

Site-Specific RIS Deployment in Cellular Networks via Calibrated Ray Tracing

Sina Beyraghi^{†*}, Javad Shabanpour[‡], Giovanni Geraci^{*}, Paul Almasan[†], and Angel Lozano^{*}

^{*}Univ. Pompeu Fabra, Barcelona, Spain [†]Telefónica Research, Barcelona, Spain

[‡]Univ. Aalto, Espoo, Finland

Abstract—This work introduces a fully-automated RIS deployment strategy validated through a digital twin, powered by Sionna ray tracing, of a UK city. On a scene calibrated with measured data, the method jointly optimizes RIS placement, orientation, configuration, and BS beamforming across 4G, 5G, and hypothetical 6G frequencies. Candidate RIS sites are identified via scattering-based rays, while user clustering reduces deployment overhead. Results show that meaningful coverage enhancement requires dense, large-aperture RIS deployments, raising questions about the practicality and cost of large-scale RIS adoption.

I. INTRODUCTION

Consistent connectivity is a key challenge in modern wireless systems, especially in dense urban environments where obstructions and multipath propagation degrade performance [1]. These issues worsen at higher frequencies due to increased pathloss and limited diffraction. Reconfigurable intelligent surfaces (RIS) have emerged as a promising solution. By reflecting signals toward underserved areas without additional transmit power, RIS offer a potentially energy-efficient means to enhance coverage [2], [3]. However, large-scale adoption by mobile network operators (MNOs) remains uncertain due to unclear cost-performance tradeoffs.

Indeed, RIS placement involves a high-dimensional combinatorial problem, with many candidate sites, orientation constraints, and site-specific propagation. Trial-and-error methods are unscalable. Tackling this challenge requires an automated and scalable framework, and an enticing enabler of such a framework is a digital twin providing a precise digital replica of the radio environment [4].

Previous studies on RIS placement and design offer valuable insights but suffer from key limitations. Many focus on single-user equipment (UE) or single-cell scenarios using idealized or stochastic models that are unable to capture site-specific details [5]–[9]. Others use simplified reflection models

that ignore polarization and angular losses [2], [10]–[12], or neglect practical constraints like building geometry, material properties, and orientation feasibility [13], [14]. Few combine ray tracing with empirical data, and none offer a scalable, automated solution across multiple frequencies and network layouts.

This paper presents a fully automated, data-driven framework for site-specific RIS deployment in cellular networks. A radio digital twin is applied to jointly optimize RIS placement, orientation, phase configuration, and base station (BS) beamforming. Unlike analytical or heuristic methods, the framework supports multi-cell, multi-band deployments and incorporates a calibrated ray-tracing engine to model material- and geometry-specific electromagnetic interactions.

The framework is validated using data from a UK commercial deployment. Simulations span 4G, 5G, and a hypothetical 6G system at 2 GHz, 3.5 GHz, and 10 GHz, respectively, using the open-source Sionna engine [15]. Material properties are calibrated with measured UE received power. The main contributions are as follows:

- Using measured UE data, material properties are calibrated to improve ray-tracing accuracy.
- A ray-based method jointly optimizes RIS placement, phase configuration, and BS beamforming, incorporating orientation and geometry constraints.
- Coverage improvements are evaluated across frequencies as a function of RIS aperture, element count, and density.
- A cost-performance analysis reveals that significant coverage gains require extensive deployment of large RIS units, questioning economic viability.

To foster reproducibility, the full simulation framework and RIS deployment algorithms are available in open source.¹

II. NETWORK, RIS, AND PROPAGATION MODELS

A. Cellular Network Model

The digital twin replicates the layout of a commercial cellular network deployed by a leading MNO. The focus area contains 12 BSs ranging from 18 to 56 m in height, each hosting three sectors for a total of 36 cells. The area spans 1340 m × 1390 m in the UK. Figure 1 shows a 3D rendering of the environment.

This work was in part supported by H2020-MSCA-ITN-2020 META WIRELESS (Grant Agreement: 956256), by the SNS JU Horizon Europe Project under Grant Agreement No. 101139161 (INSTINCT), by the 6G-Machine Intelligence based Radio Access Infrastructure (6G-MIRAI) project under Grant Agreement No. 101192369, by the Spanish Research Agency through grants PID2021-123999OB-I00 and CNS2023-145384, by ICREA, by the Maria de Maeztu Units of Excellence Programme (CEX2021-001195-M), and by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union NextGenerationEU through UNICO-5G I+D projects TSI-063000-2021-138 (SORUS-RIS) and TSI-063000-2021-59 (RISC-6G).

¹<https://github.com/Telefonica-Scientific-Research/DDRDR>




Fig. 1: 3D visualization of the area produced with OpenStreetMap, with black pins indicating BS sites.

1) *Antennas*: Each BS features a planar array of vertically polarized elements. The configuration of each BS is specified by its vertical tilt angle and its horizontal bearing angle, both determined by the MNO. A single radio-frequency chain feeds each antenna element and it abides by the 3GPP radiation pattern, with half-power beamwidths of 65° in azimuth and 10° in elevation [16].

2) *Radio Deployments*: Three systems are considered:

- 4G at 2 GHz (FR1)
- 5G at 3.5 GHz (FR1)
- 6G at 10 GHz (FR3)

For each, the parameters (adopted from commercial products, either existing or under development) are detailed in Table I. Note that only single polarization is assumed, resulting in element counts that are half those for dual polarization. Reusing existing site grids is essential for the incorporation of new spectrum to be economically feasible, as further sites would lead to considerable expenses and prolong deployment timelines [17]. Hence, The same BS coordinates are applied to all three systems.²

3) *Beamforming Codebook*: Each BS adopts a two-dimensional discrete Fourier transform (DFT) beamforming codebook per system. This codebook specifies a set of orthogonal beams on given angular directions; each such beam is characterized by the Kronecker product of two one-dimensional DFT vectors, one for azimuth and one for elevation, namely

$$\mathbf{w}_{m_h, m_v} = \mathbf{w}_{m_h} \otimes \mathbf{w}_{m_v}, \quad (1)$$

where m_h, m_v are the horizontal and vertical beam indices.

²The 6G system operates at a higher frequency and with more bandwidth than its 4G/5G counterparts, but it is configured with a lower transmit power. This reflects a role for 6G as a non-standalone capacity-boosting layer with possibly discontinuous coverage.

TABLE I: System features [18]

Feature	4G	5G	6G
Carrier frequency [GHz]	2	3.5	10
Bandwidth [MHz]	20	100	200
Sectors per site	3	3	3
Polarization	Vertical	Vertical	Vertical
Planar array topology	2x2	4x8	4x16
Beamforming codebook size	4	32	64
TX power per cell [dBm]	43	49	44
TX power per subcarrier [dBm]	12.2	13.85	8.85
Number of subcarriers	1200	3276	3276
Noise power per subcarrier [dBm]	-132	-129	-126

4) *UEs*: UEs are deployed on a grid of $2\text{ m} \times 2\text{ m}$ tiles at a height of 1.5 m, with the received power averaged per tile. The focus is on outdoor UEs as the main use case for cellular networks, though the method extends to indoor scenarios given floor plans and outdoor-to-indoor ray tracing support. Each UE employs a single vertically polarized isotropic antenna and is served individually per time-frequency resource, isolating the effects of scheduling.

B. Physically consistent RIS Model

The physically consistent RIS model in [19] is adopted, which treats the RIS as a continuous reradiating surface with a spatial modulation function based on ray optics. This captures finite surface size, angular reradiation, and energy conservation. Precisely, the spatial modulation function $\Gamma(x_R, y_R)$ describes the complex transformation applied by the RIS at each surface point (x_R, y_R) , combining amplitude and phase modulation into [20]

$$\Gamma(x_R, y_R) = R\sqrt{\eta}A(x_R, y_R)e^{j\varphi(x_R, y_R)}, \quad (2)$$

where R accounts for roughness and loss, η is the surface efficiency, and $A(x_R, y_R)$ controls reradiated power.

C. Site-specific Ray Tracing and Coverage Calculation

The propagation among BSs, UEs, and RIS units is digitally replicated using Sionna [21], an open-source ray-tracing simulator. To balance realism and efficiency, the ray tracing employs a Fibonacci shoot-and-bounce method with 10^7 rays per cell, up to four bounces per ray, and it includes specular reflection, diffraction, and scattering.

To compute the large-scale channel gain between each BS and UE, the set of corresponding rays is considered. As advanced, the region of interest is partitioned into square tiles $C_{p,q}$ of size $2\text{ m} \times 2\text{ m}$, where each tile corresponds to a UE. For each BS sector t and DFT beam index m , the directional channel gain at tile (p, q) is computed as

$$G_{t,m,p,q} = \frac{1}{4} \iint_{C_{p,q}} |\mathbf{h}_{t,p,q}^*(x, y) \mathbf{w}_{t,m}|^2 dx dy, \quad (3)$$

where $\mathbf{h}_{t,p,q}(x, y)$ is the channel vector at position (x, y) , obtained from all ray types (LoS, reflected, diffracted, and

scattered), $w_{t,m}$ is the DFT precoding vector, and 4 is the tile area.

The reference signal received power (RSRP) for tile (p, q) , served by BS sector t and beam index m , is then

$$\text{RSRP}_{p,q} = G_{t,m,p,q} P_t, \quad (4)$$

where P_t is the transmit power per subcarrier, as specified in Table I. The UE at the center of each tile selects its serving BS and beam by maximizing the RSRP across all sectors and beam indices.

III. SCENE MATERIAL CALIBRATION

Accurate ray tracing requires not only precise geometry but also realistic electromagnetic material properties—something particularly difficult at city scale. While prior work focused on small, controlled environments [21], we target dense urban areas where default material settings lead to significant mismatches. Using a large outdoor measurement dataset from a UK city—covering BS configurations (e.g., orientation, tilt) and spatial received power—material parameters are calibrated to align simulated and experimental radio coverage.

Specifically, building material parameters are calibrated by minimizing the difference between simulated and measured RSRP values. Using the Sionna ray tracer and the actual BS configurations, coverage is first simulated assuming all buildings are made of concrete. Three material properties—relative permittivity, conductivity, and surface scattering—are treated as learnable variables. UE measurement data is aggregated over $10 \times 10 \text{ m}^2$ regions with at least 20 samples, enabling stable RSRP comparisons. For each region, surrounding buildings within 100 meters are jointly calibrated by comparing average simulated and measured RSRP. Gradients are computed using an Adam optimizer, and materials are iteratively updated. The optimization runs for 600 steps per cell, gradually covering the city while preserving already calibrated areas. To ensure meaningful updates, outliers and regions with persistent mismatches are filtered out. Region merging and error thresholds are applied to manage complexity and adapt to urban density variations.

The design of this calibration approach is shaped by both modeling simplifications and practical data limitations. Since the focus is on minimizing RSRP error rather than retrieving physically exact material values, the optimizer adjusts parameters to best match the model's predictions with measurements. Calibration is performed over spatial regions instead of individual samples to mitigate device-related RSRP variability. Regions with extreme mismatches are excluded, as these are typically due to missing scene geometry. The optimizer reacts adaptively: if simulated RSRP overshoots, material properties are driven toward free-space conditions; if it undershoots, reflectivity is increased; in both cases, within constrained bounds. Persistent errors at the extremes trigger region exclusion. And, to manage urban complexity, adjacent buildings are merged, and region-specific thresholds are applied based on density, acknowledging that calibration tolerance varies across the cityscape.

A. Validation

Material calibration was conducted over a central urban area measuring $1122 \text{ m} \times 710 \text{ m}$. To assess its effectiveness, 70 representative target regions were selected across the area. In each region, the average measured RSRP was compared against simulated values, both before and after calibration. A detailed analysis of prediction errors is provided in Fig. 2, with outlier regions excluded. Prior to calibration, RSRP predictions tended to underestimate coverage, with a mean error -5.69 dB , median of -5.37 dB , and standard deviation of 5.71 dB . After calibration, the error distribution narrowed substantially, with a near-zero mean of -0.32 dB , median of -0.13 dB , and standard deviation of 2.57 dB , confirming the effectiveness of the calibration approach.




Fig. 2: Distribution of RSRP prediction errors.

IV. DATA-DRIVEN RIS DEPLOYMENT

This section presents the data-driven strategy for large-scale RIS deployment on the digital twin.

A. Coverage Evaluation and UE Clustering

The baseline network coverage is evaluated by deploying BSs and distributing UEs on a grid, as described in Sec. II. Fig. 3 shows the RSRP heatmaps. UEs with an RSRP below -100 dBm are declared in outage [22], resulting in outage rates of 2.85%, 1.65%, and 6.07% at 2 GHz, 3.5 GHz, and 10 GHz, respectively. The goal of strategic RIS deployment is to improve RSRP at these outage locations, but deploying a dedicated RIS for each UE is infeasible due to the large number and distribution of outage UEs.

To address this scalability challenge, nearby outage UEs are grouped into clusters so that each RIS serves a cluster instead of an individual UE. A hierarchical BIRCH algorithm [23] is employed, building compact clustering feature trees. Each cluster is represented by

$$\left[U, \sum_{i=1}^U \mathbf{s}_i, \sum_{i=1}^U \|\mathbf{s}_i\|^2 \right], \quad (5)$$

where U is the number of UEs and $\mathbf{s}_i \in \mathbb{R}^2$ denotes the 2D position of UE i (center of a tile). The threshold T is a key design parameter that defines the maximum allowable Euclidean distance (in meters) between data points within a



Fig. 3: RSRP heatmaps across the considered urban area for 4G, 5G, and 6G.




Fig. 4: Number of clusters formed by the BIRCH algorithm vs. parameter T across the considered frequency bands.

cluster in the BIRCH algorithm. A smaller T yields more clusters and RISs, while a larger T reduces the count but enlarges coverage per RIS. Fig. 4 illustrates how the cluster count changes with T across frequencies.

B. RIS Candidate Locations

This section presents the algorithm for the identification of candidate RIS locations. For each cluster of outage UEs produced by BIRCH, one RIS is deployed to enhance coverage by aligning with the reflection point of relevant scattered rays. This study assumes RIS units can be deployed on any building surface with a valid scattering point, providing an upper bound on performance by ignoring practical constraints such as accessibility or regulations.

The approach identifies candidate RIS sites by activating only single-bounce scattering rays and increasing ray shootings to 3×10^7 for dense angular sampling.

Among the rays that reach a given cluster centroid, the ones retained are those that (i) reach it after one scattering event and (ii) create a valid LoS path from the BS via the scatter point. Candidate RIS locations are then determined from the corresponding scatter points. If multiple candidates are found, the one with the shortest 3D distance to the cluster

centroid is selected. If the centroid's RSRP increases, UE-level gains within the cluster are assessed. A cluster is marked as RIS-effective if at least 60% of its UEs experience an RSRP improvement. This balances RIS location efficiency with the computational complexity of the optimization.

C. RIS Configuration and BS Beamforming

Once a candidate location is selected, the RIS must be oriented, configured, and incorporated into the network.

- Each RIS is installed along the building wall at the designated location, face outwards.
- For each combination of BS, UE, and RIS locations, the spatial modulation function in Sec. II-B is computed.
- The serving BS selects the beam that maximizes the RSRP at the RIS location by iterating over all available DFT beams.

D. Re-Clustering and Re-Association

For the remaining outage UEs for which the candidate RIS locations proved ineffective, two fallback strategies are employed: re-clustering and RIS re-association.

1) *Re-Clustering*: The remaining outage UEs are regrouped with BIRCH using a smaller threshold T to form more compact clusters. The new threshold is determined iteratively to optimize the trade-off between coverage and deployment cost. For each new cluster, the centroid point is extracted, and the method is reapplied to find the RIS location. Although lowering T improves performance, it also increases the number of RIS units, raising capital costs for the MNO.

2) *RIS Re-Association*: Some outage UEs may be in areas where existing RIS units are effective but were not initially assigned to their cluster. To avoid redundant deployments, these UEs can be reassigned to existing RISs if a LoS link is feasible. The process consists of four steps:

- 1) Determine which deployed RIS units maintain a LoS link to each outage UE.
- 2) For each potential RIS, identify the BSs that have a LoS connection to it.
- 3) From the feasible RIS options, select the one closest to the UE.




Fig. 5: RSRP enhancement from deploying the maximum number of RIS.




Fig. 6: RSRP enhancement vs. number of deployed RIS.

- 4) From the BSs linked to the chosen RIS, pick the one nearest to it.

This process refines the association of UEs, RISs, and BSs, boosting RSRP while avoiding further RIS deployments.

V. SIMULATION RESULTS

For preliminary testing of the deployment method, a very large RIS aperture (11.24×11.24 m) is considered, with half-wavelength element spacing. Outage UEs are clustered using BIRCH with a threshold of $T = 15$ m.

1) *4G at 2 GHz*: Outage UEs represent 2.85% of the total population. With $T = 15$ m, 246 clusters are formed, each associated with one RIS. This initial RIS placement brings 68.85% of outage UEs back into coverage. Re-clustering the residual UEs with $T = 10$ m recovers an extra 2.73%, while RIS re-association further improves coverage by 6.61%, resulting in a total recovery of 78.97% of outage UEs. Figure 5a illustrates the CDF of the enhanced RSRP values.

2) *5G at 3.5 GHz*: Outage UEs constitute 1.65% of the entire network. Using the same method, RIS deployment initially restores coverage for 66.10% of these UEs. Applying re-clustering with $T = 10$ m adds another 3.80%, while RIS re-association contributes an additional 5.67%, resulting in an overall recovery of 75.25%. Figure 5b shows these improvements, accomplished with 197 deployed RIS units.

3) *6G at 10 GHz*: At this frequency, 393 RIS units are deployed to address the 6.07% of UEs initially in outage. The RIS deployment alone restores coverage for 46.02% of these UEs, with re-clustering adding another 6.02%. In particular, RIS re-association is especially effective in this scenario, recovering an additional 12.55%, for a total of 64.59% of outage UEs brought back into coverage, as shown in Fig. 5c.

While these initial findings show notable coverage gains, these entail a very large RIS aperture and many units, as each cluster is assigned a separate RIS. The remainder of this section evaluates how restricting aperture size and deploying RISs for only a subset of clusters influences performance.

A. Impact of RIS Density

In the previous analysis, each cluster was associated with a separate RIS, resulting in 246, 197, and 393 RIS units at 2, 3.5, and 10 GHz, respectively. However, many clusters are quite small—over half contain fewer than six UEs—reflecting the spatial sparsity of outage areas. To address this, clusters are now ranked by their UE count, and RIS units are incrementally allocated to the top- N clusters as N varies. Figures 6a, 6b, and 6c show the resulting CDFs of RSRP for different values of N , showing that prioritizing larger clusters helps balance coverage gains against deployment costs. Nevertheless, even with prioritization, achieving significant improvements still needs deploying a substantial number of RIS units.




Fig. 7: Percentage of outage UEs recovered with different RIS aperture sizes in various radio deployments.

B. Impact of RIS Apertures

Next, we study the impact of varying the RIS aperture, which directly determines the number of RIS elements. The element spacing remains fixed at half a wavelength. This evaluation is conducted for the maximum number of RIS units previously identified, namely, 246, 197, and 393. Figure 7 presents the percentage of outage UEs recovered as a function of the aperture, with the corresponding number of RIS elements also indicated. The results show that performance is more robust to variations in aperture than to changes in RIS deployment density.

VI. CONCLUSION

This work has introduced an automated, data-driven framework for evaluating large-scale RIS deployments in cellular networks. By integrating site-specific ray tracing, outage user clustering, and ray-based heuristics, the framework jointly determines RIS placement, orientation, configuration, and BS beamforming. Performance assessments in 4G, 5G, and 6G scenarios, grounded in a calibrated digital twin of an urban environment, highlight the tradeoff between coverage gains and deployment costs. Meaningful coverage enhancement necessitates deploying many large RIS units per square kilometer, raising questions about their cost-effectiveness in wide-area outdoor systems.

REFERENCES

- [1] S. Islam, A. Zainab Abdulsalam, B. Anil Kumar, M. Kamrul Hasan, R. Kolandaisamy, and N. Safie, "Mobile networks toward 5G/6G: Network architecture, opportunities and challenges in smart city," *IEEE open j. Commun. Soc.*, vol. 6, pp. 3082–3093, 2025.
- [2] M. Di Renzo, A. Zappone, M. Debbah, M.-S. Alouini, C. Yuen, J. de Rosny, and S. Tretakov, "Smart radio environments empowered by reconfigurable intelligent surfaces: How it works, state of research, and the road ahead," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 11, pp. 2450–2525, 2020.

- [3] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah, and C. Yuen, "Reconfigurable intelligent surfaces for energy efficiency in wireless communication," *IEEE Trans. Wireless Commun.*, vol. 18, no. 8, pp. 4157–4170, 2019.
- [4] W. Liu, Y. Fu, Z. Shi, and H. Wang, "When digital twin meets 6G: Concepts, obstacles, and research prospects," *IEEE Commun. Mag.*, vol. 63, no. 3, pp. 16–22, 2025.
- [5] C. You and R. Zhang, "Wireless communication aided by intelligent reflecting surface: Active or passive?" *IEEE Wireless Commun. Letters*, vol. 10, no. 12, pp. 2659–2663, 2021.
- [6] P.-Q. Huang, Y. Zhou, K. Wang, and B.-C. Wang, "Placement optimization for multi-IRS-aided wireless communications: An adaptive differential evolution algorithm," *IEEE Wireless Commun. Letters*, vol. 11, no. 5, pp. 942–946, 2022.
- [7] H. Lu, Y. Zeng, S. Jin, and R. Zhang, "Aerial intelligent reflecting surface: Joint placement and passive beamforming design with 3D beam flattening," *IEEE Trans. Wireless Commun.*, vol. 20, no. 7, pp. 4128–4143, 2021.
- [8] W. Huang, W. Ding, C. Kai, Y. Yi, and Y. Huang, "Joint placement and beamforming design for IRS-enhanced multiuser MISO systems," *IEEE Trans. Commun.*, vol. 70, no. 10, pp. 6678–6692, 2022.
- [9] G. Brancati, E. F. Pupo, O. Chukhno, N. Chukhno, M. Murrone, and G. Araniti, "Reconfigurable intelligent surface deployment and orientation in beyond 5G multicast networks," in *2023 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, 2023, pp. 1–6.
- [10] J. Shabanpour, V. Lenets, G. Leroose, S. Tretakov, and C. Simovski, "Engineering of intelligent reflecting surfaces: Reflection locality and angular stability," *IEEE Trans. Antennas Propag.*, vol. 72, no. 5, pp. 4198–4211, 2024.
- [11] S. Sandh, H. Radpour, B. Rainer, M. Hofer, and T. Zemen, "Ray tracing algorithm for reconfigurable intelligent surfaces," *arXiv preprint arXiv:2402.13034*, 2024.
- [12] H. Na, M. M. Taygur, and T. F. Eibert, "A multiple Huygens surface-based ray tracing framework with GPU acceleration," *IEEE Trans. Antennas Propag.*, vol. 72, no. 1, pp. 183–196, 2024.
- [13] C. You, B. Zheng, W. Mei, and R. Zhang, "How to deploy intelligent reflecting surfaces in wireless network: BS-side, user-side, or both sides?" *J. Commun. Net.*, vol. 7, no. 1, pp. 1–10, 2022.
- [14] Z. Kang, C. You, and R. Zhang, "IRS-aided wireless relaying: Deployment strategy and capacity scaling," *IEEE Wireless Commun. Letters*, vol. 11, no. 2, pp. 215–219, 2022.
- [15] J. Hoydis, F. A. Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna RT: Differentiable ray tracing for radio propagation modeling," *arXiv:2303.11103*, 2023.
- [16] 3GPP Technical Report 38.901, "Study on channel model for frequencies from 0.5 to 100 GHz (Release 16)," Dec. 2019.
- [17] H. Holma, H. Viswanathan, and P. Mogensen, "Extreme massive MIMO for macro cell capacity boost in 5G-advanced and 6G," *Nokia Bell Labs, White paper*, 2021.
- [18] D. López-Pérez, N. Piovesan, and G. Geraci, "Capacity and power consumption of multi-layer 6G networks using the upper mid-band," *arXiv preprint arXiv:2411.09660*, 2024.
- [19] V. Degli-Esposti, E. M. Vitucci, M. D. Renzo, and S. A. Tretakov, "Reradiation and scattering from a reconfigurable intelligent surface: A general macroscopic model," *IEEE Trans. Antennas Propag.*, vol. 70, no. 10, pp. 8691–8706, 2022.
- [20] E. M. Vitucci, M. Albani, S. Kodra, M. Barbiroli, and V. Degli-Esposti, "An efficient ray-based modeling approach for scattering from reconfigurable intelligent surfaces," *IEEE Trans. Antennas Propag.*, vol. 72, no. 3, pp. 2673–2685, 2024.
- [21] J. Hoydis, F. A. Aoudia, S. Cammerer, F. Euchner, M. Nimier-David, S. T. Brink, and A. Keller, "Learning radio environments by differentiable ray tracing," *IEEE Trans. Mach. Learn. Commun. Netw.*, vol. 2, pp. 1527–1539, 2024.
- [22] Teltonika Networks Wiki, "RSRP and RSRQ." [Online]. Available: https://wiki.teltonika-networks.com/view/RSRP_and_RSRQ
- [23] T. Zhang, R. Ramakrishnan, and M. Livny, "BIRCH: An efficient data clustering method for very large databases," in *ACM SIGMOD Int'l Conf. Management of Data*, 1996, p. 103–114.