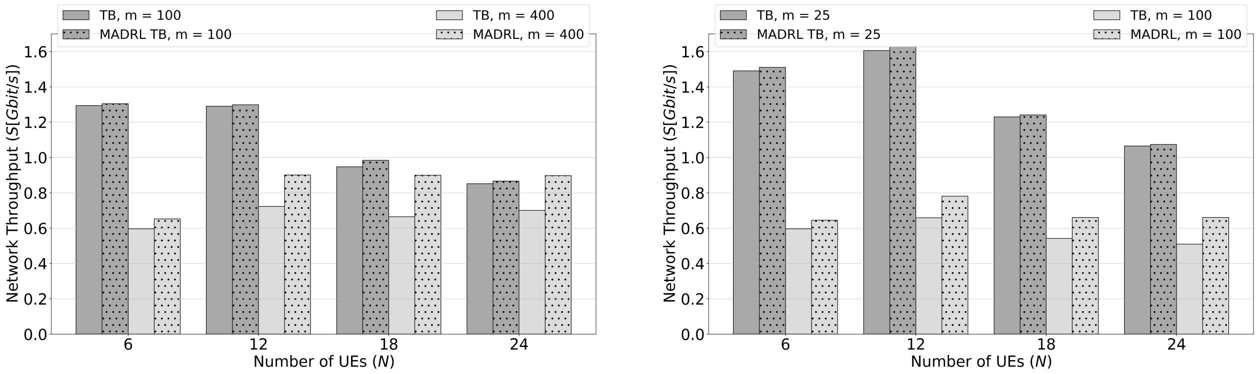

(a) TTL = 3 (b) TTL = 12

Figure 33: Network Throughput, S, as a function of the number of UEs, N, for TB and MADRL in a dynamic scenario with m movements.

scenarios with a similar impact of mobility), we empirically set $m = [100; 400]$ when TTL = 3, and $m = [25; 100]$ when TTL = 12.

Figure 33 shows the network throughput, S, as a function of the number of UEs, N, for both the TB and MADRL approaches, under TTL = 3 (with $m = [100; 400]$) and TTL = 12 (with $m = [25; 100]$). In all configurations, S increases with N, reaching a peak at $N = 12$, and then gradually decreases as N grows further. This decline is mainly due to increased channel overhead and a higher probability of collisions, which degrade overall network performance[8]. As expected, higher mobility (i.e., higher m) leads to lower throughput, as UEs have less time to discover new neighbors before moving again. Overall, MADRL outperforms TB for the same values of N and m, with a performance gap widening as mobility increases, reaching a maximum percentage increase of 36% with $N = 18$, $m = 400$, and TTL = 3. This improvement stems from the DRL-based approach's ability to rapidly adapt to topology changes by selecting optimally between unicast and broadcast transmissions, thereby accelerating neighbor table reconfiguration and enhancing throughput under dynamic conditions. The reader can also anyze the results for $m = 100$ with both TTL = 3 (Figure 33(a)) and TTL = 12 (Figure 33(b)). It can be seen that a shorter TTL leads to faster neighbor table updates, which in turn improves network throughput in dynamic scenarios.

Figure 34 shows the average latency, $\overline{L}$, as a function of the number of UEs, N, for both the TB and MADRL approaches, under TTL = 3 (with $m = [100; 400]$) and TTL = 12 (with $m = [25; 100]$). As expected, $\overline{L}$ increases with N due to higher channel contention and retransmissions, though it remains below 1.15 μs. With TTL = 3 and TTL = 12, the trends differ. In the first case, the $\overline{L}$ values are always higher than the corresponding case with TTL = 12. This occurs because with TTL = 3, the neighbor table is updated more frequently, increasing the likelihood that UEs farther from the BS discover the correct path and transmit successfully. Focusing on the comparison between the TB and the MADRL, when TTL = 12 and m is highest, the latter approach consistently shows higher $\overline{L}$ than the former for every N. The only exception is when $m = 100$ and $N = 6$, where MADRL results in lower average latency at this TTL. For $N = 18, 24$ and both TTL = 3 and TTL = 12, the trends are similar and can be attributed to the speed at which the neighbor table updates, which depends on the TTL value. However, at high N, the increased number of collisions leads to more unsuccessful transmissions, which may cause the TTL to expire even before the topology changes due to movement. This phenomenon, combined with MADRL, leads to a faster reconfiguration of the neighbor table. For lower values of N (i.e., $N = 6, 12$) and the lowest m, both approaches exhibit similar behavior. In contrast, with the highest m, the trend diverges due to the capabilities of the DRL algorithm. At lower TTL values, collisions more frequently cause the TTL to expire, making the algorithm more effective in anticipating new path discoveries. Since the

---

[8] An exception is observed with TB for TTL = 3 and $m = 400$, where S continues to grow slightly beyond $N = 18$. Here, high mobility combined with a short TTL enables frequent neighbor table updates. This leads to more reconfigurations and a broader set of paths to the BS, helping maintain or slightly improve throughput even at higher N.

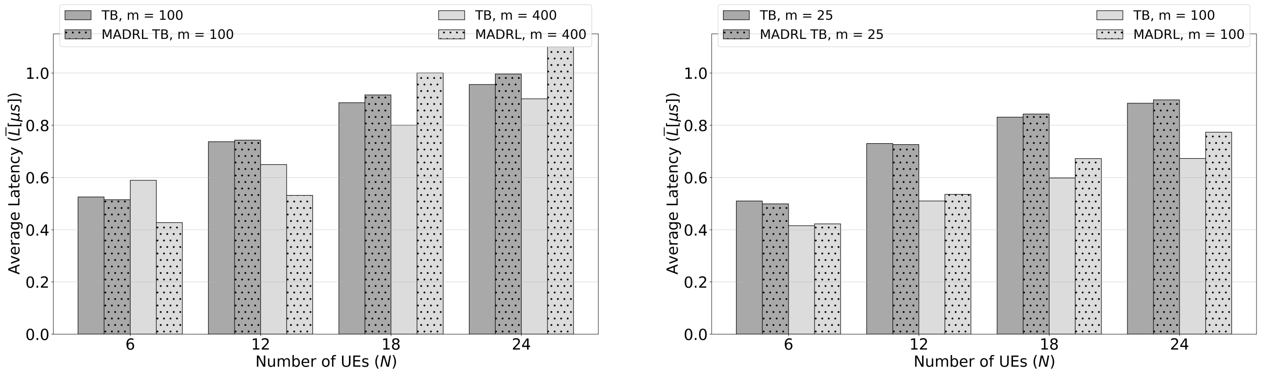

(a) TTL = 3

(b) TTL = 12

Figure 34: Average latency, $\overline{L}$, as a function of the number of UEs, N, for TB and MADRL in a dynamic scenario with m movements.

time between two consecutive movements is only slightly longer than the TTL expiration time, UEs located at the first hop are often able to reconfigure themselves and transmit successfully. Because $\overline{L}$ is calculated over all packets successfully delivered to the BS, a lower number of hops taken by users that reach the BS results in a lower average latency.

Overall, MADRL yields higher latency than TB due to its ability to optimally select the most suitable transmission type depending on the network conditions and this allows more distant UEs to reach the BS. While this improves network fairness, it also leads to longer multi-hop paths, increasing $\overline{L}$. To compare the impact of TTL, we can focus on $m = 100$, where $\overline{L}$ is higher with TTL = 3 because more UEs located farther from the BS can successfully transmit, thanks to their faster neigbors' table reconfiguration, increasing the average delay.

In summary, TTL = 3 allows faster reconfiguration of the neighbor table with respect to TTL = 12, making the network more robust to variability and mobility in the scenario. The MADRL approach improves network throughput, reaching an increase of 36%, by allowing more UEs to successfully discover paths to the BS, but this also leads to higher average latency due to an increased number of UEs able to find a new path to reach the BS with $N = 18, 24$ and both TTL values. Instead, with a fewer number of UEs the trend is different due to the simpler network topology. In particular, it worth note the difference with $N = 6, 12$ between TTL = 3 and TTL = 12 looking to the results with the highest m. With TTL = 3, $\overline{L}$ is lower with the MADRL, reaching the maximum percentage decrease of 30% with $N = 6$, whereas with TTL = 12 the average latency results similar to the TB case, but the MADRL presents higher values.

5.2 AI-based MAC optimization for multi-goal scenarios

Next-generation mobile radio networks are designed to support a wide range of applications and use cases. However, they are expected to face significant challenges in meeting the stringent, heterogeneous, time-varying, and often conflicting requirements of emerging services [87]. These challenges have led to the emergence of multi-goal wireless networks, communication systems in which device requirements evolve dynamically in both space and time [88]. As a result, such networks aim to jointly optimize multiple, and potentially conflicting, performance metrics (e.g., maximizing throughput while minimizing energy consumption) [88, 89]. Addressing these needs requires the network to understand and adapt to the temporal and spatial variability of device requirements. However, the high degree of variability and dynamic behavior in industrial settings makes it difficult and often inefficient to rely on traditional MAC protocols based on static optimization. In this work, we explore the use of AI, and specifically Reinforcement Learning (RL), as a more flexible and effective alternative. Several recent studies have investigated the use of AI techniques to improve MAC protocols and optimize key performance metrics [90–94]. These works primarily focus on AI-driven parameter tuning, such as optimizing the Congestion Window (CW) in random access protocols like slotted ALOHA and Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) [90–94]. However, existing solutions are typically limited in scalability, often considering a small number of nodes, and are not designed for IIoT environments. Moreover, these studies are not applied to high frequency ranges, such as THz frequencies [95], in which effects that were negligible at low frequencies, such as propagation delays, now have a great impact and cannot be neglected [96–98]. Furthermore, they do not address the complexities of multi-goal scenarios, where diverse and evolving requirements must be jointly satisfied. This is, indeed, quite typical in an industrial plant's scenario.

To address the aforementioned gaps, we develop an RL model for MAC optimization at THz frequencies in industrial settings with heterogeneous requirements, that is, a multi-goal environment. By understanding the variability of the industrial processes and their requirements, the proposed RL-based algorithm, located at the BS, is capable of selecting the appropriate CW to assign to each sensor to satisfy its specific requirements.

In particular, we consider an industrial plant equipped with machines, sensors, and a central BS. Sensors sense the surrounding environment and acquire telemetry data (e.g., including temperature measures), which have to be sent to the BS to control and optimize the production process. The communication relies on the Carrier Sense Multiple Access (CSMA) as MAC protocol. Depending on the data content (e.g., measured temperature), the packet acquires a different priority, imposing different requirements in terms of latency. Since industrial scenarios are dynamic by nature, it is not possible to predict the variability of the data content and assign the proper CW to each sensor. Therefore, we propose a centralized RL model, running at the BS, which learns the variability and the requirements of the industrial processes to dynamically select the optimal CW to be assigned to each sensor.

5.2.1 System Model

Scenario We consider a representative industrial plant of size $F_x \times F_y \times F_z$ m^3 comprising M machines modeled as cubes of size $M_x \times M_y \times M_z$ m^3 for goods production, and a BS located in the top center of the facility, as depicted in Figure 35. C sensors are uniformly distributed within the machines' volume. Sensors are equipped with communication capabilities and are responsible for monitoring the surrounding environment by acquiring measurement data, denoted T_m, which must be sent to the BS. Data packets sent by the sensors have a size P_{data}, while ACKs have a size P_{ack}.

Time is discretized into fixed-length intervals, referred to as ticks, each with duration t_{tick} [s]. Sensors generate data periodically and synchronously. The interval between two successive data generation events is referred to as a time step, with duration t_{step} [s]. We index the time steps using an integer variable t. An episode is defined as a burst of N data packets generated by each sensor and potentially transmitted to the BS. Accordingly, each episode consists of $t = 1, 2, \dots, N$ time steps, leading to a total episode duration of $t_{\text{ep}} = N \cdot t_{\text{step}}$ [s].

For simulation purposes, we set the tick duration to match the time required to transmit the smallest packet,

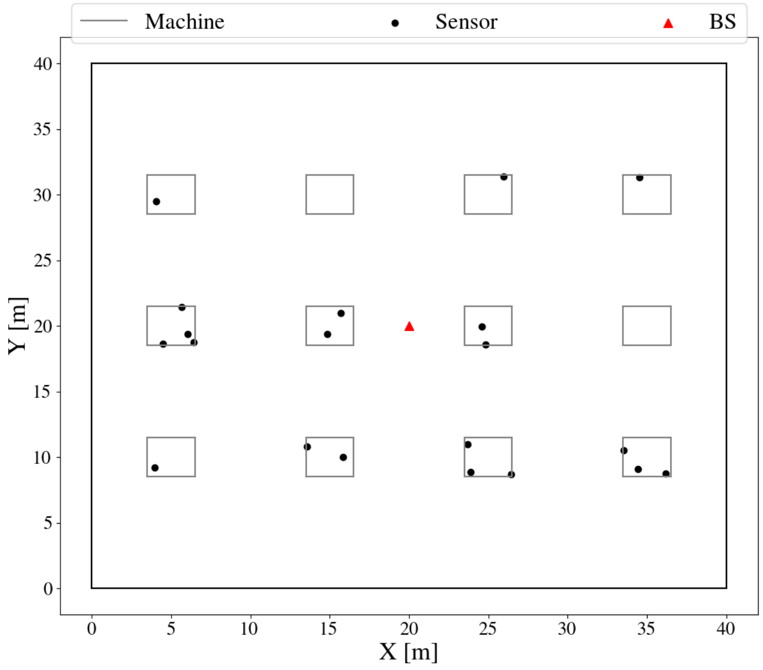

Figure 35: The considered industrial plant comprising machines, sensors, and a BS.

i.e., an ACK. Hence, $t_{\text{tick}} = P_{ack} \cdot 8/R_b$, where R_b denotes the bit rate. The time step duration is then defined as $t_{\text{step}} = N_{\text{tick}} \cdot t_{\text{tick}}$, where N_{tick} is an integer.

Multi-Goal Traffic Depending on the data content, packets may have different goals, resulting in different requirements in terms of latency. Data packets that have more stringent requirements in terms of latency are denoted as high priority (HP), while those having less stringent latency requirements are denoted as low priority (LP). Given $T_{m,t,c}$ the data measured at time step t by the generic sensor c and $T_{m,th}$ the corresponding threshold, two cases can occur:

- $T_{m,t,c} \geq T_{m,th}$: the sensor c is in HP, imposing more stringent requirements in terms of latency $\tau_{t,c}$, i.e.,$\tau_{t,c} \leq \tau_{th,hp}$;
- $T_{m,t,c} < T_{m,th}$: the sensor c is in LP, imposing more relaxed requirements in terms of latency, i.e., $\tau_{t,c} \leq \tau_{th,lp}$,

where $\tau_{th,hp}$ and $\tau_{th,lp}$ are the maximum latencies allowed for LP and HP data, respectively, with $\tau_{th,hp} < \tau_{th,lp}$. The variability of data content throughout an episode is modeled as follows: at the initial time step ($t = 1$), each sensor has a 50% probability of being in either the HP or LP state (i.e., with 50% of probability a sensor generate a data whose content exceeds $T_{m,th}$ and is thus considered HP, or vice-versa). Throughout the episode, each sensor may change its priority state according to a two-states Markov chain, as illustrated in Figure 36. If a sensor is in a given state (HP or LP) at time step t, it remains in the same state at the next time step with probability q, or switches to the other state with probability $1 - q$. In detail, q is defined as:

$$q = \begin{cases} 0.5 & \text{if } t \geq t_{c,P} \\ 1 & \text{if } t < t_{c,P} \end{cases} \tag{148}$$

Specifically, $q = 0.5$ if $t \geq t_{c,P}$, with t being the actual time step and $t_{c,P}$ being a random time step in the range $[1, N]$. While $q = 1$ otherwise.

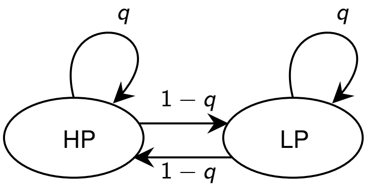

Figure 36: The two-states Markov chain that characterizes the temporal evolution of the multi-goal traffic.

MAC protocol We consider the CSMA as MAC protocol, whose flow diagram is illustrated in Figure 37. Under the assumption that all sensors are synchronized in time, each sensor is initially in an idle state, with $n_{sens} = 0$ and $n_{tx} = 0$, where n_{sens} represents the number of times a sensor sensed the channel, and n_{tx} is the number of times the sensor transmitted the same packet. If data is generated, the sensor enters a BO state; otherwise, it remains in an idle state. The sensor remains in the BO state for a duration T_B, which corresponds to an integer value, randomly chosen in the range $[1; W_i]$, where W_i is the CW to be used. At the end of the BO phase, the sensor performs sensing. If the channel is sensed as free, the sensor transmits data; otherwise, if the channel is sensed as busy, the sensor goes again in a BO phase for a given maximum number of times $n_{sens,max}$, after which the packet is discarded. After transmitting data, the sensor goes into a reception state. If the ACK is received, the packet is considered successfully delivered; otherwise, after a τ_{out} interval of time, the sensor tries again to transmit the same data for a given number of times $n_{tx,max}$, after which the data is discarded. Since we assume a fully connected network, i.e., no channel errors are introduced, data is considered successfully received at the BS if no collisions, even partial, occur. In case the sensor tries again to transmit a packet, it enters the BO state, and n_{sens} is reset.

It is worth mentioning that the time step duration is configured to allow sensors to receive updated CW values from the BS for the upcoming data transmission. Subsequently, if data are generated, sensors perform channel sensing, transmit the acquired data to the BS, and receive the corresponding ACK.

5.2.2 Reinforcement Learning Model

With the aim of properly selecting the CW that each sensor has to use to access the network, while satisfying the requirements imposed by each sensor's priority state (i.e., its goal), we exploit a centralized RL model developed at the BS.

Background on the Chosen RL Model RL algorithms enable an agent to learn task execution through repeated trial-and-error interactions within a dynamic environment [79, 90, 99], where the agent acts as a decision-maker and the environment encompasses external uncontrollable factors (Figure 38). The agent's objective is to optimize its decision-making policy to achieve a specific goal. The selected RL algorithm is the Advantage Actor-Critic (A2C) [100], chosen for its particular architecture, which effectively reduces variance by stabilizing and fastening the training process, and thereby this feature is a good fit for our multi-goal and dynamic scenario. A2C is an on-policy DRL algorithm that uses a synchronous gradient descent to optimize a DNN. It consists of two sub-networks: the actor and the critic. By integrating both policy-based and value-based methods, A2C facilitates more effective exploration of the environment.

The learning process along an episode evolves as follows. At each time step t, the agent is in an initial state S_t and explores the surrounding environment. Given the initial state S_t, it performs an action A_t according to its policy $\pi_\theta(S_t, A_t)$, which is the mapping from states to actions. The action execution leads to a transition to a new state S_{t+1}. The critic observes the action and provides feedback, developing the advantage function. The advantage function quantifies the benefit obtained by taking that action in that state with respect to the average value of the state. It is formulated as $A(S_t, A_t) = Q(S_t, A_t) - V(S_t)$, which is typically approximated as $A(S_t, A_t) = R_t + \gamma V(S_{t+1}) - V(S_t)$, with $Q(S_t, A_t)$ being the action-value function, $V(S_t)$ the state-value function, R_t the reward at current time step t, and $\gamma \in [0, 1)$ the discount factor, a parameter balancing the importance of immediate and future rewards. The advantage function is used by the actor to update its policy parameters, promoting training stability. The actor's parameters are updated as follows: $\Delta\theta = \eta \nabla_\theta (\log \pi_\theta(S_t, A_t)) A(S_t, A_t)$,

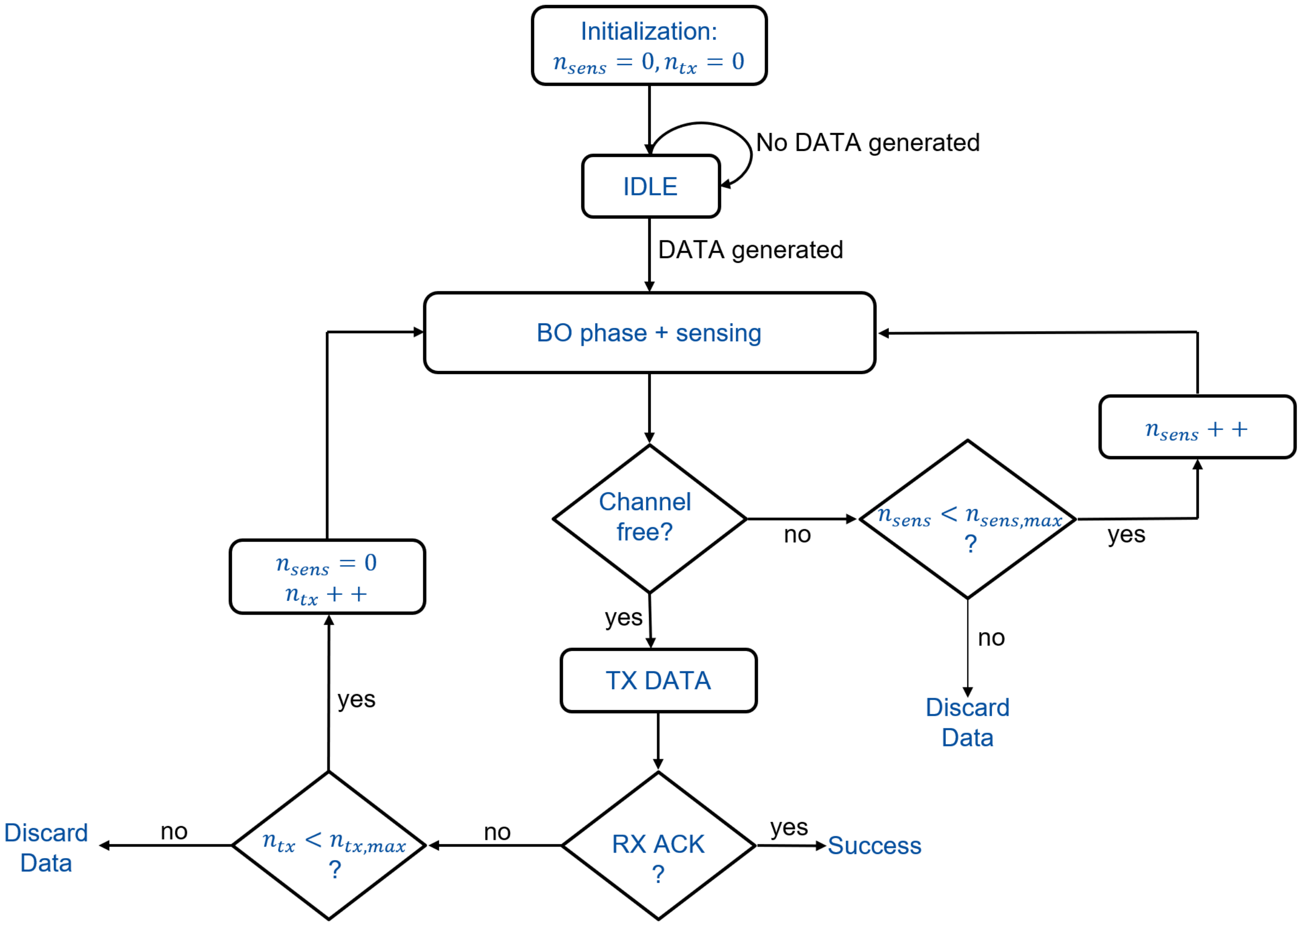

Figure 37: CSMA flow diagram.

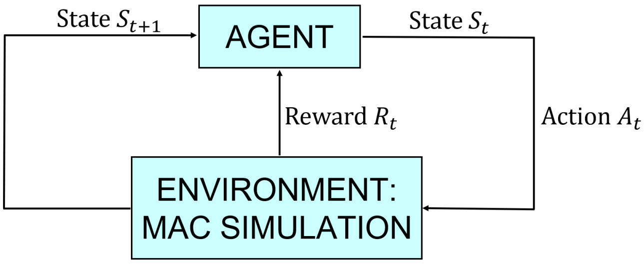

Figure 38: Illustrative block diagram of the RL operation.

with η being the learning rate. Those of the critic are updated exploiting the following formulation $\Delta\phi = \eta A(S_t, A_t)\nabla_\phi V(S_t)$. Then, the agent obtains from the environment a reward R_t, which rates the goodness of the action performed. The learning process is generally designed using a MDP, described by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R})$.

Problem Formulation We develop a centralized RL model, where the A2C agent, comprising both the actor and the critic, is located at the BS and the environment comprises the industrial plant involving machines and sensors. The MDP is formulated as follows:

- $\mathcal{S}$ is the state space, and represents the set of all possible states $S_t \in \mathcal{S}$ in which the agent can be. It is defined by $S_t = \{\mathbf{A}_{\mathbf{t-1}}, \mathbf{p}_{\mathbf{t-1}}, n_{sat}, n_{rx}\}$. In the following, we detail each entry of the state space. First,

$$R_{t,c} = \begin{cases} 1 & \text{if } n_{rx,c} \geq \overline{n_{rx,c}} \wedge n_{sat,c} \geq \overline{n_{sat,c}} \\ 0.5 & \text{if } n_{rx,c} < \overline{n_{rx,c}} \wedge n_{sat,c} \geq \overline{n_{sat,c}} \wedge n_{rx,c} \neq 0 \\ 0.25 & \text{if } n_{rx,c} \geq \overline{n_{rx,c}} \wedge n_{sat,c} < \overline{n_{sat,c}} \wedge n_{sat,c} \neq 0 \\ 0 & \text{if } n_{rx,c} < \overline{n_{rx,c}} \wedge n_{sat,c} < \overline{n_{sat,c}} \vee n_{sat,c} = 0 \vee n_{rx,c} = 0 \end{cases} \tag{3}$$

$\mathbf{A_{t-1}} = (A_{t-1,1}, A_{t-1,2}, ..., A_{t-1,C})$ represents the vector of C actions selected by the agent (BS) in the previous time step $t-1$, meaning the CW assigned to each sensor, with C being the number of sensors in the scenario. $\mathbf{p_{t-1}} = (p_{t-1,1}, p_{t-1,2}, ..., p_{t-1,C})$ is a vector of C booleans indicating the priority level of the sensors in the previous time step $t-1$, with value 1 for HP and 0 for LP sensors. Then, n_{sat} represents the number of packets successfully received at the BS that satisfies the latency requirements (i.e., $\tau_{t,c} \leq \tau_{th,hp}$ if sensor c is in HP and $\tau_{t,c} \leq \tau_{th,lp}$ if sensor c is in LP) calculated over the previous k transmissions. Finally, n_{rx} indicates the number of packets successfully received at the BS calculated over the previous k transmissions.

- $\mathcal{A}$ is the action space, and represents the set of all possible actions $\mathbf{A_t} \in \mathcal{A}$ that the agent can perform. It is defined as $\mathbf{A_t} = (A_{t,1}, A_{t,2}, ..., A_{t,C})$ with the generic action $A_{t,c}$, performed at time step t and referred to the sensor c, being $A_{t,c} \in \{W_1, W_2, ..., W_D\}$. The generic action W_i is selected from a set of D actions ranging from W_1 to W_D with a step of W_{step}. Each action corresponds with the choice of a CW with duration $t_{tick} \cdot W_i$, on which randomly chose a BO time that each sensor has to wait before accessing the network.
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \to \mathcal{S} \times (0, 1)$ is the transition probability function. It models the environment's dynamics by returning the probability of transitioning to a new state S_{t+1} given the current state S_t and action A_t. In an MDP, the probability of transitioning to the next state depends only on the current state.
- $R_t \in \mathcal{R}$ is the reward function with $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}$. It is a critical component for the learning process, offering intermediate feedback to the agent [90, 101]. The reward obtained at the generic time step t is the sum of the rewards achieved by each sensor c, therefore:

$$R_t = \sum_{c=1}^{C} R_{t,c}, \tag{149}$$

where the reward $R_{t,c}$ of the generic sensor is specifically designed to account for the historical evolution of the sensor's priority state and requirement. It is computed in a step-wise fashion and is formally defined in Eq. 3.

Specifically, $n_{rx,c}$ denotes the number of packets successfully received by the BS from sensor c, computed over the last k transmissions. In contrast, $\overline{n_{rx,c}}$ represents the average number of successful receptions, averaged over each k transmissions, carried out by sensor c along the actual episode. $n_{sat,c}$ denotes the number of packets transmitted by sensor c and successfully received at the BS that also satisfy the latency requirements defined by the sensor's current priority level (HP or LP), computed over the last k transmissions. Conversely, $\overline{n_{sat,c}}$ represents the number of such packets, averaged over each k transmissions, performed by sensor c along the actual episode.

The goal of the RL algorithm is to identify the optimal policy $\pi_\theta(S_t, A_t)$ which maximizes the cumulative discounted return at time step t, that is, $G_t = R_t + \gamma R_{t+1} + \gamma^2 R_{t+2} + \cdots + \gamma^{N-t} R_N$.

According to the defined MDP, at each time step t, the BS observes the current state and selects one action per sensor, specifying the CW to be used. This information is then communicated to each sensor for use in the upcoming data transmission. Subsequently, each sensor follows the CSMA MAC protocol described in Section 5.2.1: it senses the wireless channel and, if appropriate, transmits the acquired measurement data to the BS, using the assigned CW, before waiting for the reception of the corresponding ACK.

5.2.3 Performance Evaluation

Performance Metrics The performance metrics used to validate the proposed RL-based model include average latency and success probability.

- *Average Latency* is defined as the interval of time between the generation of the first bit of a packet (on the sensor side) and the reception of the last bit of the corresponding ACK (on the sensor side). It is formally computed as:

$$L = \frac{1}{N}\sum_{t=1}^{N} L_t = \frac{1}{N}\sum_{t=1}^{N}\frac{1}{C}\sum_{c=1}^{C} L_{t,c}, \tag{4}$$

  where $L_{t,c}$ is the latency of each packet transmitted by each sensor c, averaged over the total number of time steps composing an episode, N, and the total number of sensors, C.

- *Success Probability* is the probability that a packet sent by a sensor is successfully received at the BS. It is formulated as:

$$p_s = \frac{1}{N}\sum_{t=1}^{N} p_{s,t} = \frac{1}{N}\sum_{t=1}^{N}\frac{N_{ack,t}}{N_{data,t}}, \tag{5}$$

  where $N_{ack,t}$ and $N_{data,t}$ represent the number of ACKs received and the number of data transmitted during the total number of N time steps composing an episode, respectively.

Numerical Results Numerical results are obtained using the set of parameters reported in Table 10. The A2C algorithm is implemented exploiting the Stable-baselines3 library [102]. Both the actor and the critic Neural Networks (NNs) receive as inputs the observation space. These observational inputs are processed by a dedicated feature extractor and passed to two separate fully connected NNs, comprising a linear layer and two hidden layers consisting of 64 neurons each for policy and value estimation. The output of the actor includes $C \cdot D$ neurons, whereas the critic has a single neuron responsible for value function estimation. The proposed RL-based MAC protocol, which dynamically adapts the CWs to the different multi-goal traffic sources, is compared to a benchmark in which a fixed CW is assigned to all sensors, i.e., W_i, set to $W_i = 10, 250, 500$. This ensures a fair comparison, since our proposal selects the CW from a set ranging from $W_1 = 10$ to $W_D = 500$.

Figs. 39(a), 39(b) illustrate the average latency, L, as a function of the number of sensors operating in HP and LP modes, respectively. Similarly, Figs. 39(c), 39(d) illustrate the success probability, p_s, as a function of the number of sensors working in HP and LP modes, respectively. For all plots, we set $P_{data} = 60$ B. It is worth noting that both metrics are averaged according to the percentage of time each sensor belongs to a given priority level (i.e., either HP or LP), in accordance with the multi-goal traffic model described in Sec. 5.2.1. The horizontal lines in the plots represent the performance requirements for HP and LP traffics, which takes inspiration from typical industrial scenarios [88]. In particular:

- $\tau_{th,hp} = 1 \cdot 10^{-6}$ s is the latency requirement for HP sensors;
- $\tau_{th,lp} = 2 \cdot 10^{-6}$ s is the latency requirement for LP sensors (reflecting a lower value w.r.t the HP traffic, in agreement with what described in Section 5.2.1);
- $p_{s,th} = 0.9$ is the success probability requirement for both HP and LP sensors.

By looking at Figure 39(a), as expected, one can notice that the average latency increases with the number of sensors, due to a higher probability of packet collisions. Most importantly, it can be seen that the proposed RL-based CSMA model with adaptive CW selection significantly outperforms the fixed CW versions, achieving a lower average latency while meeting the imposed latency requirement. Specifically, the RL-based model reduces average latency (averaged over $C = 10, 20, 30$) by 16.36% compared to the fixed $W_i = 250$ configuration, and by 34.97% relative to the fixed $W_i = 500$ configuration. However, the average latency remains 41.68% higher than that achieved with the fixed $W_i = 10$ configuration, which corresponds to the smallest possible

Table 10: Simulation Parameters

Parameter Type	Parameter	Notation	Value
Scenario	Industrial plant size	$F_x \times F_y \times F_z$	$40 \times 40 \times 10$ m^3
	Machines size	$M_x \times M_y \times M_z$	$3 \times 3 \times 3$ m^3
	Number of machines	M	12
	Number of sensors	C	{10, 20, 30}
Radio Configuration	Carrier frequency	f_c	300 GHz
	Bit rate	R_b	50 Gbit/s
MAC	ACK size	P_{ack}	10 byte
	Data size	P_{data}	{20, 40, 60} byte
	Number of data generated by each sensor	N	400
	Total number of CWs	D	50
	Value of the CW W_1	W_1	10
	Value of the CW W_D	W_D	500
	Difference in the value of two adjacent CWs	W_{step}	10
	Tick duration	t_{tick}	1.6 ns
	Sensors periodicity	N_{tick}	10.000
	Number of times the channel is sensed	$n_{sens,max}$	3
	Number of transmission attempts	$n_{tx,max}$	2
	Number of transmissions used for computing $\overline{n_{rx,c}}$ and $\overline{n_{sat,c}}$	k	5
RL	Neural network architecture	NN	64×64
	Learning rate	η	0.0005
	Discount factor	γ	0.9
	Number of simulated episodes	N_{ep}	5.000
Application Requirements	Success probability	$p_{s,th}$	0.9
	HP latency	$\tau_{th,hp}$	$1 \cdot 10^{-6}$ s
	LP latency	$\tau_{th,lp}$	$2 \cdot 10^{-6}$ s

CW. At first glance, this result may suggest that the optimal approach is to statically assign a CW of 10 to all sensors. Nevertheless, this conclusion does not hold when the success probability is also taken into account. By observing both Figure 39(a) and 39(c), it can be seen that the RL-based model is able to jointly meet the latency and success probability requirement across all device counts, whereas a static CW assignment with $W_i = 10$ fails to do so as it never satisfies the success probability requirement. By comparing the two approaches, the difference is quite substantial. For instance, with $C = 30$ devices, the improvement in success probability reaches up to 71.93%.

From Figure 39(b) emerges that the RL-based model provides a higher average latency compared to two fixed CW configurations, e.g., $W_i = 10$ and $W_i = 250$, but is still capable of satisfying the LP latency requirements. Specifically, the average latency obtained with the RL model is 56.62% higher than that achieved with a fixed $W_i = 10$ configuration and 6.63% higher than that achieved with a fixed $W_i = 250$ configuration. Moreover, the average latency is 17.33% smaller than that of the fixed $W_i = 500$ configuration. This outcome highlights the RL model's ability to differentiate service levels by allocating larger CWs to LP sensors, given their more relaxed latency requirements, and shorter CWs to HP sensors, to meet the respective latency constraints and mitigate collision occurrences.

Finally, Figure 39(d) confirms the superior performance of the proposed RL-based MAC protocol over fixed CW configurations for LP sensors. When $C = 30$, the model achieves an improvement in success probability of up to 75.52%, compared to the fixed $W_i = 10$ configuration. Moreover, the RL model successfully meets the LP latency requirements while ensuring a minimum success probability of 97.78%. These results prove the robustness of the RL-based model, which, regardless of the number of sensors, is able to guarantee a high success probability.

To summarise, the results reveal important limitations associated with the use of fixed CW configurations. Specifically, setting $W_i = 10$ results in minimal latency, but significantly compromises the success probability, making it unsuitable for reliable communication since it does not satisfy the success probability requirements. Conversely, $W_i = 500$ ensures a high success probability but incurs excessive latency for dense networks (i.e.,

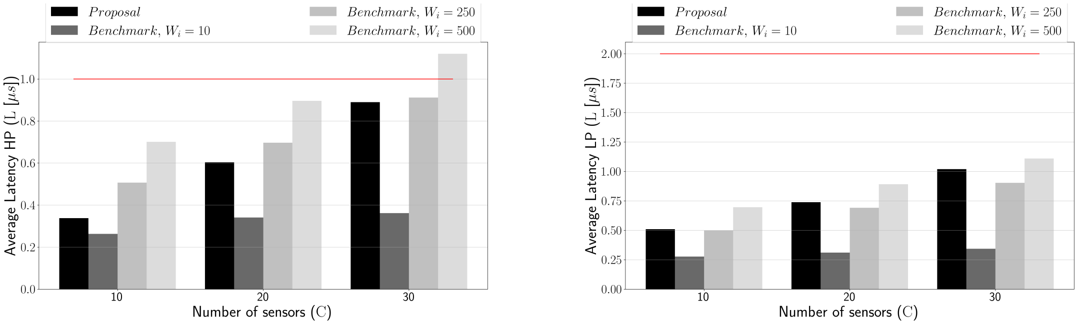

(a) Latency performance of HP sensors (b) Latency performance of LP sensors

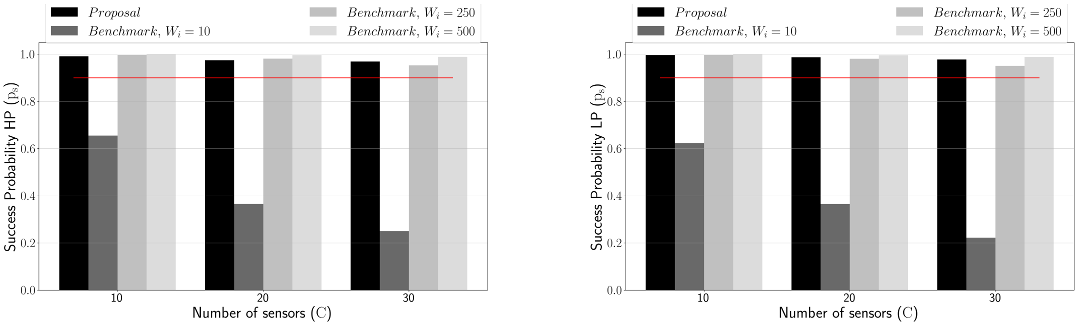

(c) Success probability of HP sensors (d) Success probability of LP sensors

Figure 39: Average latency (L) and success probability (p_s) as a function of the number of HP and LP sensors with the proposed RL-based MAC protocol and fixed CW. The horizontal lines represent the performance requirements for HP and LP traffics. We set $P_{data} = 60$ B.

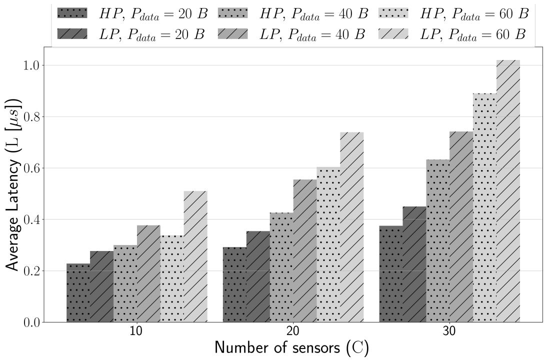

Figure 40: Average latency (L) as a function of the packet size (P_{data}) for HP and LP sensors with the proposed RL-based MAC protocol.

for $C = 30$), failing to meet latency constraints. An intermediate configuration, $W_i = 250$, not only results in intermediate values of latency and success probability, but also requires numerous simulation trials before being selected as the optimal value. In contrast, the proposed model dynamically adapts the CW selection in one shot, effectively identifying the optimal configuration. This adaptive strategy achieves a balanced compromise between latency reduction and success probability maximization, while simultaneously meeting heterogeneous and evolving sensors' priority requirements typical of multi-goal scenarios.

Finally, the adaptability of the RL model to varying packet sizes is examined. Figure 40 illustrates the average latency experienced by HP and LP sensors as a function of the number of sensors and for different data sizes (e.g., $P_{data} = 20, 40, 60$ B). As expected, the average latency increases with both the number of sensors and the packet size, due to increased collisions due to denser traffic. Notably, in all cases, the average latency experienced by HP sensors is always lower with respect to that of LP sensors, demonstrating the RL model's ability to effectively differentiate sensors depending on their priority level (i.e., multi-goal setting) and assign the appropriate CW accordingly. The average latency discrepancy between HP and LP sensors (averaged over $C = 10, 20, 30$) is 17.29% for $P_{data} = 20$ B, 19.45% for $P_{data} = 40$ B, and 21.58% for $P_{data} = 60$ B. This means that a larger packet size also leads to an increased latency discrepancy between HP and LP sensors. The obtained results confirm the model's effectiveness in satisfying diverse and evolving requirements within a multi-goal wireless network environment.

5.3 Predictive network management and RIS configuration via radar-based sensing

The adoption of RIS technology introduces significant challenges, particularly regarding real-time configuration and management. This is particularly true in IIoT scenarios, where communication disruptions caused by the latency between blockage detection and subsequent RIS and BS reconfiguration are often unacceptable. While increasing the localization rate can help reduce the duration of such outages, this approach is not always feasible or cost-effective.

To fill this gap, in this section we propose a novel approach that exploits radar-based localization data to proactively configure both the BS and the RISs, ensuring reliable communication links at THz frequencies despite blockages caused by moving objects. Specifically, our approach predicts the blockage event before it occurs, thus giving time to the system to react properly and allowing for a reduction of localization rate. We validate our solution through network-level simulations of uplink data communications between User Equipments (UEs) and the BS in a dynamic industrial scenario.

5.3.1 Scenario

Although the proposed approach is general, for the sake of explanation and performance assessment, we refer to a specific industrial scenario. It consists of a 2D industrial environment measuring 9m × 13m, whose layout is depicted on the left-hand side of Figure 41. We assume that broadband cellular connectivity is provided in this area by a BS operating at 300 GHz, serving two stationary UEs, namely UE0 and UE1. These devices may represent, for instance, sensing units collecting data to be forwarded to a control center monitoring ongoing industrial processes.

Two moving machines (e.g., pallet conveyors) are also considered, represented by dark grey boxes in Figure 41. These machines follow a straight trajectory from the top to the bottom of the light grey area highlighted in the same figure, referred to as the machine zone, which measures 3m × 5m. With this setup, UE0 always maintains a LoS connection with the BS, whereas UE1 initially has a LoS link but later experiences blockages due to the movement of the machines.

To avoid connectivity failures, two RISs are strategically positioned to maintain LoS links with the BS at all times, while also providing LoS coverage to areas that may be shadowed by the moving machinery, ensuring continuous connectivity between the BS and UE1 when the direct link is unavailable.

Figure 41 also highlights, using distinct colors, the area where a LoS link can always be established between the BS and the UEs, as well as the area where this condition might be lost, potentially requiring the activation of the most suitable RIS among the two positioned in the top-right and bottom-right corners.

5.3.2 Radar-based localization at THz frequencies

To acquire information on the industrial environment, the BS employs a Frequency-Modulated Continuous-Wave (FMCW) radar with a single transmit antenna and M_r receive antennas, operating at a carrier frequency f_c. The radar periodically transmits a burst of chirp signals to sense the surrounding environment. The received signal is mixed with the transmitted signal to produce an Intermediate Frequency (IF) signal containing frequency and phase differences. The data cube $\mathbf{X} \in \mathbb{C}^{S \times M_r \times N}$ is a three-dimensional tensor representing raw radar data: S is the number of samples per chirp, M_r is the number of receive antennas, and N is the number of chirps per frame. It is obtained by sampling the IF signal at rate f_s with an analog-to-digital converter to get S samples per chirp.

From the raw radar data $\mathbf{X} \in \mathbb{C}^{S \times M_r \times N}$, a FFT is performed across all dimensions of $\mathbf{X}$. The FFT along the first dimension (fast-time) transforms the IF signal into the frequency domain, yielding the beat frequency $f_b = \mu\tau$, where $\mu = B/T_c$ is the chirp slope, B is the bandwidth, T_c is the chirp duration, and $\tau = 2r/c$ is the round-trip delay for a target at range r, with c being the speed of light. The range is then computed as $r = cf_b/(2\mu)$. Next,

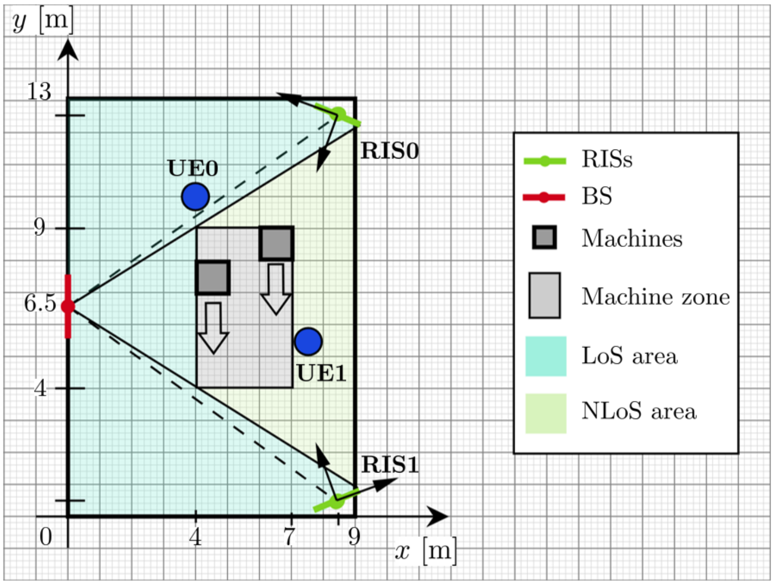

Figure 41: The considered industrial scenario.

an FFT across the chirp dimension (slow-time) resolves the Doppler shift f_d, enabling radial velocity calculation via $v = cf_d/(2f_c)$, where f_c is the carrier frequency. Finally, an FFT across the antenna dimension extracts phase differences, providing azimuth angular information. Then, range-Doppler and range-azimuth maps can be constructed for target localization.

A distance-dependent threshold is applied to the radar maps to produce a detection point cloud. This is achieved using an exponential decay threshold defined as $\gamma = \gamma_0 e^{-\alpha r}$, where γ_0 is the initial threshold value, r is the target range, and α is a tunable parameter to adjust the threshold's sensitivity to distance. Range-azimuth and range-Doppler point clouds are merged based on common range, ensuring target distinction when they are close in angle but differ in velocity, or vice versa. We employ Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [103] to cluster targets by density. DBSCAN uses two main parameters: *epsilon* and *min_pts*. *Min_pts* sets the minimum number of detected points to declare a valid cluster.

A Kalman filter is applied to track the positions of detected targets over time. The Kalman filter, implemented with a state vector $[x, y, v_x, v_y]^T$, is initialized with positions from the first radar frame and velocities derived from positional differences. The state transition matrix F incorporates a finer time step than the radar update rate. By denoting I_n as a $n \times n$ identity matrix, the measurement matrix H observes the position, with process noise covariance $Q = \sigma_q I_4$ and measurement noise covariance $R = \sigma_r I_2$, where σ_q and σ_r are tuned parameters. The filter predicts positions at each fine step, updating every radar measurement with radar centroids. The resulting fine-grained data include positions, velocities, and sizes for each detected target. Additionally, a LoS blockage check is conducted at each step, leveraging bounding box vertices of detected machines to evaluate obstructions between the BS and each UE, updating the blockage condition accordingly.

5.3.3 RIS-based Communication at THz Frequencies

In this section, we describe how we model the radio channel between K single-antenna UEs, M RISs comprising $N_\mathrm{m} = N_\mathrm{m,x} \times N_\mathrm{m,y}$ elements, and a BS equipped with an antenna array of N_r elements.

UE-BS radio channel modeling: The channel vector $\mathbf{h}_\mathrm{k} \in \mathbb{C}^{N_\mathrm{r} \times 1}$ for the k-th UE-BS link considers free space propagation between each transmitting and receiving antenna element. This assumption is justified for THz communications, where highly selective beamforming in LoS conditions makes the free-space model a valid

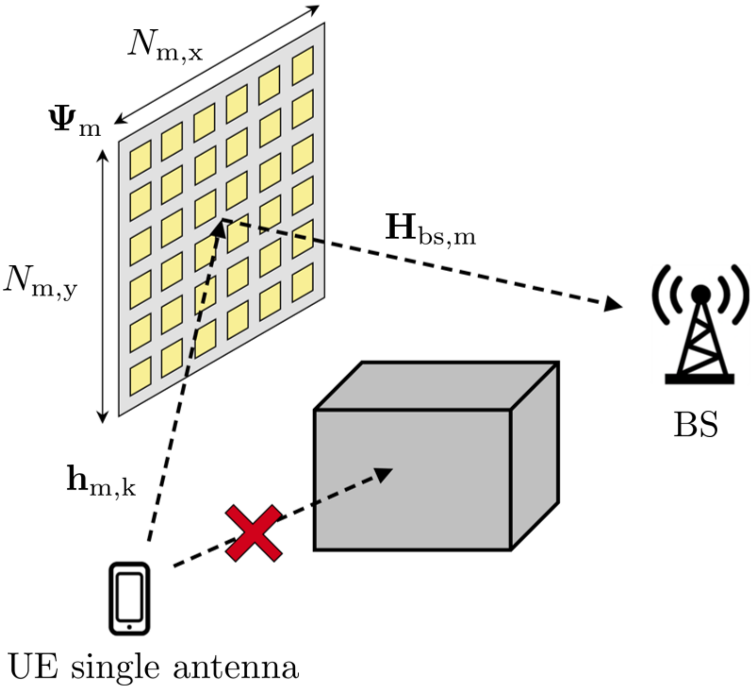

Figure 42: Representation of the UE-RIS-BS link in case of a blockage event.

approximation. The generic element of the channel vector is thus defined as:

$$h_{\mathrm{k,n}} = \sqrt{G_{\mathrm{k}} G_{\mathrm{n}}} \frac{\lambda}{4\pi d_{\mathrm{k,n}}} e^{-\mathrm{j}\frac{2\pi}{\lambda} d_{\mathrm{k,n}}}, \tag{6}$$

with $h_{\mathrm{k,n}}$ representing the channel coefficient between the k-th UE and the n-th BS antenna element, $d_{\mathrm{k,n}}$ denoting the Euclidean distance, λ is the wavelength, G_{k} and G_{n} are the gain of the elements of the UE and BS, respectively.

As described in Section 5.3.1, we consider a scenario with moving machines that may introduce blockages between the transmitter and receiver. In our channel model, we assume that the signal is completely obstructed by these obstacles, resulting in a null $\mathbf{h}_{\mathrm{k}}$ when the direct path is blocked.

UE-RIS-BS radio channel modeling: The channel vector $\mathbf{h}_{\mathrm{k}}$ that characterizes the UE-RIS-BS link for the k-th UE and the m-th RIS ($m \in \{1, ..., M\}$) is expressed as follows:

$$\mathbf{h}_{\mathrm{k}} = \mathbf{H}_{\mathrm{bs,m}} \mathbf{\Psi}_{\mathrm{m}} \mathbf{h}_{\mathrm{m,k}}, \tag{7}$$

where $\mathbf{h}_{\mathrm{m,k}} \in \mathbb{C}^{N_{\mathrm{m}} \times 1}$ is the channel vector between the k-th UE and the m-th RIS, $\mathbf{H}_{\mathrm{bs,m}} \in \mathbb{C}^{N_{\mathrm{r}} \times N_{\mathrm{m}}}$ is the channel matrix between the BS and the RIS, while $\mathbf{\Psi}_{\mathrm{m}} \in \mathbb{C}^{N_{\mathrm{m}} \times N_{\mathrm{m}}}$ denotes the scattering matrix of the m-th RIS according to the model described in Section 3.3.1. Specifically, $\mathbf{\Psi}_{\mathrm{m}}$ can be written as:

$$\mathbf{\Psi}_{\mathrm{m}} = \mathrm{diag}(e^{-\mathrm{j}\psi_{\mathrm{m},1}}, e^{-\mathrm{j}\psi_{\mathrm{m},2}}, ..., e^{-\mathrm{j}\psi_{\mathrm{m},N_{\mathrm{m}}}}), \tag{8}$$

where $\psi_{\mathrm{m}} \triangleq (\psi_1, \psi_2, ..., \psi_{N_{\mathrm{m}}})$ is the vector of the phase shifts. Each RIS is assigned with a codebook of vectors $\psi_{\mathrm{m}}(\Theta_{\mathrm{inc}}, \Theta_{\mathrm{ref}})$, where the incident angle is $\Theta_{\mathrm{inc}} = (\theta_{\mathrm{inc}}, \phi_{\mathrm{inc}})$ and the reflective angle is $\Theta_{\mathrm{ref}} = (\theta_{\mathrm{ref}}, \phi_{\mathrm{ref}})$. Assuming a fixed Θ_{ref} for the RIS-BS link and dividing both the elevation and azimuth planes into 90 discrete angles, a codebook of 90×90 phase shift vectors can be computed [26] for each possible Θ_{inc} in the UE-RIS link. We say that the RIS is *activated* when one of the possible ψ_{m} is chosen.

Based on the relative positions of UEs, RISs, obstacles, and the BS, the radio channel vector $\mathbf{h}_{\mathrm{k}}$, which models the direct or indirect link between the k-th UE and the BS, can be classified into three distinct cases:

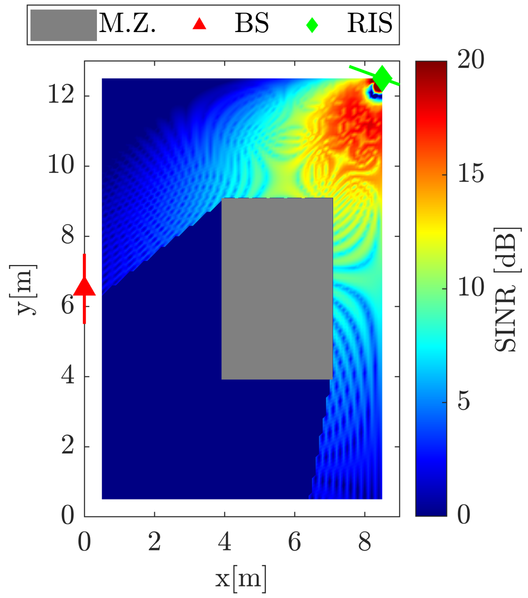

Figure 43: Exemplary SINR distribution for the UE-RIS-BS links obtained with a custom Matlab simulator, when considering the upper RIS.

- Direct link only: A single LoS path exists between the UE and the BS. In this case, each element of $\mathbf{h}_k$ follows (6);
- Indirect link only: Due to obstructions, only a indirect UE-RIS-BS path is available if the RIS is properly configured (as illustrated in Figure 42), in which case $\mathbf{h}_k$ follows (7);
- Combined direct and indirect links: Both direct and indirect paths are present, resulting in an overall channel vector given by the sum of the respective components.

Given $\mathbf{h}_k$, it is possible to compute the SINR for each $k \in \{1, ..., K\}$, where we assume the Maximum Ratio Combining (MRC) technique at the receiver. In particular, the expression of SINR_k is as follows:

$$\text{SINR}_k = \frac{\|\mathbf{h}_k\|^2 P_k}{\sigma_w^2 + I}, \tag{9}$$

where the numerator represents the useful received power from the k-th UE that transmits with power P_k, σ_w^2 is the thermal noise power, and the interference power I is given by

$$I = \sum_{i \neq k}^{K} \left| \frac{\mathbf{h}_k}{\|\mathbf{h}_k\|} \mathbf{h}_i \right|^2 P_i, \tag{10}$$

where P_i is the transmit power of the i-th interfering UEs, $\|\mathbf{h}_k\|$ denotes the norm of the channel vector $\mathbf{h}_k$, while $|\cdot|^2$ represents the square modulus; under the assumption that both noise and interference follow a Gaussian distribution.

Figure 43 illustrates an example of the uplink SINR spatial distribution for a single UE-RIS-BS link within the considered scenario (see Section 5.3.1) with parameters listed in Table 11. For simplicity and to reduce execution time, the figure was generated using a custom Matlab simulator, which models only the two-ray radio channel described, instead of relying on the complex and more accurate tools presented in Section 3.2.4. The results show that, with an appropriate configuration of $\mathbf{\Psi}_m$, the coverage can be significantly extended behind the machine zone, achieving an SINR range of 8–15 dB within a 3m^2 square region where the RIS is positioned at the top-right corner (see the red–yellow–green area near the RIS in Figure 43).

5.3.4 Proposed Approach

In this section, we present the proposed algorithm for predictive network management at THz frequencies, where both the BS and RISs are proactively configured based on localization information provided by the radar mounted on the BS.

Specifically, with a periodicity of τ, the radar provides estimates of the coordinates and dimensions of detected targets, which are subsequently classified as either machines or UEs based on their respective sizes: if the largest dimension is less than 1 m, the target is classified as a UE; otherwise, it is identified as a machine. Using this estimated scenario information, the MRC precoding vector at the BS for a k-th UE can be computed as

$$\mathbf{v}_{\mathrm{k}} = \frac{\mathbf{h}_{\mathrm{k}}^{\mathrm{H}}}{\|\mathbf{h}_{\mathrm{k}}\|}, \tag{11}$$

with $\mathbf{h}_{\mathrm{k}}$ given by (7) by assuming perfect CSI estimation. Additionally, based on the estimated positions of the targets, we leverage a predictive algorithm to forecast the occurrence of potential blockages, i.e., to estimate whether an obstruction will occur between the BS and any UEs within two radar updates. Specifically, each Kalman filter (see Section 5.3.2) updates the state at each radar update, while in between two measurements, it predicts the state using a constant velocity model. Then, a linear interpolator considers the past five time instances to predict the machine's position at the next step, enabling the verification of potential blockages. Then, based on the current and predicted positions, the possibility of blockage events is assessed for each UE-machine pair. If a blockage is predicted, the network proactively activates the appropriate RIS[9] by selecting the proper phase shift vector $\psi_m(\Theta_{\mathrm{inc}}, \Theta_{\mathrm{ref}})$, to ensure stable communication for UEs that are expected to be obstructed before the next radar measurement. Since we consider only two RISs and assume that each RIS can serve at most one UE at a time, we adopt a greedy approach, prioritizing the UEs closest to the RISs. Similarly, the BS's multiple antenna is configured to direct its beam toward the newly activated RIS.

Based on this proactive configuration of both the BS and RISs, UEs perform uplink transmissions to the BS following an Unslotted Aloha MAC layer protocol accounting for the propagation delay. In fact, at THz frequencies the propagation delay is typically in the same order of magnitude as the transmission time thus making useless any MAC relying on carrier sensing.

In particular, the MAC layer protocol works as follows. Whenever an UE has new DATA to be transmitted, it initiates a BO phase for $T_{\mathrm{B}} = T_{\mathrm{BO}}\,\xi$, where T_{BO} is the minimum BO duration; ξ is a uniform random number in the range $[1; 2^i C]$, with C being an integer defining the duration of the contention window and i being an integer counting the number of transmission attempts (starting at 0 for the first attempt). At the end of the BO period, the UE sends the DATA and enters reception mode for a maximum period, T_{WAIT}. Specifically, we set $T_{\mathrm{WAIT}} = T_{\mathrm{ACK}} + 2\,\tau_{\mathrm{p_{max}}}$, since it is the time needed to receive the corresponding ACK, where $T_{\mathrm{ACK}} = \frac{8\,P_{\mathrm{A}}}{R_{\mathrm{b}}}$, P_{A} is the number of bits forming an ACK, $R_{\mathrm{b}} = B\log_2(M)$ is the bit rate (where we assume that the symbol rate equals the bandwidth), M is the modulation order considering an M-QAM modulation scheme, and $\tau_{\mathrm{p_{max}}}$ is the maximum possible propagation delay in the considered IIoT scenario. If an ACK is received during T_{WAIT}, the UE goes back to BO, as we consider the worst-case condition where all UEs have always a new DATA to be transmitted. Conversely, if no ACK is received during T_{WAIT}, the UE retries the transmission up to a maximum number of attempts, R.

The successful reception of a DATA or ACK transmission depends on the radio channel quality. For the k-th link, this quality is measured in terms of SINR_k, which is computed as in (9), based on the proactive configuration of the BS and RISs, updated every τ seconds. In this way, given a specific error correction code and coding rate $R_{\mathbf{c}}$, SINR_k is mapped to a packet error probability P_w[10]. A transmission is thereby considered successfully received if a randomly generated number between 0 and 1 exceeds P_w; otherwise, it is deemed lost.

Finally, to evaluate the overall network performance of our solution, we use the network throughput, defined as:

---

[9] For the sake of simplicity, we assume an ideal and instantaneous communication channel between the BS and the RISs, as well as that the network has a-priori knowledge of the true coordinates of the RISs.

[10] For simplicity, we assume that each transmission consists of a single codeword.

Table 11: Simulation parameters.

Parameter	Symbol	Value
Carrier frequency	f_c	300 GHz
Bandwidth	B	3 GHz
BS antenna element gain	G_n	0 dB
BS noise figure	N_F	8 dB
BS antenna elements	N_r	256×256
UE antenna element gain	G_k	0 dB
Transmit Power	P_t	0 dBm
RIS antenna elements	N_ris	128×128
Modulation order	M	4
Coding Rate (BCH code)	R_c	0.9
DATA size	P	20 byte
Minimum Back-Off time	T_BO	13.3 ns
Simulation time	T_S	1 ms
ACK size	P_A	10 byte
Maximum propagation delay	τ_{Pmax}	72.7 ns
Contention window	C	5
Maximum retransmission attempts	R	3

$$S = \frac{8PN_\mathrm{R}}{T_\mathrm{S}}, \tag{12}$$

where P is the number of bytes forming the DATA, N_R is the number of DATA successfully received at the BS within the simulation time T_S.

5.3.5 Numerical Results

The proposed approach has been implemented in a custom Matlab-Python simulator, where Matlab is used to model the PHY-layer radio channel (as described in Section 5.3.3), and Python is used to simulate the proposed approach along with the corresponding benchmarks (see Section 5.3.4). Simulation parameters, if not otherwise specified, are listed in Table 11. Notably, to manage the huge number of RIS elements required to provide a reasonable link budget in non-line-of-sight (NLoS) condition, the simulator uses the large RIS approximation from [27], which significantly lowers the computational complexity when deriving the UE-RIS-BS channel. Additionally, each RIS element is modeled like a patch antenna, as detailed in [26]. Regarding the BS and the UEs, both employ omnidirectional antenna elements with 0 dB gain. However, the BS is equipped with an array of N_r elements, resulting in an overall array gain of 48 dB (very close to the 50 dB gain of the high-gain antenna produced by Anteral for WP5), whereas each UE is equipped with only a single antenna element.

To evaluate the overall effectiveness of our proposal, Figure 44 depicts the network throughput as a function of the radar update interval within a simulation time of $T_\mathrm{S} = 1$ ms. We compare our radar-based approach, applied using two configurations of receive antenna elements at the radar, namely $M_\mathrm{r} = 32$ and $M_\mathrm{r} = 128$, against an *oracle*-based approach that has perfect and continuous knowledge of obstacles and UEs positions. The performance of the latter, introduced here as a benchmark, has been evaluated with and without the presence of RISs in the scenario.

As expected, network throughput increases with more frequent radar updates, as this enhances the accuracy

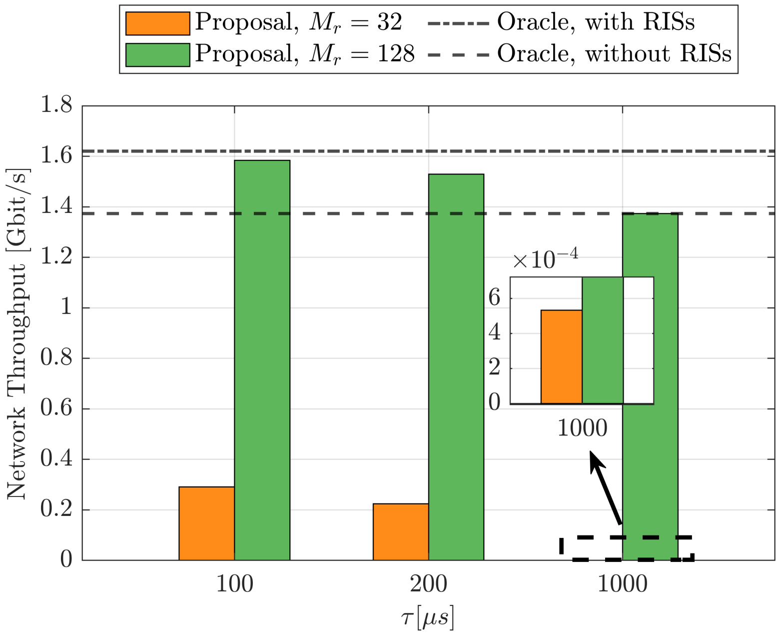

Figure 44: *Network throughput comparison between our solution and an ideal oracle which has perfect and continuous knowledge of obstacles and UEs positions, varying radar update periodicity (τ), number of radar receiving antennas (M_r), and RISs presence.*

of proactive RISs activation, thereby reducing losses during obstruction events. Additionally, the performance deteriorates significantly when the radar has an insufficient number of receiving antennas due to inaccurate position estimation (see the curves for $M_r = 32$ and $M_r = 128$). Indeed, the performance of our approach matches that of the oracle without RISs only by setting $M_r = 128$. With a single radar update, $\tau = 1$ ms, the network acquires the estimated map only at the start of the simulation, preventing timely RISs activation and thereby resulting in performance far below that of the oracle with RISs. Conversely, with five and ten radar updates, i.e., $\tau = 200$ and $\tau = 100$ μs, respectively, our solution is able to closely approximate the ideal oracle, even when the oracle employs RISs as soon as needed.

5.3.6 Performance Analysis with WP3 Channel Model

The aim of this section is to compare the numerical results obtained using the theoretical models for the communication channel and the RIS, described in the previous sections, with those obtained from a more realistic model obtained by raytracing outcomes including real measurements from the WP3/WP5 participants. To ensure a fair comparison, the same scenario (see Figure 41) was reproduced by WP3 members (TUBS in particular) and simulated using the SiMoNe ray-tracing tool (see Section 4.6), as also detailed in TIMES D2.2 [18]. Specifically, both the RIS model (see Section 3.3.1) and the BS configuration, originally based on channel vector computation and MRC at the receiver, both implemented in our custom Matlab-based PHY simulator, were replaced by the results of raytracing simulations including measurements of the RIS model provided by Anteral from WP5, and considering that the BS is equipped with the high-gain antenna that will be employed in the POCs of WP6.

Since deterministic ray-tracing methods are typically applied to static scenarios, the mobility of the machines was modeled by generating a sequence of scenario snapshots over time, each separated by a defined time step of 100 μs. For each snapshot, WP3 provided the path gain values for all UE–BS and UE–RIS–BS links, under both LoS and NLoS conditions. Since the UE transmit power is set to 0 dBm in this study (see Table 11), the received power, denoted as P_r, can be easily derived from the path gains according to the well-known link budget formula. Therefore, Tables 12 and 13 report a comparison of the P_r and SNR values at the BS, respectively, for the theoretical model used in the previous section (i.e., WP4-related) and the raytracing

Table 12: Comparison of the received power at the BS between the theoretical channel model used in our study (WP4-related) and the values obtained from ray-tracing (SiMoNe) analysis including models from real measurements (WP3-related), as a function of the visibility condition.

Link status	**WP3 model, P_r [dBm] from UE0**	**WP3 model, P_r [dBm] from UE1**	**WP4 model, P_r [dBm] from UE0**	**WP4 model, P_r [dBm] from UE1**
LoS	-44.45	-47.44	-48.33	-50.7
NLoS (RIS1)	n/a	-111.67	n/a	-59.1

Table 13: Comparison of the SNR at the BS between the theoretical channel model used in our study (WP4-related) and the values obtained from ray-tracing (SiMoNe) analysis including models from real measurements (WP3-related), as a function of the visibility condition.

Link status	**WP3 model,** SNR **[dB] from UE0**	**WP3 model,** SNR **[dB] from UE1**	**WP4 model,** SINR **[dBm] from UE0**	**WP4 model,** SINR **[dB] from UE1**
LoS	24.19	20.75	19.85	17.47
NLoS (RIS1)	n/a	-43.47	n/a	9.09

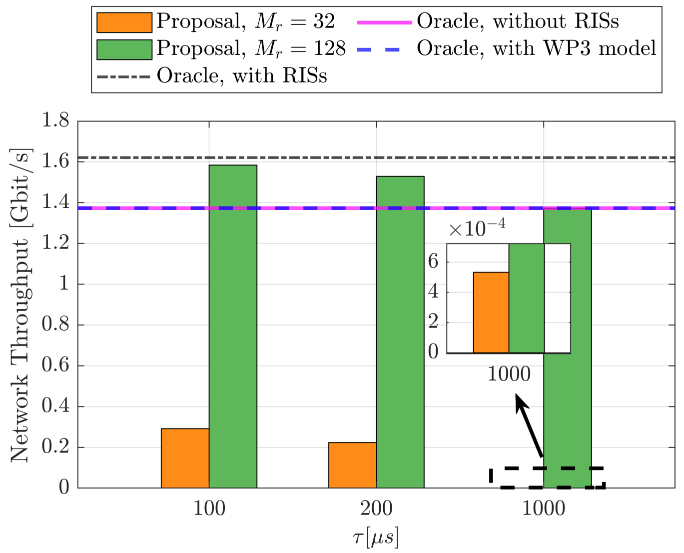

Figure 45: Network throughput comparison between our solution and an ideal oracle which has perfect and continuous knowledge of obstacles and UEs positions, leveraging either the theoretical channel model from WP4 or the ray-tracing analysis with real measurements from WP3, under varying radar update periodicity (τ), number of radar receiving antennas (M_r), and the presence of RISs.

analysis including measurements of real RIS and antenna model (WP3/WP5-related), only considering the links under investigation:

- LoS between BS and UE0;
- LoS between BS and UE1;
- NLoS between BS and UE1, with an alternative path through RIS1 (see Figure 41),

where it can be clearly seen that the values in LOS conditions are comparable, while the theoretical model significantly overestimates the performance in NLOS when RIS1 is activated.

The final results are shown in Figure 45, which, compared to Figure 44, has an additional blue dashed horizontal line illustrating the resulting network throughput obtained with the SiMoNe ray-tracing input compared with the theoretical results. For a fair comparison, we only hard-coded the received power levels provided by the WP3 members in our simulations, while we preserved the scenario (see Figure 41), and simulation parameters (see Table 11). In terms of algorithms, we decided to apply the raytracing-based channel model on the oracle approach rather than our proposal, that is applying it on exact positions than estimated positions, because this ensures an exact performance comparison, as the objective of this new curve is to showcase the network throughput difference when considering a more realistic channel model.

As can be seen, however, this curve overlaps with the oracle-based case simulated with WP4 model and without activating the RIS, reflecting the severe path loss observed in the WP3 simulations when exploiting the UE–RIS–BS link. This means that the oracle-based approach with RISs simulated with the theoretical models, represented by the dotted-dashed line in Figure 45, tends to overestimate the network throughput by approximately 16% for the studied industrial scenario.

As shown in Table 13, the P_r for the UE1–RIS1–BS link, when the direct path is obstructed, is significantly low (–111.67 dBm). To achieve a bit error probability of, e.g., 10^{-2} with the modulation order considered in this study (see Table 11), P_r should be approximately –60.4 dBm [11]. This explains why the network throughput obtained with the WP3 model is essentially equivalent to the WP4 results without RIS.

---

[11] To achieve $P_\mathrm{r} = -60.4$ dBm with the WP3 model, the scenario would need to be scaled down by a factor of approximately 19.

5.4 Network-aided blockage prediction for reliable THz communications

The establishment of reliable links for seamless and uninterrupted connectivity comes with a series of challenges, in particular for higher frequency bands such as millimeter wave and THz [104, 105]. In particular, severe spreading loss, molecular absorption, and diffuse scattering result in severe path loss that can be compensated by highly-directive antennas and efficient beam management techniques [106]. Ensuring seamless connectivity requires proactive reasoning; therefore, being able to predict the blockage status of a link in advance is a paramount prerequisite. In this way, a heavily obstructed link can be replaced by (e.g.), dynamically configuring a RIS within the communication range.

Given the lack of rigorous mathematical models to address link blockage, many recently proposed solutions in the literature revolve around the use of Machine Learning (ML) techniques. In [107] and [108], observations of beamforming vectors adopted at the BS side are used to monitor and predict single/multi-directional link blockages within the next few time frames in a mmWave/sub-THz MIMO system. Anders et al. [109] introduce a centralized Recurrent Neural Network (RNN)-based predictor, trained using meta-learning to overcome challenges related to adaptability, few-shot learning, and generalization. While these solutions effectively address blockage prediction and offer novel approaches, they fall short in terms of scalability to large network scenarios. Specifically, the prediction task relies on a centralized controller that, unlike our proposed approach, processes *unstructured* data, limiting the ability to handle complex and expansive network environments. In [110], the idea of predicting future blockages from the observation of in-band wireless signal is presented. By exploiting RNNs and Convolutional Neural Networks (CNNs) for sequence processing (limited to the scope of individual devices), the authors show that it is possible to exploit self-interference due to multipath components to infer a link's future blockage status. Similarly, the authors from [111] consider the same task employing Liquid Time Constant Network (LTC) networks [112] as a light, physics-inspired, and adaptable solution for sequence processing tasks. While the proposed solutions are effective in identifying short-term blockages, their performance quickly degrades with increasing prediction windows. This is due to devices partial observability, which is bounded by the phenomenon of self-interference manifesting shortly before the occurrence of a blockage.

Finally, as blockage prediction lends itself to being addressed through the observation of heterogeneous measurements, different contributions in the literature consider the use of data sources beyond radio channel measurements, such as camera images [113, 114], range-angle LiDAR maps [115], vehicular sensors [116] or multimodal data [117].

In this section, we consider a scenario where multiple sensors are connected to an access points through wireless links, which can be obstructed by moving obstacles. Our goal is to predict in advance the blockage status of the links by effectively exploiting structured information collected at the sensors. Leveraging recent advancements in Machine Learning on Graphs (GML), we harness the inherent spatial and temporal (non-linear) dependencies of the problem at hand. This approach enables us to address the above-mentioned open challenges related to scalability (**C1**), partial observability (**C2**), and long-term predictions (**C3**).

5.4.1 System model

Our system considers a set $\mathcal{M}$ ($|\mathcal{M}| = M$) of wireless nodes connected to an access point. These nodes can represent sensing or actuation devices deployed in an industrial environment communicating with a central controller. In order to meet typical Quality of Service (QoS) requirements of industrial applications (e.g., motion control, virtual commissioning, predictive maintenance and control, mobile robot management [118–120]), and ensure operational safety, LoS condition between nodes and a central controller should be constantly maintained. To enable proactive identification of LoS blockage, all nodes $i \in \mathcal{M}$ sense the channel with periodicity T and store sequential observations. Through periodic measurement reporting, these observations are sent to the central controller, which is in charge of the training of the ML model for the prediction task. As an example, in the case of a 5G network this can be achieved by exploiting CSI-RS or SSB measurement reports. With the aim of exercising a device-agnostic solution, we represent a generic channel measurement using the notation $s_i(t)$. This measurement may range from narrowband power measurements (e.g., CSI-RSRP) to estimates of Channel Impulse Responses (CIRs), depending on the complexity of the devices. We further denote each sequence of measurements as

$$\mathbf{x}_{i,n} = \{s_i^{(t_0-(n-1)T)}, s_i^{(t_0-(n-2)T)}, \dots, s_i^{(t_0)}\} \ , \tag{13}$$

where n indicates the sequence length. We denote the generic antenna gain at the i-th node by $g_{r_i}(\phi, \psi)$, while $g_t(\phi, \psi)$ indicates the gain of the coordinator node/BS.

5.4.2 Problem Formulation

Let us introduce a graph-representation of the network $\mathcal{G} := (\mathcal{M}, \mathcal{E}, \mathcal{W}, \mathcal{F})$, where $\mathcal{M}$ denotes the set of nodes, $\mathcal{E}$ indicates the set of edges, $\mathcal{W}$ denotes the set of edge weights, and $\mathcal{F}$ denotes the set of node features. Each element $e_{i,j} \in \mathcal{E}$ represents an ordered connection of nodes $(i, j) \in \mathcal{M}$, and the set of peers for every node collectively defines the edge set $\mathcal{E}$. In our scenario, each node feature $\mathbf{x}_{i,n} = f_{i,n} \in \mathcal{F}$ corresponds to the temporal sequence of measurements collected by node i, as defined by Eq. (13). Let us further define the graph shift operator, denoted as $\mathbf{S}$, as the sparse adjacency matrix $\mathbf{S} := w_{i,j}$, for all $(i, j) \in \mathcal{E}$.

Given a graph $\mathcal{G}$, and a graph shift operator $\mathbf{S}$, it is finally possible to define the temporal process supported by the graph as the set $\mathbf{X}_n \in \mathbb{R}^{M \times n}$, such that

$$\mathbf{X}_n = \{\mathbf{x}_{1,n}, \mathbf{x}_{2,n}, \dots, \mathbf{x}_{M,n}\} \ . \tag{14}$$

To simplify our notation, we will refer to $\mathbf{X}_n$ as $\mathbf{X}$ from this point forward. We further indicate an instantaneous realization of the graph temporal process $\mathbf{X}$ as the graph signal $\mathbf{X}(t_k) = \{s_1^{(t_k)}, s_2^{(t_k)}, \dots, s_M^{(t_k)}\}$.

We formulate the blockage prediction problem as a sequence-to-sequence learning task. The goal is to minimize the cross-entropy loss, denoted as $\mathcal{L}$, by comparing the predicted future blockage status of the i-th node's LoS path with the actual ground truth sequence denoted as $\mathbf{y}_i^{(k)}$. Specifically, we express the latter as

$$\mathbf{y}_{i,k} = \sum_{s=1}^{n} a_{s,i} g(t - (s - k)\ T) \ , \tag{15}$$

where k represents the time step horizon, $a_{s,i}$ takes values in the set $\{0, 1\}$ and denotes the blockage status, and $g(t)$ is a rectangular pulse shape. The width of $g(t)$ is equal to the channel sampling time T for each edge node. Considering the set of temporal sequences for all nodes $i \in \mathcal{M}$ yields a structured sequence $\mathbf{Y}_k = \{\mathbf{y}_{1,k}, \dots \mathbf{y}_{M,k}\}$. For ease of notation, we will refer to $\mathbf{Y}_k$ as $\mathbf{Y}$ in the remainder of the section. The goal is to learn a function $\Phi(\mathbf{X}, \mathbf{S}, \theta)$, parametrized by θ, such that, from the observation of the structured sequence $\mathbf{X}$ the loss $\mathcal{L}(\Phi(\mathbf{X}, \mathbf{S}, \theta), \mathbf{Y})$ is minimized. Accordingly, the loss function $\mathcal{L}$ is defined as

$$\begin{aligned} &\mathcal{L}(\Phi(\mathbf{X}, \mathbf{S}, \theta), \mathbf{Y}) = \\ &\quad \frac{1}{M} \sum_{i=1}^{M} \frac{1}{n} \sum_{j=1}^{n} y_{i,j} \cdot \log\left(P(\hat{y}_{i,j} = 1 |\ \Phi)\right) + \\ &\qquad\qquad (1 - y_{i,j}) \cdot \log(P(\hat{y}_{i,j} = 0 |\ \Phi)) \ , \end{aligned} \tag{16}$$

where $\hat{y}_j$ represents the j-th element of the sequence predicted by Φ. We aim to solve the above-introduced problem through empirical risk minimization, i.e., by learning a function parametrization θ to minimize the empirical loss function over a training set $\mathcal{T} \sim P(\mathbf{X}, \mathbf{Y})$

$$\arg\min_{\theta} \sum_{(\mathbf{X}, \mathbf{y}) \in \mathcal{T}} \mathcal{L}\left(\Phi(\mathbf{X}, \mathbf{S}, \theta), \mathbf{Y}\right) \ . \tag{17}$$

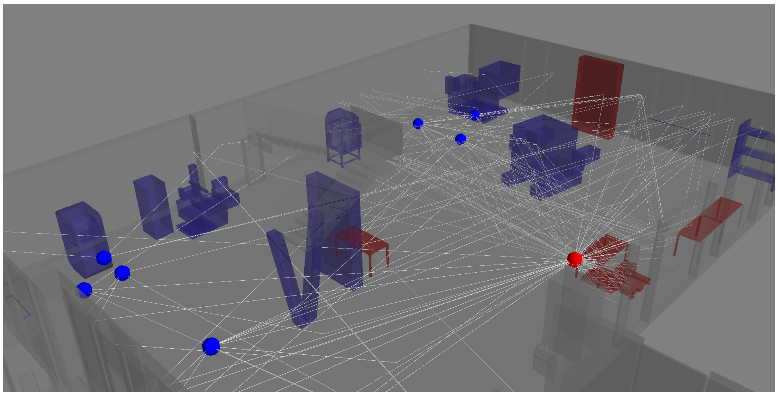

(a) Realistic industrial scenario (20×10 m^2)

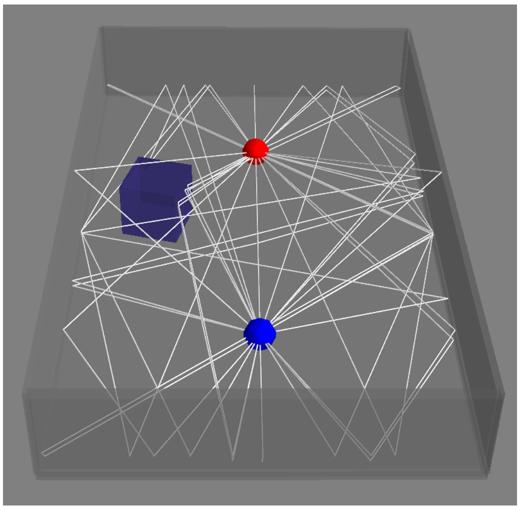

(b) Empty room scenario (5×10 m^2)

Figure 46: Ray tracing scenarios. Note: fewer receivers are depicted in the figures for visual purposes.

5.4.3 Methodology

Ray tracing simulation & scenarios We consider a set of distributed nodes according to the system model in Section 5.4.1. With the aim of advancing reproducible research, we use ray tracing simulations with the open-source tool Sionna [23] to evaluate distributed channel measurements in various environments and blocking conditions. In our simulations, we consider one transmitter and $r = 25$ receivers communicating at 300 GHz within two distinct scenarios (Figure 46), namely:

- *Empty room*: in an environment consisting of an empty room with concrete walls and receiver positions arranged in a fixed grid (Section 5.4.4), or randomly sampled from a uniform Poisson Point Process (PPP) (Section 5.4.4), we introduce a metallic cube acting as a blocker. This cube moves linearly from one point of the room to another, with its start and stop positions randomly determined.
- *Realistic industrial scenario*: we model realistic trajectories for a mobile robot navigating the industrial environment of the BI-REX smart factory in Bologna, Italy [18]. From a collection of feasible regions ensuring direct LoS with a receiver mounted on a wall within the area, we randomly select a set of receivers for every experiment.

From these environments, we generate two distinct datasets $\mathcal{T}_{\{A,B\}}$ comprising $I = 150$ independent experiments with random blocker mobility patterns. During each experiment, a set of r nodes collectively senses the wireless environment for a sequential number of $n = 80$ steps. The outcome of the ray tracing simulation can be formally described as a Directional Channel Impulse Response (DCIR) $h_i(t, \tau, \phi, \psi)$ in the time (t)-delay (τ)-angular (ϕ, ψ) domain [121], evaluated at each node $i \in \mathcal{M}$. Accordingly, the resulting CIR in the time-delay domain, assuming far-field conditions, is given by

$$h_i(t, \tau) = \int_{\phi}\int_{\psi} h_i(t, \tau, \phi, \psi) \cdot g_t(\phi, \psi) \cdot g_{r_i}(\phi, \psi)\, d\phi\, d\psi \,, \tag{18}$$

The contribution of the directional gain in Eq. (18) yields a straightforward, yet interesting tradeoff. When communicating at THz frequency, directional gain and efficient beam alignment are beneficial to overcome severe path loss and improve link spectrum efficiency. Yet, when the same waveform is used to derive environment information for blockage prediction (or any other kind of sensing task), a highly directional channel yields a less informative mean of information, as secondary Multipath Components (MPCs) are strongly attenuated. To this end, considering the graph signal of channel measurements collectively sampled by the network of industrial nodes (collected at a centralized controller) provides an effective means to avert the above-mentioned tradeoff without requiring additional radio resources for sensing (e.g., to perform periodic beam sweeping). Nonetheless, in our simulations, we opted to leave the CIR unfiltered in space to evaluate the performance gain introduced by the Spatio-Temporal Graph Neural Network (STGNN) under worst-case conditions. We anticipate that incorporating directional filtering of the CIR would widen the performance gap even further compared to other baselines, as the wireless signature phenomenon (i.e., increased self-interference due to MPCs) becomes less noticeable. However, this aspect remains a subject for future investigation.

As previously remarked, one of the objectives of our framework is to exercise a device complexity-agnostic solution. In this regard, it would be unrealistic to assume low complexity devices able to estimate CIRs and use them directly for blockage prediction. Instead, we consider node measurements $s_i(t)$ to coincide with the channel gain, which is directly proportional to the received power under the assumption of constant transmit power, as per

$$s_i(t) = \left| \int_{-\infty}^{\infty} h_i(t, \tau)\, d\tau \right|^2 \,. \tag{19}$$

Spatio-Temporal Blockage Prediction Graph-structured temporal sequences are inherently characterized by spatial and temporal dependencies. Learning from sequences of graph signals involves the flow and processing of information in the spatial/temporal domain as a means to capture and model such dependencies. Recently, a variety of deep graph-recurrent [122–124] and graph-convolutional [125] models able to process structured sequential data in both domains have been shown to achieve state-of-the-art performance in a variety of learning tasks. Here, inspired by recent advancements in the field, we consider a simple yet effective solution to tackle link blockage prediction by leveraging the collective observations gathered from a network of nodes. In our model, we sequentially process spatial information using graph-convolutional heads and temporal information through stacked recurrent layers. In this way, graph-structured signals $\mathbf{X}(t_k) \in \mathbf{X}$, with $k \in \{1, \ldots, n\}$, are convolved over the graph space before being reconstructed and processed as sequences. The considered graph operator is formally expressed by

$$\mathbf{h}_i^{(l+1,t_k)} = \sigma\left(\theta_1^{(l)}\mathbf{h}_i^{(l,t_k)} + \theta_2^{(l)} \sum_{j\in\mathcal{N}(i)} w_{i,j} \cdot \mathbf{h}_j^{(l,t_k)}\right) . \tag{20}$$

In this equation, $\mathbf{h}_i^{(l,t_k)}$ refers to the node embedding of the i-th node, at time t_k and layer l. $\mathbf{h}_*^{(1,t_k)}$ is equal to the instantaneous node observation $s_*^{(t_k)}$. The function $\sigma(\cdot)$ represents a non-linear activation function, such as ReLu or sigmoid. The model parameters at layer l are represented by $\theta_*^{(l)}$, which are calculated differently for the node embedding and its neighborhood. Finally, $w_{i,j}$ refers to an arbitrary edge weight between node i and j, and $\mathcal{N}(i)$ is the node's neighborhood. The node embeddings are computed sequentially for all time instants using the same model parameters θ_* and updated simultaneously during training. A sequence of temporal node embeddings, denoted as $\mathbf{h}_i^{(l)}$, is then reconstructed by concatenating all elements $\mathbf{h}_i^{(l,t_k)}$ by row

$$\mathbf{h}_i^{(l)} = [\mathbf{h}_i^{(l,t_1)}, \ldots, \mathbf{h}_i^{(l,t_k)}] \,. \tag{21}$$

For clarity of notation, in the following, we denote each temporal element of (21) as $\mathbf{h}^{(t)}$, implying that the embedding of each node is computed at the last graph layer. The sequence (21) is subsequently processed by stacked Long-Short Term Memory (LSTM) layers [126], according to the formulation proposed in [127], where each cell updates its latent space $\tilde{\mathbf{h}}^{(t)}$ as per

$$\begin{aligned}
in^{(t)} &= \sigma\left(\theta_{\mathbf{h},in}\mathbf{h}^{(t)} + \theta_{\tilde{\mathbf{h}},in}\tilde{\mathbf{h}}^{(t-1)} + \theta_{\mathbf{c},in} \odot \mathbf{c}^{(t-1)} + b_{in}\right), \\
f^{(t)} &= \sigma\left(\theta_{\mathbf{h},f}\mathbf{h}_t + \theta_{\tilde{\mathbf{h}},f}\tilde{\mathbf{h}}^{(t-1)} + \theta_{\mathbf{c},f} \odot \mathbf{c}^{(t-1)} + b_f\right), \\
\mathbf{c}^{(t)} &= f^{(t)} \odot \mathbf{c}^{(t-1)} + in^{(t)} \odot \tanh\left(\theta_{\mathbf{h},\mathbf{c}}\mathbf{h}^{(t)} + \theta_{\tilde{\mathbf{h}},\mathbf{c}}\tilde{\mathbf{h}}^{(t-1)} + b_{\mathbf{c}}\right), \\
o^{(t)} &= \sigma\left(\theta_{\mathbf{h},o}\mathbf{h}^{(t)} + \theta_{\tilde{\mathbf{h}},o}\tilde{\mathbf{h}}^{(t-1)} + \theta_{\mathbf{c},o} \odot c_t + b_o\right), \\
\tilde{\mathbf{h}}^{(t)} &= o \odot \tanh\left(c_t\right),
\end{aligned} \tag{22}$$

where $in^{(t)}, f^{(t)}, \mathbf{c}^{(t)}, o^{(t)}$ represent the input, forget, cell, and output gates respectively. $\tilde{\mathbf{h}}^{(t)}$ denotes the latent hidden space, and $\mathbf{h}^{(t)}$ represents the embedding space of the upstream graph convolutional layers which is supplied as model input at every time step. Finally, the model parameters for every gated unit are denoted by $\theta_{(\cdot,\cdot)}$, and $\odot$ denotes the Hadamard product.

One major advantage of the proposed formulation is its ability to facilitate scalable learning. Unlike centralized models working over unstructured data, for which augmenting the number of devices in the network incurs in the curse of dimensionality, STGNN offers a scalable solution to network size by operating on local neighborhoods of nodes, which allows for efficient processing of large graphs without requiring global information. Additionally, the model parameters θ are applied for message transformation to all nodes across the graph $\mathcal{G}$ through the operator **S**, making the number of parameters irrelevant to network size and enabling inference over graphs of varying sizes. This means that the proposed solution can be trained on a small network and seamlessly used for inference on a larger setting.

5.4.4 Experimental Results

We evaluate the predictive performance of our solution against a set of baselines. Our objective is to jointly address the challenges **C1**, **C2**, and **C3**. To that end, we conduct an ablation study, which, at the same time, involves testing our solution against representative family of models that have been recently employed in the literature, including centralized [107–109] and decentralized [110, 111] RNNs for sequence processing. Specifically, we consider:

- a stacked LSTM model with partial observability (**LSTM partial**), i.e., an LSTM model trained on individual time series collected by all nodes. This model does not consider spatial relations. However, it is naturally suited for scalable distributed inference.
- a stacked LSTM model with global observability (**LSTM global**), i.e., an LSTM model trained on multivariate time series, which accounts for all node observations concurrently. This approach enables the modeling of spatial relations by simultaneously observing all nodes.
- a stacked LSTM model with global observability and positional encoding (**LSTM global pos. enc.**). The positional encoding of each node is fed as residual information to a fully connected layer downstream of the stacked recurrent layers. Additionally, node ordering is maintained by arranging node sequences and (x, y) positions based on their relative distance to a fixed reference point in the room.
- The **STGNN** model described in section 5.4.3. This model incorporates a relational inductive bias in the learning process by handling structured data. Unlike the global LSTM models, as described in Section 5.4.3, this solution offers inherent increased scalability to network size.

We categorize our experimental results within transductive and inductive setting. The former relates to model inference on a held-out test set comprising unlabelled nodes of the same graph $\mathcal{G}$. The latter entails inference

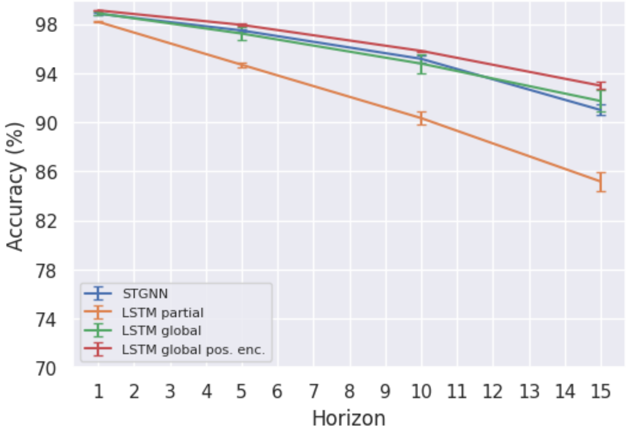

(a) Transductive scenario (empty room)

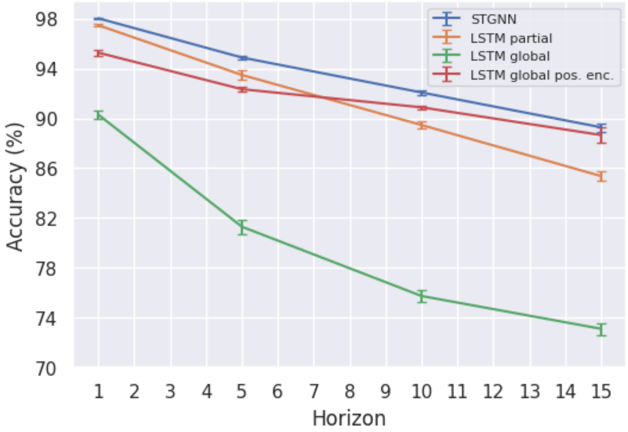

(b) Inductive scenario (empty room)

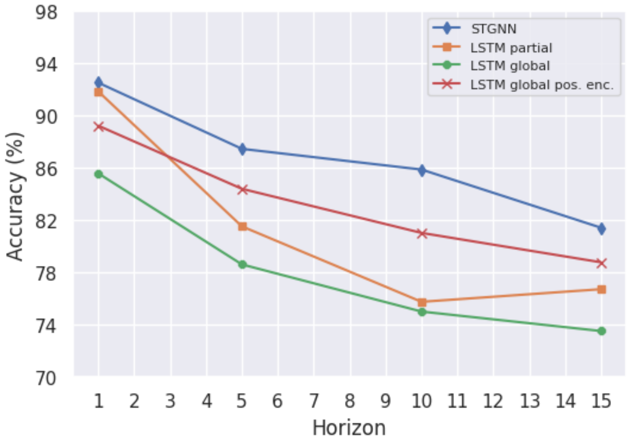

(c) Inductive scenario (empty → industrial)

Figure 47: Numerical experiments - Accuracy.

on unseen graphs during training time, thus denoting the model ability to infer and generalize beyond local observations. For illustrative numerical experiments, we consider the LSTM models composed by two stacked recurrent layers with 64 neurons each and one final fully connected layer. The STGNN model, on the other

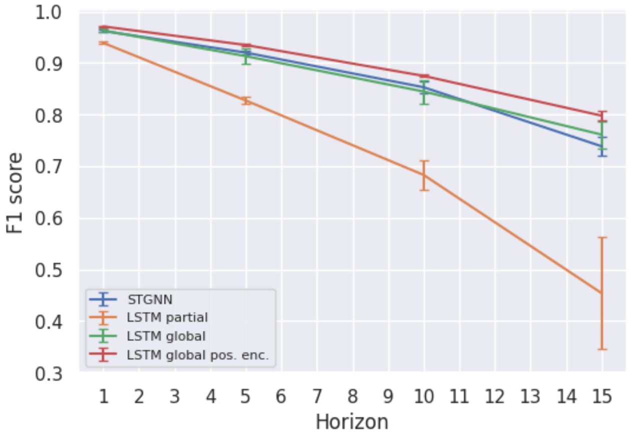

(a) Transductive scenario (empty room)

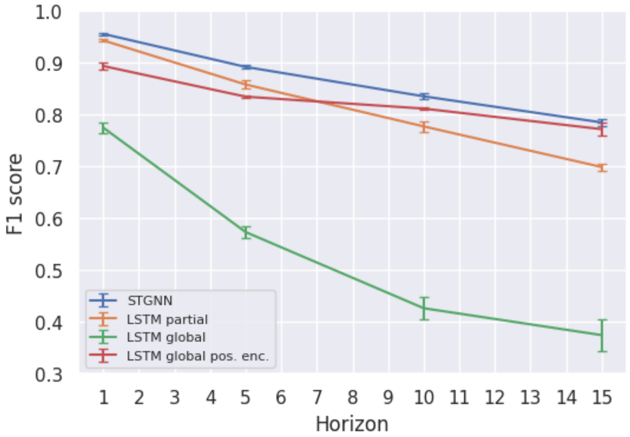

(b) Inductive scenario (empty room)

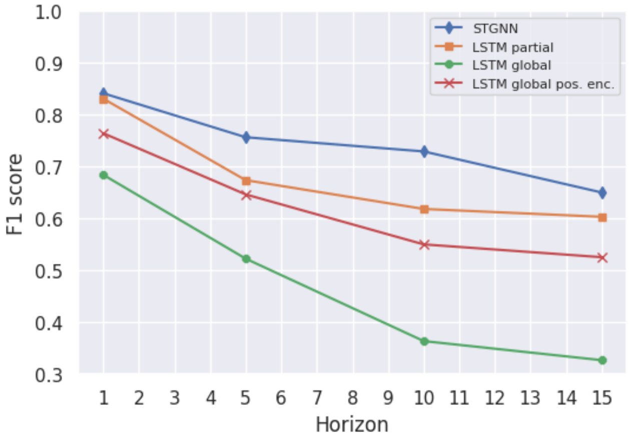

(c) Inductive scenario (empty → industrial)

Figure 48: Numerical experiments - F1 score.

hand, is composed of three graph layers with embedding dimensions corresponding to $\{25, 25, 10\}$ and the same downstream recurrent layers of the LSTM models. For the STGNN, we further consider a set $\mathcal{E}$ of binary edge weights where $e_{i,j} = 1 \iff \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \leq r$, and $r = 2.5$ [m] in the case of *empty room* scenario, whereas $r = 6$ [m] in the case of *realistic industrial scenario*.

Transductive setting We evaluate the nodes' blockage prediction performance on a held-out test set comprising unlabelled nodes on a static graph. We perform blockage simulations within the *empty room* scenario considering a total of 120 independent blockage trajectories for training, and 30 for testing. Each experiment yields 80 sequential channel observations by a set of 25 nodes. The nodes are arranged over a fixed grid and evenly spaced within the 10x5 m^2 room. The results of 32 independent training iterations with random parameter initialization have been summarized in Figure 47(a) and 48(a) in terms of overall accuracy and F1 score, computed as:

$$F1 = \frac{2 \cdot TP}{2 \cdot TP + FP + FN}, \tag{23}$$

where TP, FP, and FN represent the number of true positives, false positives, and false negatives, respectively.

In the transductive scenario, the LSTM partial heavily underperforms with respect to the centralized solutions and the STGNN. This is imputable to the fact that the model is not able to compensate for partial visibility and thus can only rely on local observations to predict blockages. This becomes especially pronounced for long-range predictions, as the influence of the blocker on the LoS channel becomes visible to the node only a few time steps before the actual blockage occurs. On the other hand, the STGNN slightly outperforms the LSTM global model and approaches the performance of the LSTM global with positional encoding. This is a noteworthy accomplishment as both baseline models benefit from explicit information in a static scenario, which makes them non-scalable approaches and unable to generalize as effectively as the STGNN in a dynamic setting, as demonstrated in the following subsection.

Inductive setting We evaluate the ability to predict blockages for scenarios beyond the model's local observations during training time. Specifically, we consider the task of blockage prediction in a dynamic setting, where inference is performed over unseen graph structures. This involves incorporating node positions from independent realizations of PPPs and blocker trajectories within the same room. Additionally, we evaluate the model performance when tasked to perform inference on the realistic scenario, while being trained on the empty room scenario.

Figure 47(b) and 48(b) depict numerical results obtained in the inductive setting, where training and inference are performed in the same scenario, i.e., *empty room*. Several notable observations can be made. Firstly, STGNN consistently outperforms all baseline models, aligning with our expectations owing to its enhanced ability to generalize across inductive scenarios compared to models relying solely on explicit information. Secondly, LSTM global exhibits notably inferior performance compared to other models. This is attributed to its lack of permutation equivariance, which, conversely, is a prominent feature of GNNs. As a consequence, its performance is affected by random sampling from the PPP, leading to inconsistent ordering in the model input space, which is perceived as independent observations. While the introduction of positional encoding and node ordering partially mitigates these issues, the resultant LSTM global pos. enc. model still falls short of STGNN's performance. On the other hand, the LSTM partial model remains unaffected by the inductive scenario in terms of performance compared to its performance in the transductive case. This is because the latter is trained on all sequences collected by the nodes but presented with one sequence at a time. Notably, the performance gap between STGNN and LSTM global pos. enc. slightly decreases by considering longer time horizons. This is likely due to the sequential averaging nature of GNNs, for which relevant information from further neighbors in the graph space gets averaged through sequential latent space updates. We anticipate that this aspect could be improved by learning goal-specific (horizon-specific) attention weights, which is left for future work.

In Figure 47(c) and 48(c), the outcomes of the last inductive experiment are presented. This experiment involves performing inference on the *realistic industrial scenario* while using models trained on the *empty room*. Testing on the new scenario involves an unseen environment, blocker shape, trajectories, as well as unseen graph structure and labels. The numerical findings reveal that STGNN surpasses the other models in terms of accuracy and F1 score, showcasing its capacity to handle unseen observations.

6 Conclusions and main findings

This deliverable 4.4 constitutes the final submission of WP4 within the TIMES project, with a particular emphasis on Tasks T4.3 and T4.4. The purpose of this document is to provide a detailed account of the main achievements of these two tasks. Specifically, Task T4.3 focuses on the exploration of smart radio environments to enable physics-aware algorithms and procedures for THz networks, while Task T4.4 addresses the design of MAC and NET layer protocols tailored for THz communications in intelligent, multi-goal mesh networks.

To this end, Section 3 presents the reference models adopted for these contributions, including RF impairments, channel characterization, and RIS modeling, drawing on existing literature that underpins the research developments described and inputs from other WPs (namely WP2, WP3 and WP5). Sections 4 and 5 then highlight the various research outcomes for T4.3 and T4.4, respectively, whose key findings are summarized as follows:

- Sect. 4.1 investigates the problem of cascaded channel estimation in RIS-aided THz wireless communications under the presence of random EMI. Unlike prior studies [34], a scenario in which the UE, BS, and interference source are all located within the NF region of the RIS, is considered. By leveraging known NF statistical characteristics, an AO-based LMMSE estimator is employed to accurately recover the uplink cascaded channel. Numerical results confirm the necessity of incorporating NF statistics for channel estimation in such a scenario, and demonstrate the clear advantage of the AO-based LMMSE estimator over existing alternatives in NF cascaded channel estimation.
- Sect. 4.2 examines the potential of RIS-enhanced deployments in factories at THz frequencies by exploiting properties of the radiative near-field under LOS conditions. In this case, the LOS MIMO channel then yields high-rank, enabling parallel communication via its eigenchannels, which can boost the link capacity significantly. These gains come at a high sensitivity of the additionally used eigenchannels, though, which has been analyzed in detail for a specific reference scenario.
- Sect. 4.3 presents a power adaptation approach to improve the communication performance of a RIS assisted link while being impaired by transmit PA nonlinearity and limitations in RIS phase shifting capabilities. The results show that allowing for the variation in transmission power allows for the improvement in the communication performance. Moreover, the results have shown that the improvement due to power adaptation is significant when phase shifting capabilities at the RIS are good, while limited performance gains can be achieved if the RIS has low resolution at its elements' phase shifters.
- Sect. 4.4 presents an in-depth analysis of the performance limits in the near-field of a SIMO monostatic ELAAs for target position (in terms of range and AoA) and velocity estimation. More specifically, the asymptotic CRLB for range and angle estimation is derived in the geometric near-field region, showing a performance saturation mainly due to the adoption of the true phase profile for range and due to the nonuniform signal strength across the array for AoA. Concerning instead velocity estimation in near-field region, the adoption of ELAAs allows the projection of the target velocity along the set of directions that connects the target and each element of the array. This permits the estimation not only of the radial component of the velocity, as in far field, but also the transverse component, which ensures a more detailed picture of the target motion.
- Sect. 4.5 addresses the joint design of transmit beamforming at the AMAF and phase shift configuration of the RIS in a multi-user MISO downlink system operating in the near field. An alternating optimization framework based on the WMMSE criterion is employed to jointly refine active and passive beamforming strategies, aiming to maximize the system sum-rate while mitigating inter-user interference. Numerical results show that the proposed heuristic RIS configuration, based on a focusing strategy, achieves per-UE and sum-rate performance nearly matching that of the MMSE-based scheme, requiring substantially lower computational burden and no CSI. Its performance also closely approaches that of the fully optimized solution, while avoiding channel estimation and significantly reducing complexity. This confirms its suitability as a low-complexity, scalable, and near-optimal strategy for near-field multi-user MISO systems in future 6G scenarios.

- Sect. 4.6 has presented raytracing-based simulations to demonstrate that reliable communication via a RIS at 285 GHz is feasible in an industrial environment over distances exceeding 50 m, provided sufficiently high RCS values are achieved. The evaluation shows that a constant RCS of at least -18 dBsm enables basic data transmission, while values above 0 dBsm allow operation with 64-QAM and high-rate FEC, reaching the maximum data rate at a bandwidth of 2.16 GHz defined in [40]. These results will be taken into account for defining the layout of the POC that will be carried out as part of WP6 in the Robopac factory hall.
- Sect. 5.1 proposes two fully distributed multi-hop routing protocols for THz networks that rely solely on user-plane data transmissions: TL and TB, and an extension of TB based on MADRL. Under static conditions, the TB protocol outperforms TL due to efficient next-hop selection and lower channel overhead, especially with high TTL values that limit neighbor table updates. However, when the overhead increases (e.g., due to a higher number of UEs or larger payloads), TB suffers from frequent neighbor losses, which can result in suboptimal next-hop selection. In contrast, TL offers a higher likelihood of reaching the BS, though often through non-optimal paths, which leads to higher average latency compared to TB. In dynamic scenarios, the MADRL-based approach outperforms classic TB throughput due to its ability to optimally choose between unicast and broadcast transmissions, enabling better adaptation to changing network conditions. However, this results in higher average latency, as distant UEs more frequently succeed in reaching the BS, increasing the number of hops in successful transmissions. Overall, shorter TTL values further enhance performance by allowing more frequent updates of the neighbor table in dynamic environments. Through this analysis, which was obtained through simulations, using a scenario, channel model, and parameters defined based on inputs from WP2, WP3, and WP5, respectively, we identify a set of trade-offs between throughput and latency which affect the choice of the best performing protocol and are influenced by network conditions such as channel overhead, UEs' density, and mobility conditions.
- Sect. 5.2 has proposed a centralized RL-based model, developed at the BS, to dynamically manage network access in a highly congested, multi-goal industrial environment. In this context, sensors generate data with different priorities and corresponding requirements that vary over time. The proposed model predicts requirements variability and enables optimal assignment of the CSMA CW values to individual sensors. The proposal has been evaluated against benchmark approaches based on traditional CSMA, with fixed CW configurations. Results demonstrate that the presented approach offers superior adaptability to dynamic and heterogeneous traffic priorities by autonomously adjusting CW assignments in response to user-specific needs. This enables the joint satisfaction of stringent latency and success probability requirements for different network conditions and/or temporal evolution. In contrast, static configurations either fail to meet these dynamic demands or require multiple trials to identify a suitable setup.
- Sect. 5.3 presented an algorithm that dynamically and proactively configures the THz network (i.e., BS and RISs) based on centimeter-level sensing data obtained from a multi-antenna radar co-located with the BS. The proposed solution is validated through network-level simulations of uplink data traffic in a realistic, dynamic industrial environment and compared against an ideal oracle with perfect, continuous environmental knowledge. The main findings of this study highlight the importance of both frequent radar updates to maintain seamless knowledge of the scenario over time and high radar resolution to achieve a proper radar target estimation at THz frequencies. This is crucial to handle the potential issue of beam misalignment due to the extremely narrow beams at these frequencies. Notably, numerical results show that our approach can closely match the oracle's network throughput when the radar has at least 128 receive antennas and generates radio maps with a periodicity of at most 100 μs. The analysis using the WP3 channel model highlights that the two PHY-layer results are in close agreement under LoS conditions. Therefore, for the studied industrial scenario, the free-space model used in WP4 can be considered a reasonable approximation for THz communications. On the other hand, the approximated RIS model in WP4 tends to overestimate the NLoS link when compared with the realistic RIS implementation, even considering the high-gain antenna of WP5 at the BS.
- Sect. 5.4 presented an AI-based solution to predict in advance the blockage status of the links by effectively exploiting structured information collected at the nodes. The experimental results validate the

effectiveness of the proposed STGNN model for LOS blockage prediction, as exemplified by predicting blockage for THz links in industrial environment. The STGNN architecture consistently outperforms baseline approaches across various settings, demonstrating its ability to leverage both spatial and temporal dependencies within the network. Importantly, the framework excels in inductive learning tasks, where the model generalizes well beyond the training data to predict blockages in unseen environments with different graph structures and blocker dynamics. The proposed solution has the potential to be deployed for low-latency decentralized inference in a mesh network with peer-to-peer communication, in accordance with the principle of Centralized Training - Decentralized Execution (CTDE).

Overall, the presented set of scientific results establishes important milestones (i) in the theoretical characterization and practical exploitation of smart propagation environments for THz networks, and (ii) in the development of both model-based and data-driven algorithms and procedures for intelligent mesh networks operating at THz frequencies. The significance of these contributions is evidenced by their direct or indirect dissemination through publications in leading first- and second-tier journals and conferences (see, e.g., [21, 22, 24, 34, 38, 39, 43, 64, 67, 90, 97, 98, 120, 128] or check the TIMES website). Moreover, some of the concepts introduced in this deliverable will be integrated into the experimental activities of the TIMES project (i.e., WP6-related) for further validation.

References

[1] TIMES. Deliverable 4.2 - intermediate report on multi-goal mesh network optimization and exploitation of smart propagation environments. Technical report.

[2] Mihai-Alin Badiu and Justin P Coon. Communication through a large reflecting surface with phase errors. *IEEE Wireless Communications Letters*, 9(2):184–188, 2019.

[3] TIMES. Deliverable 4.3 - intermediate report on phy layer enhancements for thz links supporting sensing and communication functionalities. Technical report.

[4] TIMES. Deliverable 4.1 - intermediate report on phy layer enhancements for thz links supporting sensing and communication functionalities. Technical report, TIMES SNS, 2023.

[5] Kanti V Mardia and Peter E Jupp. *Directional statistics*. John Wiley & Sons, 2009.

[6] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1948.

[7] Christoph Rapp. Effects of HPA-nonlinearity on a 4-DPSK/OFDM-signal for a digital sound broadcasting signal. *ESA Special Publication*, 332:179–184, 1991.

[8] Nokia. Realistic power amplifier model for the New Radio evaluation. *3rd Generation Partnership Project (3GPP)*, R4-163314(TSG-RAN WG4 Meeting Number 79), 2016.

[9] 3GPP. Realistic power amplifier model for the New Radio evaluation. TSG RAN WG4 Meeting 79, R4-163314, May 2016.

[10] Sumit Pratap Singh, Timo Rahkonen, Marko E Leinonen, and Aarno Pärssinen. A 290 GHz Low Noise Amplifier Operating above $f\_\{max\}/2$ in 130 nm SiGe Technology for Sub-THz/THz Receivers. In *2021 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, pages 223–226. IEEE, 2021.

[11] Ahmad Nimr et al. Radio models and enabling techniques towards ultra-high data rate links and capacity in 6G. Deliverable D2.3, March 2023.

[12] ETSI ISG THz. TeraHertz technology (THz); RF Hardware Modeling. Group Report 004, January 2025. Available on: https://www.etsi.org/deliver/etsi_gr/THz/001_099/004/01.01.01_60/gr_thz004v010101p.pdf.

[13] TIMES. Deliverable 3.1 - state-of-art channel models in industrial environments at sub-thz frequencies. Technical report, TIMES SNS, 2023.

[14] Mingyao Cui, Zidong Wu, Yu Lu, Xiuhong Wei, and Linglong Dai. Near-field mimo communications for 6G: Fundamentals, challenges, potentials, and future directions. *IEEE Commun. Mag.*, 61(1):40–46, Jan. 2023.

[15] Naveed A. Abbasi, Arjun Hariharan, Arun Moni Nair, and Andreas F. Molisch. Channel measurements and path loss modeling for indoor THz communication. In *2020 14th European Conference on Antennas and Propagation (EuCAP)*, pages 1–5, 2020.

[16] Yuanwei Liu, Zhaolin Wang, Jiaqi Xu, Chongjun Ouyang, Xidong Mu, and Robert Schober. Near-field communications: A tutorial review. *IEEE Open J. Commun. Soc.*, 4:1999–2049, Aug. 2023.

[17] Özlem Tuğfe Demir, Alva Kosasih, and Emil Björnson. Spatial correlation modeling and rs-ls estimation of near-field channels with uniform planar arrays. In *2024 IEEE 25th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 236–240, 2024.

[18] TIMES SNS. *Deliverable D2.2 - Definition of scenarios for software simulation*, 2023.

[19] Lorenz H. W. Loeser and Thomas Kürner. Towards a Heuristic Path Loss Model for RIS Links. In *2025 URSI International Symposium on Electromagnetic Theory (EMTS)*, pages 1–4, June 2025.

[20] Nils Dreyer and Thomas Kürner. An Analytical Raytracer for Efficient D2D Path Loss Predictions. In *2019 13th European Conference on Antennas and Propagation (EuCAP)*, pages 1–5, March 2019.

[21] Christoph Herold and Thomas Kürner. A Concept for the Efficient Integration of Reconfigurable Intelligent Surfaces into a Ray Tracing Framework. In *2023 48th International Conference on Infrared, Millimeter, and Terahertz Waves (IRMMW-THz)*, pages 1–2, September 2023.

[22] Georg Jensen, Bo Kum Jung, Qi Luo, and Thomas Kürner. Enhancing Wireless Communications at 300 GHz using RIS in an Industrial Scenario. In *2025 19th European Conference on Antennas and Propagation (EuCAP)*, pages 1–5, March 2025.

[23] Jakob Hoydis, Fayçal Aït Aoudia, Sebastian Cammerer, Merlin Nimier-David, Nikolaus Binder, Guillermo Marcus, and Alexander Keller. Sionna RT: Differentiable ray tracing for radio propagation modeling. *arXiv preprint arXiv:2303.11103*, 2023.

[24] Steffen Pahlke, Tommaso Zugno, Mate Boban, Diego Dupleich, and Thomas Kürner. Ray Tracing and Measurement-Based Characterization of Inter/Intra-Machine THz Wireless Channels. In *2024 18th European Conference on Antennas and Propagation (EuCAP)*, pages 1–5, 2024.

[25] Diego Dupleich, Damir Sitdikov, Alexander Ebert, and Mate Boban. Measurement-based Validation of Ray-tracing Model at sub-THz for ISAC Applications of Blockage in Industrial Scenario. In *2024 4th URSI Atlantic Radio Science Meeting (AT-RASC)*, 2024.

[26] Davide Dardari and Devis Massari. Using metaprisms for performance improvement in wireless communications. *IEEE Transactions on Wireless Communications*, 20(5):3295–3307, 2021.

[27] Marzieh Najafi, Vahid Jamali, Robert Schober, and H. Vincent Poor. Physics-based modeling and scalable optimization of large intelligent reflecting surfaces. *IEEE Transactions on Communications*, 69(4):2673–2691, 2021.

[28] Yuanwei Liu, Zhaolin Wang, Jiaqi Xu, Chongjun Ouyang, Xidong Mu, and Robert Schober. Near-field communications: A tutorial review. *IEEE Open J. Commun. Soc.*, 4:1999–2049, Aug. 2023.

[29] David Middleton. Statistical-physical models of electromagnetic interference. *IEEE Trans. Electromagn. Compat.*, EMC-19(3):106–127, Aug. 1977.

[30] Huayan Guo and Vincent K. N. Lau. Robust deep learning for uplink channel estimation in cellular network under inter-cell interference. *IEEE J. Sel. Areas Commun.*, 41(6):1873–1887, Jun. 2023.

[31] Songjie Yang, Wanting Lyu, Zhenzhen Hu, Zhongpei Zhang, and Chau Yuen. Channel estimation for near-field XL-RIS-aided mmWave hybrid beamforming architectures. *IEEE Trans. Veh. Technol.*, 72(8):11029–11034, Aug. 2023.

[32] Yijin Pan, Cunhua Pan, Shi Jin, and Jiangzhou Wang. RIS-aided near-field localization and channel estimation for the terahertz system. *IEEE J. Sel. Topics Signal Process.*, 17(4):878–892, Jul. 2023.

[33] Jiao Wu, Seungnyun Kim, and Byonghyo Shim. Parametric sparse channel estimation for RIS-assisted terahertz systems. *IEEE Trans. Commun.*, 71(9):5503–5518, Sept. 2023.

[34] Wen-Xuan Long, Marco Moretti, Andrea Abrardo, Luca Sanguinetti, and Rui Chen. MMSE design of RIS-aided communications with spatially-correlated channels and electromagnetic interference. *IEEE Trans. Wireless Commun.*, 23(11):16992–17006, Nov. 2024.

[35] Ming-Min Zhao, Qingqing Wu, Min-Jian Zhao, and Rui Zhang. Intelligent reflecting surface enhanced wireless networks: Two-timescale beamforming optimization. *IEEE Trans. Wireless Commun.*, 20(1):2–17, Jan. 2021.

[36] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena scientific optimization and computation series. Athena Scientific, 2016.

[37] Özlem Tuğfe Demir, Alva Kosasih, and Emil Björnson. Spatial correlation modeling and RS-LS estimation of near-field channels with uniform planar arrays. In *2024 IEEE 25th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 236–240, 2024.

[38] Wen-Xuan Long, Marco Moretti, Luca Sanguinetti, and Rui Chen. Channel estimation in RIS-aided communications with interference. *IEEE Wireless Commun. Lett.*, 12(10):1751–1755, Oct. 2023.

[39] Giulio Bartoli, Andrea Abrardo, Nicolo Decarli, Davide Dardari, and Marco Di Renzo. Spatial multiplexing in near field MIMO channels with reconfigurable intelligent surfaces. *IET Signal Processing*, 17(3):e12195, Mar. 2023.

[40] IEEE Standard for Wireless Multimedia Networks. *IEEE Std 802.15.3-2023 (Revision of IEEE Std 802.15.3-2016)*, pages 1–684, February 2024.

[41] Özgecan Özdogan, Emil Björnson, and Erik G Larsson. Intelligent reflecting surfaces: Physics, propagation, and pathloss modeling. *IEEE Wireless Communications Letters*, 9(5):581–585, 2019.

[42] Fabio Saggese, Victor Croisfelt, Radosław Kotaba, Kyriakos Stylianopoulos, George C Alexandropoulos, and Petar Popovski. On the impact of control signaling in ris-empowered wireless communications. *IEEE Open Journal of the Communications Society*, 2024.

[43] Hui Chen, Musa Furkan Keskin, Adham Sakhnini, Nicolò Decarli, Sofie Pollin, Davide Dardari, and Henk Wymeersch. 6G localization and sensing in the near field: Features, opportunities, and challenges. *IEEE Wireless Communications*, 31(4):260–267, Aug. 2024.

[44] Musa Furkan Keskin, Visa Koivunen, and Henk Wymeersch. Limited feedforward waveform design for OFDM dual-functional radar-communications. *IEEE Trans. Signal Process.*, 69:2955–2970, Apr. 2021.

[45] Steven M. Kay. *Fundamentals of Statistical Processing: Estimation Theory*, volume 1. Prentice-Hall Signal Processing Series, 1993.

[46] Anna Guerra, Francesco Guidi, Davide Dardari, and Petar M. Djuric. Near-field tracking with large antenna arrays: Fundamental limits and practical algorithms. *IEEE Trans. Signal Process.*, 69:5723–5738, August 2021.

[47] Benjamin Friedlander. Localization of signals in the near-field of an antenna array. *IEEE Trans. Signal Process.*, 67(15):3885–3893, June 2019.

[48] Mohammed Nabil El Korso, Rémy Boyer, Alexandre Renaux, and Sylvie Marcos. Conditional and unconditional Cramér–Rao bounds for near-field source localization. *IEEE Trans. Signal Process.*, 58(5):2901–2907, May 2010.

[49] Zhaolin Wang, Xidong Mu, and Yuanwei Liu. Near-field velocity sensing and predictive beamforming. *IEEE Trans. Veh. Technol.*, pages 1–6, 2024.

[50] Saeid K. Dehkordi, Lorenzo Pucci, Peter Jung, Andrea Giorgetti, Enrico Paolini, and Giuseppe Caire. Multistatic parameter estimation in the near/far field for integrated sensing and communication. *IEEE Trans. Wireless Commun.*, 23(12):17929–17944, Dec. 2024.

[51] Nicolò Decarli and Davide Dardari. Communication modes with large intelligent surfaces in the near field. *IEEE Access*, 9:165648–165666, Dec. 2021.

[52] Mingyao Cui, Zidong Wu, Yu Lu, Xiuhong Wei, and Linglong Dai. Near-field MIMO communications for 6G: Fundamentals, challenges, potentials, and future directions. *IEEE Commun. Mag.*, 61(1):40–46, January 2023.

[53] Ahmed H Abdelrahman, Fan Yang, Atef Z Elsherbeni, Payam Nayeri, and Constantine A Balanis. *Analysis and design of transmitarray antennas*. Springer, 2017.

[54] Krishan K Tiwari and Giuseppe Caire. RIS-based steerable beamforming antenna with near-field eigenmode feeder. In *Proc. of IEEE Int. Conf. Commun.*, pages 1293–1299. IEEE, Jun. 2023.

[55] Krishan K Tiwari and Giuseppe Caire. On the behavior of the near-field propagation matrix between two antenna arrays, with applications to RIS-based over-the-air beamforming. In *Proc. of IEEE VTC-Spring*, pages 1–6. IEEE, Jun. 2022.

[56] Krishan K Tiwari and Giuseppe Caire. Power transfer between two antenna arrays in the near field. In *Proc. of IEEE VTC-Spring*, pages 1–6. IEEE, Jun. 2024.

[57] Krishan Kumar Tiwari and Giuseppe Caire. A new old idea: Beam-steering reflectarrays for efficient sub-thz multiuser MIMO. *Authorea Prepr.*, 2023.

[58] Krishan K Tiwari and Giuseppe Caire. A modular pragmatic architecture for multiuser MIMO with array-fed RIS. In *Proc. of. IEEE Int. Workshop Signal Process. Adv. Wireless Commun.*, pages 556–560. IEEE, Sep. 2024.

[59] Vahid Jamali, Antonia M Tulino, Georg Fischer, Ralf R Müller, and Robert Schober. Intelligent surface-aided transmitter architectures for millimeter-wave ultra massive MIMO systems. *IEEE Open J. Commun. Soc.*, 2:144–167, Dec. 2020.

[60] Vahid Jamali, Antonia M. Tulino, Georg Fischer, Ralf Müller, and Robert Schober. Reflect- and transmit-array antennas for scalable and energy-efficient mmwave massive MIMO. Feb. 2019.

[61] Giulia Torcolacci, Malte Schellmann, Ozgur Ercetin, and Davide Dardari. RIS-aided beamfocusing for active feeder architectures in near-field MU-MIMO. *Manuscript in preparation*, 2025.

[62] Emil Björnson, Özlem Tuğfe Demir, and Luca Sanguinetti. A primer on near-field beamforming for arrays and reconfigurable intelligent surfaces. In *Proc. of Asilomar Conf. Signals Syst. Comput.*, pages 105–112. IEEE, Oct. 2021.

[63] Constantine A Balanis. *Antenna theory: analysis and design*. John Wiley & Sons, 2016.

[64] Giulia Torcolacci, Anna Guerra, Haiyang Zhang, Francesco Guidi, Qianyu Yang, Yonina C Eldar, and Davide Dardari. Holographic imaging with XL-MIMO and RIS: Illumination and reflection design. *IEEE Journal of Selected Topics in Signal Processing*, Jun. 2024.

[65] David Tse and Pramod Viswanath. *Fundamentals of wireless communication*. Cambridge university press, 2005.

[66] Hyuckjin Choi, A Lee Swindlehurst, and Junil Choi. WMMSE-based rate maximization for RIS-assisted MU-MIMO systems. *IEEE Trans. Commun.*, Mar. 2024.

[67] Giulio Bartoli, Andrea Abrardo, Nicolo Decarli, Davide Dardari, and Marco Di Renzo. Spatial multiplexing in near field MIMO channels with reconfigurable intelligent surfaces. *IET Signal Processing*, 17(3):e12195, Mar. 2023.

[68] Ezio Biglieri, Robert Calderbank, Anthony Constantinides, Andrea Goldsmith, Arogyaswami Paulraj, and H Vincent Poor. *MIMO wireless communications*. Cambridge University Press, 2007.

[69] TIMES SNS. *Deliverable D2.3 - Definition of scenarios and KPI for hardware demonstration and PoC*, 2023.

[70] Anamaria Moldovan, Prasanth Karunakaran, Ian F Akyildiz, and Wolfgang H Gerstacker. Coverage and achievable rate analysis for indoor terahertz wireless networks. In *2017 IEEE International Conference on Communications (ICC)*, pages 1–7. IEEE, 2017.

[71] Akram Shafie, Nan Yang, Zhuo Sun, and Salman Durrani. Coverage analysis for 3D terahertz communication systems with blockage and directional antennas. In *2020 IEEE International Conference on Communications Workshops (ICC Workshops)*, pages 1–7. IEEE, 2020.

[72] Eiman Alotaibi and Biswanath Mukherjee. A survey on routing algorithms for wireless ad-hoc and mesh networks. *Computer networks*, 56(2):940–965, 2012.

[73] Djohara Benyamina, Abdelhakim Hafid, and Michel Gendreau. Wireless mesh networks design — a survey. *IEEE Communications Surveys & Tutorials*, 14(2):299–310, 2012.

[74] Qing Xia and Josep Miquel Jornet. Multi-hop relaying distribution strategies for terahertz-band communication networks: A cross-layer analysis. *IEEE Transactions on Wireless Communications*, 21(7):5075–5089, 2021.

[75] Zhengying Lou, Baha Eddine Youcef Belmekki, and Mohamed-Slim Alouini. Coverage analysis of hybrid RF/THz networks with best relay selection. *IEEE Communications Letters*, 27(6):1634–1638, 2023.

[76] Sara Cavallero, Kristi Qirjako, Roberto Verdone, and Chiara Buratti. A Multi-Hop Industrial IoT Network at THz Bands Using Contention-Based Access. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–6. IEEE, 2023.

[77] Parul Goyal, V. Parmar, and Rahul Rishi. Manet: Vulnerabilities, challenges, attacks, application. *International Journal of Computational Engineering Management*, 11:32–37, 01 2011.

[78] Qing Xia, Zahed Hossain, Michael Medley, and Josep Miquel Jornet. A link-layer synchronization and medium access control protocol for terahertz-band communication networks. *IEEE Transactions on Mobile Computing*, 20(1):2–18, 2019.

[79] Miguel Morales. *Deep Reinforcement Learning*. Manning Publications Co., 2020.

[80] Drago Hercog. *Communication protocols: principles, methods and specifications*. Springer Nature, 2020.

[81] Qing Xia and Josep Miquel Jornet. Cross-layer analysis of optimal relaying strategies for terahertz-band communication networks. In *2017 IEEE 13th international conference on wireless and mobile computing, networking and communications (WiMob)*, pages 1–8. IEEE, 2017.

[82] Saim Ghafoor, Noureddine Boujnah, Mubashir Husain Rehmani, and Alan Davy. MAC protocols for terahertz communication: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(4):2236–2282, 2020.

[83] Sven Gronauer and Klaus Diepold. Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review*, 55:895–943, 2022.

[84] Aurélien Géron. *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow*. " O'Reilly Media, Inc.", 2022.

[85] Mohit Sewak. *Deep Reinforcement Learning: Frontiers of Artificial Intelligence*, pages 95–108. Springer Singapore, 2019.

[86] Miguel Morales. *Grokking deep reinforcement learning*. Manning Publications, 2020.

[87] Dinh C. Nguyen et al. 6G Internet of Things: A Comprehensive Survey. *IEEE Internet of Things J.*, 9(1):359–383, 2022.

[88] Giampaolo Cuozzo, Enrico Testi, Salvatore Riolo, Luciano Miuccio, Gianluca Cena, Gianni Pasolini, Luca De Nardis, Daniela Panno, Marco Chiani, Maria-Gabriella Di Benedetto, Enrico Buracchini, and Roberto Verdone. Research directions and modeling guidelines for industrial internet of things applications. *IEEE Communications Standards Magazine*, pages 1–1, 2025.

[89] Muhammad Iqbal et al. Wireless sensor network optimization: Multi-objective paradigm. *Sensors*, 15(7):17572–17620, 2015.

[90] Alessia Tarozzi et al. Reinforcement Learning Based Backoff Management for Industry 5.0. In *Proc. IEEE Network Operations and Management Symposium (NOMS)*, pages 1–6, 2025.

[91] Muhammad Awais Jadoon et al. Deep Reinforcement Learning for Random Access in Machine-Type Communication. In *Proc. IEEE Wireless Commun. and Netw. Conf. (WCNC)*, pages 2553–2558, 2022.

[92] Taegyeom Lee and Ohyun Jo. Learning Backoff: Deep Reinforcement Learning-Based Wireless Channel Access. *IEEE Syst. J.*, 18(1):351–354, 2024.

[93] Zhichao Zheng et al. An adaptive backoff selection scheme based on Q-learning for CSMA/CA. *Wirel. Netw.*, 29(4):1899–1909, 2023.

[94] Mostefai Mohammed, Zibouda Aliouat, and Yacine Slimani. Intelligent CW Selection Mechanism Based on Q-Learning (MISQ). *Ingénierie des Systèmes D Information*, 25:803–811, 2020.

[95] Christina Chaccour, Mehdi Naderi Soorki, Walid Saad, Mehdi Bennis, Petar Popovski, and Mérouane Debbah. Seven Defining Features of Terahertz (THz) Wireless Systems: A Fellowship of Communication and Sensing. *IEEE Communications Surveys & Tutorials*, 24(2):967–993, 2022.

[96] Sara Cavallero, Nicolò Decarli, Giampaolo Cuozzo, Chiara Buratti, Davide Dardari, and Roberto Verdone. Terahertz networks for future Industrial Internet of Things. *ITU J. Futur. Evol. Technol.*, 2023, 4, 196–208.

[97] Sara Cavallero, Chiara Buratti, Alexey Tsarev, Giampaolo Cuozzo, Emil Khayrov, Yuliya Gaidamaka, and Roberto Verdone. Applying Carrier Sense Multiple Access to Industrial IoT at Terahertz Frequencies. *IEEE Internet of Things Journal*, 11(7):11986–11999, 2024.

[98] Sara Cavallero, Kristi Qirjako, Roberto Verdone, and Chiara Buratti. A Multi-Hop Industrial IoT Network at THz Bands Using Contention-Based Access. In *IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–6, 2023.

[99] Kai Arulkumaran et al. A Brief Survey of Deep Reinforcement Learning. *IEEE Signal Process. Mag., Special Issue on Deep Learn. for image understanding*, 2017.

[100] Volodymyr Mnih et al. Asynchronous Methods for Deep Reinforcement Learning. *CoRR*, abs/1602.01783, 2016.

[101] Doina Precup David Silver, Satinder Singh and Richard S. Sutton. Reward is enough. *ScienceDirect*, 12 May 2021.

[102] Antonin Raffin et al. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *J. of Machine Learning Research*, 22(268):1–8, 2021.

[103] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.

[104] Chong Han, Yongzhi Wu, Zhi Chen, and Xudong Wang. Terahertz communications (TeraCom): Challenges and impact on 6G wireless systems. *arXiv preprint arXiv:1912.06040*, 2019.

[105] Ho-Jin Song and Namyoon Lee. Terahertz communications: Challenges in the next decade. *IEEE Transactions on Terahertz Science and Technology*, 12(2):105–117, 2021.

[106] M. Qurratulain Khan, Abdo Gaber, Philipp Schulz, and Gerhard Fettweis. Machine learning for millimeter wave and terahertz beam management: A survey and open challenges. *IEEE Access*, 11:11880–11902, 2023.

[107] Ahmed Alkhateeb, Iz Beltagy, and Sam Alex. Machine learning for reliable mmwave systems: Blockage prediction and proactive handoff. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1055–1059, 2018.

[108] Syed Hashim Ali Shah, Manali Sharma, and Sundeep Rangan. LSTM-based multi-link prediction for mmwave and sub-thz wireless systems. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2020.

[109] Anders E. Kalør, Osvaldo Simeone, and Petar Popovski. Prediction of mmwave/thz link blockages through meta-learning and recurrent neural networks. *IEEE Wireless Communications Letters*, 10(12):2815–2819, 2021.

[110] Shunyao Wu, Muhammad Alrabeiah, Chaitali Chakrabarti, and Ahmed Alkhateeb. Blockage prediction using wireless signatures: Deep learning enables real-world demonstration. *IEEE Open Journal of the Communications Society*, 3:776–796, 2022.

[111] Martin H Nielsen, Chia-Yi Yeh, Ming Shen, and Muriel Médard. Blockage prediction in directional mmwave links using liquid time constant network. *arXiv preprint arXiv:2306.04997*, 2023.

[112] Ramin Hasani, Mathias Lechner, Alexander Amini, Daniela Rus, and Radu Grosu. Liquid time-constant networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7657–7666, 2021.

[113] Gouranga Charan and Ahmed Alkhateeb. Computer vision aided blockage prediction in real-world millimeter wave deployments. In *2022 IEEE Globecom Workshops (GC Wkshps)*, pages 1711–1716. IEEE, 2022.

[114] Ahsan Raza Khan, Iftikhar Ahmad, Lina Mohjazi, Sajjad Hussain, Rao Naveed Bin Rais, Muhammad Ali Imran, and Ahmed Zoha. Latency-aware blockage prediction in vision-aided federated wireless networks. *Frontiers in Communications and Networks*, 4:1130844, 2023.

[115] Shunyao Wu, Chaitali Chakrabarti, and Ahmed Alkhateeb. Lidar-aided mobile blockage prediction in real-world millimeter wave systems. In *2022 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 2631–2636. IEEE, 2022.

[116] Roman Alieiev, Jiri Blumenstein, Roman Maršalek, Thorsten Hehn, Andreas Kwoczek, and Thomas Kürner. Sensor-based predictive communication for highly dynamic multi-hop vehicular networks. In *2017 25th European Signal Processing Conference (EUSIPCO)*, pages 633–637, 2017.

[117] Yongjun Ahn, Jinhong Kim, Seungnyun Kim, Kyuhong Shim, Jiyoung Kim, Sangtae Kim, and Byonghyo Shim. Toward intelligent millimeter and terahertz communication for 6G: Computer vision-aided beamforming. *IEEE Wireless Communications*, 30(5):179–186, 2023.

[118] 5G-ACIA White Paper. *Key Use Cases and Requirements*, 2020.

[119] TIMES SNS. *Deliverable D2.1 - Definition of use cases, KPIs, and scenarios for channel measurements*, 2023.

[120] Tommaso Zugno, Lutfi Samara, Mate Boban, Per Hjalmar Lehne, and Thomas Kürner. Terahertz communications for industrial manufacturing: a use case analysis. In *2023 IEEE Conference on Standards for Communications and Networking (CSCN)*, pages 136–141, 2023.

[121] Andreas F Molisch, Henrik Asplund, Ralf Heddergott, Martin Steinbauer, and Thomas Zwick. The COST259 directional channel model-part I: Overview and methodology. *IEEE Transactions on Wireless Communications*, 5(12):3421–3433, 2006.

[122] Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017.

[123] Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *Neural Information Processing: 25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13-16, 2018, Proceedings, Part I 25*, pages 362–373. Springer, 2018.

[124] Luana Ruiz, Fernando Gama, and Alejandro Ribeiro. Gated graph recurrent neural networks. *IEEE Transactions on Signal Processing*, 68:6303–6318, 2020.

[125] Jinyin Chen, Xueke Wang, and Xuanheng Xu. GC-LSTM: Graph convolution embedded LSTM for dynamic network link prediction. *Applied Intelligence*, pages 1–16, 2022.

[126] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[127] Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.

[128] Andrea Abrardo, Alberto Toccafondi, and Marco Di Renzo. Optimal ris design by using s-parameter multiport network model. In *2024 IEEE INC-USNC-URSI Radio Science Meeting (Joint with AP-S Symposium)*, pages 50–50, 2024.