

Fully Differentiable Ray Tracing via Discontinuity Smoothing for Radio Network Optimization

J erome Eertmans, Laurent Jacques, Claude Oestges

ICTEAM, Universit  catholique de Louvain, Louvain-la-Neuve, Belgium, jerome.eertmans@uclouvain.be

Abstract—Recently, Differentiable Ray Tracing has been successfully applied in the field of wireless communications for learning radio materials or optimizing the transmitter orientation. However, in the frame of gradient-based optimization, obstruction of the rays by objects can cause sudden variations in the related objective functions or create entire regions where the gradient is zero. As these issues can dramatically impact convergence, this paper presents a novel Ray Tracing framework that is fully differentiable with respect to any scene parameter, but also provides a loss function continuous everywhere, thanks to specific local smoothing techniques. Previously non-continuous functions are replaced by a smoothing function, that can be exchanged with any function having similar properties. This function is also configurable via a parameter that determines how smooth the approximation should be. The present method is applied on a basic one-transmitter-multi-receiver scenario, and shows that it can successfully find the optimal solution. As a complementary resource, a 2D Python library, DiffeRT2d, is provided in Open Access, with examples and a comprehensive documentation.

Index Terms—Ray Tracing, Differentiability, Wireless Communications, Optimization.

I. INTRODUCTION

The field of radiocommunications has witnessed unprecedented growth over the past few decades, largely fueled by the increase of data-intensive applications, as well as the need to access the Internet from nearly everywhere. This accessibility constraint, in addition to the other constraints faced by telecom operators, makes the task of placing antennas even more complicated. With the emergence of dynamic scenarios, such as vehicle-to-vehicle communications, it is also important to understand how to update the current configuration, *e.g.*, the orientation of the transmitting antenna, with respect to object displacements.

Traditionally, the design, analysis, and optimization of communication networks have relied on mathematical models, simulations, and measurements. Channel models are mainly either deterministic or stochastic. Based on measurements, empirical models often try to capture the general impact of the scene parameters on a given metric. For example, one could use the COST 231 Hata model [1] to estimate the path loss between two nodes in an urban environment. This model only needs a limited amount of inputs, and considers a very general urban scenario, with no specific assumption on how the buildings are placed. Conversely, deterministic models, like Ray Tracing (RT) or the Method of Moments, require a very precise input model, *e.g.*, the location of every building in

the scene, to provide a very accurate simulation of how radio waves propagate.

While these approaches have been invaluable in the development of modern communication systems, they face limitations in simulating evolving environments, such as iteratively searching for the best antenna location or analyzing a radio link involving mobile objects like cars. Empirical models, owing to their simplicity, cannot be used for site-specific network optimization. Deterministic channel models, on the other hand, often have such a large computational cost that performing a simulation for each small variation in scene parameters could be incredibly time-consuming.

In optimization problems, it is increasingly common to utilize gradient descent methods, particularly since the emergence of automatic differentiation (AD) libraries like JAX [2]. This enables differentiation of nearly every programmable function in relation to one or multiple input parameters. As such, applying AD to RT may lead to the development of a Differentiable Ray Tracing (DRT) program, hence enabling network parameter optimization.

However, the convergence of an optimizer can be dramatically impacted by the properties of the computed gradient; for instance, if it vanishes in certain regions of the parameter space, the optimization can be stuck in a suboptimal solution. In DRT, this phenomenon is directly connected to the geometry of the scene, *e.g.*, object occlusion can create discontinuous parameter transitions or zero-gradient regions. This paper proposes to address this problem by introducing a novel smoothing technique that, when applied to RT, enables to programmatically avoid those issues.

This paper is organized as follows. First, Sec. II summarizes the work in DRT and the main challenges. Next, Sec. III provides a mathematical and precise definition of the problem addressed by this work. After having motivated the use of smoothing techniques in RT in Sec. IV, we present in Sec. V our DRT framework, published under the DiffeRT2d Python package¹. Finally, in Sec. VII, we demonstrate the capacity of our framework to improve the convergence rate of network optimization problems, before providing a few research perspectives in the concluding Sec. VIII.

II. RELATED WORK

In computer graphics (CG), DRT is a widely studied topic, particularly in the context of reconstructing scene parameters

¹GitHub repository: <https://github.com/jeertmans/DiffeRT2d>.

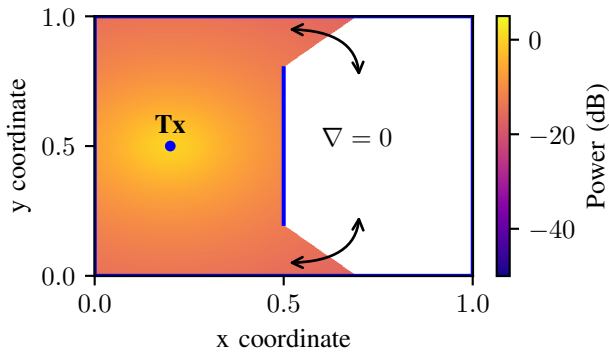


Fig. 1. Illustration of gradient discontinuities, as shown by the two arrows, and a zero-gradient region ($\nabla = 0$) in the objective function. The color gradient indicates the received power for every receiver coordinate in the scene.

based on images. In this field, many advances have been made, for example concerning the question of the presence of zero-gradient regions [3], also referred to as *plateaus*, or how to correctly carry out RT in order to take into account the abrupt transitions created by objects extremities [4].

DRT was mainly made possible thanks to AD and the development of toolboxes, such as JAX. Introduced by Wengert in 1964 [5] and actively developed for machine learning (ML) algorithms, AD is an alternative to symbolic calculation which, conversely to finite differences, computes the true derivatives up to machine precision [6].

While CG uses RT to perform optimization-based inverse rendering, communication operators rely on RT for channel simulation, as an alternative to actual measurements (often complicated to obtain). This difference explains why DRT is still a fairly young subject in the field of radiocommunications.

Recently, a novel tool, named Sionna [7], has been developed to provide DRT for radio propagation modeling. Sionna integrates AD within a ready-to-use RT tool that provides many features for radio channel modeling. Unfortunately, such a DRT tool can still suffer from vanishing, discontinuous gradients, mainly because RT programs partly rely on conditional statements (such as `if`, `else`), e.g., to characterize object occlusions. Fig. 1 illustrates this fact on a simple 2-D scene by showing how the change in visibility can cause discontinuities near the ends of a central wall. Moreover, all locations in the blank area receives zero contribution from the transmitting antenna, assuming only line-of-sight (LOS) path.

In the deep learning literature, the problem of vanishing gradients is well studied and has been circumvented with several strategies. For instance, the evaluation of the gradient can be artificially corrupted by noise to help the convergence in zero-gradient parameter domain [8]. A similar effect is obtained in the context of reverse rendering with DRT by blurring the rendered image with a specific kernel [3].

In this work, we propose another solution suited to radio wave propagation. We replace all discontinuous functions used in RT (e.g., in conditions or occlusions) with a smoothed, parametrized function. A direct advantage of this technique is

the possibility to progressively insert the objects in the scene by varying this smoothing parameter (see Sec. VII).

III. PROBLEM DEFINITION

Ray Tracing techniques approximate radio waves propagation by tracing rays between two or more communicating nodes. They can consider for instance one transmitting antenna (Tx) and one receiver antenna (Rx), bouncing rays off objects whenever an object is encountered (using reflection or diffraction models), and selecting all paths that finally reach a desired target. By computing all valid paths, geometrical optics approximations can be applied to determine, e.g., the received electrical field at a given position. Because there can be many paths, it is common to limit the maximum number of allowed interactions, e.g., to only simulate LOS and a limited number of reflections.

When searching for optimal parameters γ to a given problem, such as finding the best antenna position x to serve a fixed amount of users on multiple locations, a common approach is to use a gradient-based iterative optimizer maximizing a specific objective function $\mathcal{F}(\gamma)$, or minimizing its related loss function $\mathcal{L}(\gamma) = -\mathcal{F}(\gamma)$. A key element in this process is computing the gradient of \mathcal{L} with respect to the input parameters. Leveraging automatic-differentiation, computing this gradient becomes straightforward, as long as the objective function is differentiable.

However, limiting the number of paths in RT (e.g., to reduce the computational complexity) or not simulating certain types of interactions such as diffraction impact an object visibility and may introduce discontinuities in the objective function (see Fig. 1), i.e., \mathcal{L} is then only piecewise differentiable.

These discontinuities, as well as potential zero-gradient regions, present a real challenge to optimization. The next two sections explain how to build a fully DRT model and an innovative method to avoid discontinuities and zero-gradient regions. Overall, our approach improves the convergence of RT parameter optimization, as shown numerically in Sec. VII.

IV. SMOOTHING LOCAL DISCONTINUITIES

The use of smoothing was first motivated by the observation that a slight inaccuracy in the position of objects in a scene can introduce a sudden variation in the number of paths linking two nodes. In urban scenarios, where the precision is often of the order of a meter [9], this poses a real problem.

An obstacle, whose frontiers can be locally modeled with a unit step function

$$\theta(x) = \begin{cases} 1, & \text{if } x > 0, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

e.g., with $\theta(\mathbf{n}^\top \mathbf{x})$ in a given direction \mathbf{n} when the 3-D position \mathbf{x} varies, creates discontinuities in the scene. To overcome this, we replace θ with an approximating smoothed, parametrized function s .

In the context of a Ray Tracing program, this function models comparisons to eliminate discontinuities. The outputs 1 and 0 of the step function, therefore of the approximation




Fig. 2. Various smoothing functions plotted, for a fixed α parameter.

function, are associated, respectively, with *true* and *false* Boolean values.

A. Smoothing Properties

The approximation function, s , is not unique, and anyone can craft its own function, granted that it satisfies the following properties driven by the needs of our approach.

First, the function s must be parameterized with a parameter $\alpha \in \mathbb{R}^+$ whose value increases or decreases the smoothness of the function. In particular, we impose that $s : (x; \alpha) \in \mathbb{R} \times \mathbb{R}^+ \mapsto s(x; \alpha) \in [0; 1]$ respects

$$\lim_{\alpha \rightarrow \infty} s(x; \alpha) = \theta(x). \quad (2)$$

Additionally, the function must meet the following criteria:

- [C1] $\lim_{x \rightarrow -\infty} s(x; \alpha) = 0$ and $\lim_{x \rightarrow +\infty} s(x; \alpha) = 1$;
- [C2] $s(\cdot; \alpha)$ is monotonically increasing;
- [C3] $s(0; \alpha) = \frac{1}{2}$;
- [C4] and $s(x; \alpha) - s(0; \alpha) = s(0; \alpha) - s(-x; \alpha)$.

Imposing $s \in [0, 1]$ and C1 simplifies the conversion between floating-point numbers and Boolean values. The output can either be true ($s(x; \alpha) = 1$), false ($s(x; \alpha) = 0$), or fall in between these two extremes ($0 < s(x; \alpha) < 1$). To determine whether the output should be assigned to the true or false category, a threshold variable is used. Property C2 prevents local minima in minimization problems, while C3 and C4 are key symmetry properties for the framework presented in the next Section.

B. Examples of Smoothing Functions

With some inspiration from the ML community², the present framework proposes two smoothing functions that satisfy all constraints C1-C4, the sigmoid and the hard sigmoid, whose α -parametrization is reached by scaling

$$s(x; \alpha) = s(\alpha x). \quad (3)$$

²In ML, and in the DiffeRT2d Python package, the term *activation* function is used in place of smoothing function.




Fig. 3. Power map showing the normalized received power at every coordinate in the scene. In the case of approximation ($\alpha = 50$ and $s(x) = \text{hard_sigmoid}(x)$), one can observe that some zero-gradient regions have vanished, and that there is a small transition region right after the central wall. Only LOS and first order reflection paths are simulated.

The sigmoid is defined with a real-valued exponential

$$\text{sigmoid}(x; \alpha) = \frac{1}{1 + e^{-\alpha x}}, \quad (4)$$

and the hard sigmoid is the piecewise linear function defined by

$$\text{hard sigmoid}(x; \alpha) = \frac{\text{relu6}(\alpha x + 3)}{6}, \quad (5)$$

where

$$\text{relu6}(x) = \min(\max(0, x), 6). \quad (6)$$

The hard sigmoid smoothing function is much cheaper to evaluate than the sigmoid, but at the cost of a non-continuous derivative at $x = \pm 3/\alpha$. Choosing an appropriate smoothing function will be dictated by the trade-off between the desired smoothing (see Fig. 2) and its related computational cost.

V. RAY TRACING FRAMEWORK

Among the different possible RT implementations, the framework presented here aims at simulating the radio channel between pairs of points, the details of which are introduced in [10]. In contrast to Ray Launching, the objective here is to exactly determine all paths up to a certain order.

Therefore, based on a given scene and a visibility matrix between all objects, an exhaustive list of all possible path candidates is created. For each of these candidates, the path

can be calculated, for example using the image method (IM) or a more complete method (see Sec. V-B).

As discussed in [10], RT frameworks often separate the path tracing step into two consecutive tasks: determining path coordinates in a relaxed environment, *e.g.*, by assuming infinitely large objects, and using an appropriate path finding method (*e.g.*, IM). Then, the path is further validated, checking for objects intersection, valid reflection angles, and so on. By default, this validity check contains many discontinuities, caused by visibility changes, among others. Leveraging the mentioned approximation framework, and thus the use of an approximation function everywhere it is needed, a fully continuous path validity check can be constructed.

In a channel propagation context [10], the received electrical field in Rx can be computed from

$$\mathbf{E} = \sum_{\mathcal{P} \in \mathcal{S}} V(\mathcal{P}) (\overline{\mathbf{C}}(\mathcal{P}) \cdot \mathbf{E}(\mathcal{P}_1)), \quad (7)$$

where \mathcal{P} is one path among the set of all possible paths \mathcal{S} , $V(\mathcal{P})$ is a scalar value, usually 0 or 1, denoting the validity of a given path, and \cdot represents the matrix product between the dyadic coefficient, $\overline{\mathbf{C}}(\mathcal{P})$, and the received field at the first interaction point on the path, $\mathbf{E}(\mathcal{P}_1)$. The matrix $\overline{\mathbf{C}}(\mathcal{P})$ is the product of all Fresnel coefficients, path attenuation factors, polarization changes, and so on, from the first interaction to Rx. In LOS, $\overline{\mathbf{C}}(\mathcal{P})$ simplifies to the identity matrix.

In (7), two components notably contain many discontinuities: the validity function, $V(\mathcal{P})$, and the path itself. The following sections each discuss how to implement a continuous, but approximate, version. Fig. 3 illustrates the impact of approximation on the normalized received power.

A. Continuous Validity Check

A path can be considered invalid for many reasons: an object is blocking part of it, some coordinates are not correctly placed on the interacting objects, caused by first assuming infinitely large objects, and so on. Checking for this condition ultimately results in a Boolean, discontinuous, output.

To avoid this, the aforementioned smoothing function is applied in place of any conditional check. For example, the usual *greater than* comparison, $x_0 > x_1$, is replaced with $s(x_0 - x_1; \alpha)$ for some $\alpha > 0$. Using this, the validity check now becomes a function that returns a continuum of values between 0 and 1, where values close to 1 map to a valid path, and values close to 0 to an invalid one.

B. Continuous Path Finding

Depending on the simulated types of interaction (reflection, diffraction, scattering, etc.), the desired computational time, and other criteria, a variety of path finding methods exist. The present framework implements the following three methods:

- 1) IM, that can only handle specular reflections on planar surfaces;
- 2) Fermat-Path-Tracing³ (FPT) [11], that can simulate any number of diffractions or reflections;

³FPT is a personal designation, used to indicate minimization according to Fermat's principle, as opposed to MPT which minimizes another cost function.

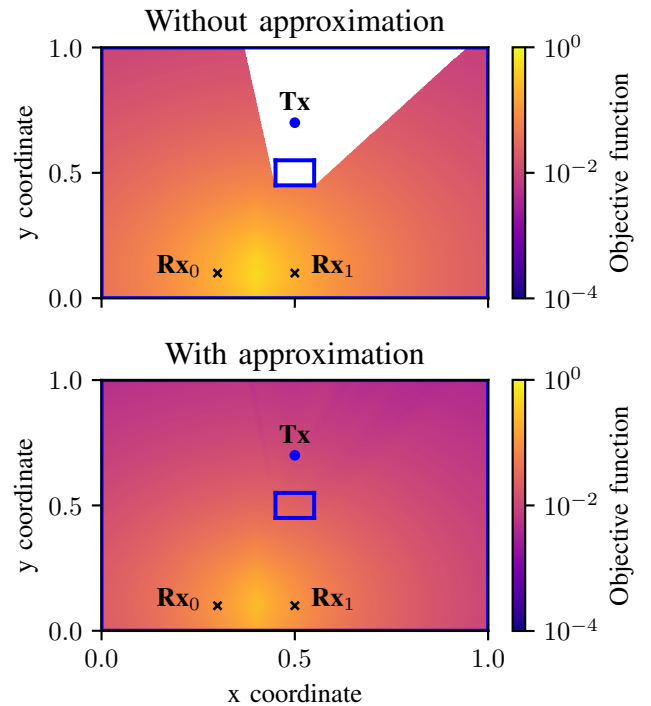


Fig. 4. Illustration of the optimization problem studied, starting with $\alpha = 1$ for the approximation case. Tx is at its initialization point. For each plot, the objective function (8) has been evaluated on all scene coordinates.

- 3) and Min-Path-Tracing (MPT) [10], that can simulate any number and any type of interactions.

All three methods find a path using an iterative approach, and have only one discontinuity: the convergence check. For this purpose, they all use the MPT loss function, as it evaluates to zero when the path is valid. Again, comparison to zero is approximated with the smoothing function s . More details can be found either in the paper describing the method [10], or in the code provided by the DiffeRT2d toolbox.

VI. DIFFERT2D AND 3D RAY TRACING

As mentioned in the introduction, the present framework is made available in open source on GitHub, and offers a self-contained Python module, DiffeRT2d, available on the Python Package Index. For demonstration purposes, the RT model is developed on two-dimensional (2D) scenes. However, it is important to note that the method is fully compatible with 3D scenes, with no assumptions made regarding the number of dimensions involved. Implementing a 3D DRT based on this framework is straightforward, although it would lengthen computation time, which is why proposing an efficient 3D implementation of this framework is a future work.

VII. APPLICATION TO ANTENNA LOCATION OPTIMIZATION

In Fig. 4, a basic one-transmitter-two-receiver optimization scenario is studied. To keep this example as simple as possible, only LOS paths have been simulated. The aim is to find the

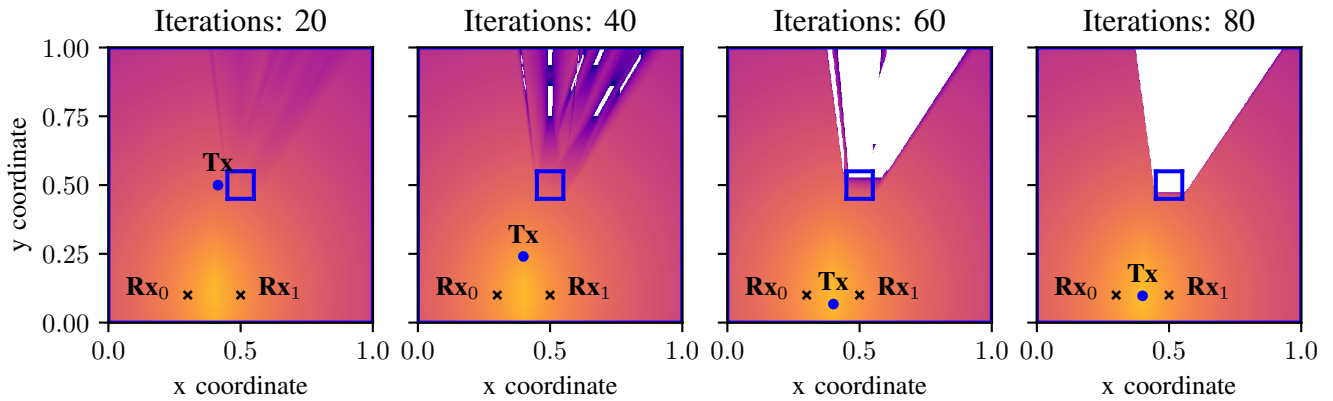


Fig. 5. Illustration of the different iterations converging towards the maximum of the objective function (8). The corresponding α values are: 2.54, 6.43, 16.30 and 41.32. Here, α follows a geometric progression, from 1 (see Fig. 4) to 100, with a number of values equal to the number of iterations (100).

optimum location for the transmitting antenna, to best serve both receivers. The objective function is therefore:

$$\mathcal{F}(x, y) = \min(P_{\text{Rx}_0}(x, y), P_{\text{Rx}_1}(x, y)), \quad (8)$$

where (x, y) are the Tx coordinates, $P_{\text{Rx}_i}(x, y)$ is the total power received by receiver i for a given antenna location, and the min operator takes the minimal values between the two powers, to avoid favoring one receiver.

As shown in Fig. 4, the problem has a zero-gradient region, which can cause convergence problems, especially if the starting point lies in the latter. By applying a sufficiently large smoothing, it is possible to reduce or even eliminate the zero-gradient region. This is the case when α is sufficiently small, as shown in the second plot on Fig. 4.

Thanks to this, it is possible to perform a gradient descent and reach the maximum of the objective function, as illustrated in Fig. 5. As this smoothing artifice is not realistic, it is important to increase the value of the α parameter after each iteration, in order to obtain a value of the objective function very close to that obtained without approximation.

By repeating the optimization on random scene configurations, with two, three or four RXs, results show that using our approximation can increase the convergence success rate by 1.5 to 2. Regarding the cases where convergence was already reached without approximation, using approximating still successfully converge between 92% and 98% of the time.

VIII. CONCLUSION

In this paper, a novel approach to Ray Tracing was proposed to avoid discontinuity-related issues, especially in optimization problems. By smoothing previously discontinuous conditions, the framework increases the convergence success.

The present solution offers a very customizable experience, where both the smoothing function, but also the smoothing rate, can be easily parametrized. In itself, the approximation enables to progressively introduce objects in the scene, as shown in Fig. 5. Here, all objects and functions are using the same α value, but a future research direction would be to

study the impact of this smoothing in every part of the code, to determine where it benefits the most. Also, the optimizer uses a very idiomatic progression for α values between steps, but one could imagine a more intelligent algorithm, that only applies smoothing when the gradient is too small.

Finally, it is possible to utilize the smoothing technique to simulate semi-transparent objects, as well as refraction and attenuation through them. Specifically, the use of an intermediary α value permits a path to pass through an object while simultaneously reflecting on it.

REFERENCES

- [1] E. Commission, D.-G. for the Information Society, and Media, "Propagation prediction models," in *COST Action 231 – Digital mobile radio towards future generation systems: Final Report*. Publications Office, 1999, ch. 4, pp. 134–135.
- [2] R. Frostig, M. Johnson, and C. Leary, "Compiling machine learning programs via high-level tracing," in *SysML Conference 2018*, 2018.
- [3] M. Fischer and T. Ritschel, "Plateau-reduced differentiable path tracing," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4285–4294.
- [4] T.-M. Li, M. Aittala, F. Durand, and J. Lehtinen, "Differentiable monte carlo ray tracing through edge sampling," *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, vol. 37, no. 6, pp. 222:1–222:11, 2018.
- [5] R. E. Wengert, "A simple automatic derivative evaluation program," *Communications of the ACM*, vol. 7, no. 8, pp. 463–464, 1964.
- [6] M. Bartholomew-Biggs, S. Brown, B. Christianson, and L. Dixon, "Automatic differentiation of algorithms," *Journal of Computational and Applied Mathematics*, vol. 124, no. 1, pp. 171–190, 2000, numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations.
- [7] J. Hoydis *et al.*, "Sionna rt: Differentiable ray tracing for radio propagation modeling," 2023.
- [8] A. Neelakantan, L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens, "Adding gradient noise improves learning for very deep networks," 2015.
- [9] M. Brovelli, M. Minghini, M. Molinari, and G. Zamboni, "Positional accuracy assessment of the openstreetmap buildings layer through automatic homologous pairs detection: The method and a case study," *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLI-B2, pp. 615–620, 06 2016.
- [10] J. Eertmans, C. Oestges, and L. Jacques, "Min-path-tracing: A diffraction aware alternative to image method in ray tracing," in *2023 17th European Conference on Antennas and Propagation (EuCAP)*, 2023, pp. 1–5.
- [11] F. Puggelli, G. Carluccio, and M. Albani, "A novel ray tracing algorithm for scenarios comprising pre-ordered multiple planar reflectors, straight wedges, and vertexes," *IEEE Transactions on Antennas and Propagation*, vol. 62, no. 8, pp. 4336–4341, 2014.