

Comparing Differentiable and Dynamic Ray Tracing: Introducing the Multipath Lifetime Map

J rome Eertmans*, Enrico Maria Vittuci , Vittorio Degli-Esposti , Laurent Jacques*, Claude Oestges*

*ICTEAM, Université catholique de Louvain, Louvain-la-Neuve, Belgium, jerome.eertmans@uclouvain.be

[†]DEI, University of Bologna, Bologna, Italy

Abstract—With the increasing presence of dynamic scenarios, such as Vehicle-to-Vehicle communications, radio propagation modeling tools must adapt to the rapidly changing nature of the radio channel. Recently, both Differentiable and Dynamic Ray Tracing frameworks have emerged to address these challenges. However, there is often confusion about how these approaches differ and which one should be used in specific contexts. In this paper, we provide an overview of these two techniques and a comparative analysis against two state-of-the-art tools: 3DSCAT from UniBo and Sionna from NVIDIA. To provide a more precise characterization of the scope of these methods, we introduce a novel simulation-based metric, the Multipath Lifetime Map, which enables the evaluation of spatial and temporal coherence in radio channels only based on the geometrical description of the environment. Finally, our metrics are evaluated on a classic urban street canyon scenario, yielding similar results to those obtained from measurement campaigns.

Index Terms—ray tracing, differentiable, dynamic, multipath lifetime, stationarity, propagation, comparison, modeling.

I. INTRODUCTION

As telecommunication systems operate in increasingly dynamic environments, such as those found in Vehicle-to-Vehicle (V2V) communications, the need for radio propagation models capable of handling rapidly varying channel conditions has become essential. These dynamic scenarios introduce significant challenges for propagation models, which must not only maintain high accuracy but also adapt efficiently to continuously changing environments.

Historically, accurate modeling techniques like Ray Tracing (RT) and the Method of Moments are static, requiring recalculations for even slight changes in scene configuration, such as the displacement of a receiver (RX). To address the challenges imposed by dynamic scenarios, Dynamic RT (DynRT) was introduced as a way to extrapolate future time instants of a moving scene based on a static Ray Tracing simulation and manually derived derivatives. More recently, automatic differentiation techniques have enabled the development of Differentiable RT (DiffRT) software. Although DynRT and DiffRT share some conceptual similarities, they differ significantly in design and objectives, which often leads to confusion, especially as **few papers directly compare the two approaches**. In regard to the scope for which dynamic scene extrapolations remain valid, this is primarily estimated from measurements, e.g., with the stationarity analysis in V2V scenarios, but this **does not extend well to complex changes in the scene or in the radio materials**.

In alignment with the aforementioned challenges, we summarize our major contributions as follows.

- 1) We provide a **first qualitative comparison of two state-of-the-art tools**, 3DSCAT [1], a RT tool developed at UniBo that incorporates DynRT [2], and Sionna [3], an open-source simulation framework from NVIDIA that features a DiffRT engine for radio propagation modeling;
- 2) We introduce a **novel simulation-based visual tool**, the Multipath Lifetime Map (MLM), along with two metrics, to accurately determine when DynRT can be applied. In contrast to decisions based on measurements, MLM provides a solution that is independent of frequency and accounts only for the geometry of the scene. This makes it possible to extend the solution to any radio-material properties.

The remainder of this paper is organized as follows. First, we review the key components of modern ray tracing techniques and the fundamental principles underlying them. We define the terms DynRT and DiffRT as they are used in this work and provide relevant historical context. Next, we conduct a comparative analysis of two state-of-the-art RT frameworks, 3DSCAT and Sionna, evaluating qualitative aspects based on their respective implementation. We then introduce our novel visual comparison tool, MLM, along with its associated metrics, and demonstrate its application in an urban street canyon scenario. Finally, we summarize the results presented and suggest future directions for research.

II. RAY TRACING FOR RADIO PROPAGATION

In the section, we briefly recall the important concepts behind RT in the context of radio propagation. Then, we review the two methods of interest, DynRT and DiffRT, by layering the fundamental principles.

A. Fundamentals

From a virtual representation of a given environment, known as a scene, a RT tool simulates wave propagation by dividing an emitted wavefront into a set of rays. These rays interact with objects in the scene, and by applying principles from Geometrical Optics (GO) and the Uniform Theory of Diffraction (UTD) [1], the exact coordinates of the ray paths can be determined. Depending on the desired balance between accuracy and computational efficiency, there are two primary families of methods for solving ray paths.




Fig. 1. Simplified two-step procedure taken by most RT software. First, the rays are traced based on GO and UTD. Then, the radio materials and EM properties are used to compute the EM fields.

For higher accuracy, exhaustive RT methods attempt to enumerate all possible ray paths using exact techniques, such as the image method for specular reflection or minimization-based methods for more complex interactions like diffraction [1], [4], [5]. On the other hand, ray-launching offers greater computational efficiency by emitting multiple rays from a specified transmitter (TX) node [1]. Each ray is allowed a fixed number of interactions with the scene’s objects before being terminated. The received signal is then approximated from the ray paths that fall within the vicinity of each RX.

For both approaches, the electromagnetic (EM) fields are typically computed in a separate step (see Fig. 1), as the ray path directions are independent of the frequency or antenna characteristics. A full recalculation of all ray paths for a given scene is referred to as a **snapshot**.

In dynamic environments, where ray path coordinates may vary smoothly with respect to changes in object positions, there is potential to bypass the need for full snapshot recalculations. Instead, local derivatives could be used to extrapolate the changes in the perceived EM fields. This concept will be explored further in the following sections.

B. Dynamic and Differentiable Ray Tracing

Although the first occurrence of the term “*dynamic ray tracing*” seems to date back to the early 1980s [6], when it was used to model the propagation of seismic waves, the use of DynRT as we define it in this article appeared only in the early 2020s [2], [7]. Meanwhile, the term has also been borrowed by both the radio propagation [8] and computer graphics [9] communities, but for very different purposes. Finally, while RT used to study earthquakes involved solving differential equations, RT presented today is derived from GO and UTD and accurately models wave propagation.

The primary motivation behind DynRT comes from a simple observation: in dynamic scenarios, such as when some RXs are moving down a street, most of the environment remains static. Traditional RT tools do not take this into account, requiring a complete recalculation of a new snapshot every time the scene changes. However, RT is structured so that the types of ray paths affecting a particular region of the scene remain consistent within a local neighborhood. The size and shape of these regions are determined by the shadowing and visibility effects of objects in the environment. To address this, [2] proposed a method for extrapolating new RT simulations from a static snapshot as long as the user (or any dynamic object)

remains within the same visibility region, i.e., the region where multipath structure remains the same.

As RT and EM field calculations can be treated as separate processes, DynRT often reduces to manually deriving the first—and sometimes second-order—derivatives of the ray path coordinates with respect to the variables of interest, such as the position of a moving RX. These derivatives can then be used in a Taylor expansion to predict future changes in the ray paths.

Although the concept of automatic differentiation (AD) is not new [10], the recent rise of DiffRT software has been propelled by the development of highly efficient, GPU-optimized Machine Learning (ML) libraries such as PyTorch and TensorFlow for Python. At its core, AD applies the chain rule and known derivatives to recursively compute exact gradients. These libraries provide a differentiable framework, allowing users to compute the gradient of **any output** with respect to **any input** parameter at a computational cost typically proportional to the function being evaluated. This makes them highly effective for ML tasks, but they have also gained attention in the radio propagation community, most notably with the release of Sionna [3].

While DiffRT could theoretically be used to compute local derivatives for extrapolating future snapshots, similar to DynRT, its primary applications lie in optimizing networks through gradient-based methods, and learning material properties of the radio environment from simulation data [11].

C. Determining the Scope of Snapshot Extrapolation

As previously stated, the duration over which a snapshot can be extrapolated on the basis of a previous snapshot is closely linked to the concepts of coherence time, denoted as T_c , and stationarity analysis [2]. However, these quantities are typically obtained through measurements, which makes the process time-consuming, costly, and not particularly scalable to large scenes. Moreover, the results are dependent on the frequency utilized, the type of antenna, and other variables that may vary during a simulation, making it challenging to identify the scope of application of these extrapolations. In Sec. IV, we present a novel approach for determining the scope of application of these extrapolations based solely on a geometry-based RT simulation.

III. COMPARISON

In this section, we present a qualitative comparison between DynRT and DiffRT, using 3DSCAT and Sionna as representative software for each method. Since both tools theoretically produce identical derivatives, a quantitative comparison would be less informative. Furthermore, performance differences are more likely influenced by the target hardware or compiler optimizations rather than the methods themselves.

A. Dynamic Ray Tracing with 3DSCAT

The 3DSCAT software, developed at UniBo, is a RT tool written in C++ that outputs ray path coordinates based on a specified scene configuration. Snapshot extrapolation using

DynRT and subsequent EM field calculations are performed separately with MATLAB code.

B. Differentiable Ray Tracing with Sionna

Sionna is a Python-based library for link-level simulations developed by NVIDIA, built on top of TensorFlow, an efficient, GPU-compatible AD framework. This framework enables Sionna to compute derivatives for nearly any scene parameter, making it highly suitable for various radio-focused ML applications.

C. Pros and Cons

Differentiating by hand, as typically done with DynRT, is laborious and prone to human error. Additionally, derivatives are often obtained under specific assumptions, e.g., a user moving in a fixed area, and thus remain valid only within a local region, limiting their applicability in general optimization contexts.

Conversely, AD is appealing because it automates the complex and error-prone process of computing derivatives and often includes built-in solutions compatible with GPUs or TPUs. This allows derivatives to be easily recomputed after significant environmental changes and at various points throughout the scene. However, using AD often restricts developers to a limited set of operations, sometimes leading to more intricate programming solutions and requiring the same AD framework throughout the codebase. While this consistency is not an issue when developing a new tool, it does mean that existing commercial RT software cannot be made differentiable with AD. In contrast, DynRT allows differentiation to be implemented as a separate module, as demonstrated by 3DSCAT's DynRT plugin.

Finally, DiffRT lacks the interpretability provided by symbolic derivatives, which DynRT offers by using explicit symbolic differentiation.

IV. MULTIPATH LIFETIME MAP

Although it is relatively straightforward to assume that users, or any moving object, will experience the same multipath structure in their immediate vicinity, determining appropriate values for T_c and understanding how to compute it for a specific scene are less clear. Additionally, using a time-based reference instead of distance implies that the solution is tailored to a specific “displacement speed”. To address these uncertainties, we introduce a new simulation-based visual tool, named MLM, along with associated metrics to offer a deeper understanding of when DynRT is most effective.

In the following sections, we define the concepts of similar multipath structures, cells, and the relevant metrics for evaluating them. While we exemplify MLM on moving RXs, our method and the concept of lifetime **are not restricted to moving antennas**, and can be applied to **any number of dynamic objects**, like a moving reflector or an antenna performing beamforming, but specific implementation details will not be covered in this paper.

A. Multipath Structure

While the EM fields are continuous everywhere in 3D, this continuity does not always hold in RT simulations due to the nature of how ray paths are traced. Certain paths, such as the line of sight (LOS) path, only exist within specific regions, and unless some smoothing technique is applied [12], the collection of simulated paths tiles the space into continuous regions, separated by discontinuities.

Using the formalism defined in our previous work [5], we identify each ray path by a unique, ordered list of interacting objects, a list we call the “path candidate”. Each element of this list specifies, in order, the type of interaction and the object involved, such as a specular reflection off a particular wall. The length of the list corresponds to the number of interactions the path undergoes between TX and RX. We define a **multipath cell** as the set of continuous regions where the set of valid path candidates—the multipath structure—remains unchanged (see Fig. 2).

B. How to identify and color cells

The identification and coloring of multipath cells present certain implementation challenges, for which we provide the following solution.

Let the size of the scene, i.e., the number of objects (usually, triangular or rectangular facets), be N , and the number of admissible interactions be K . Then exhaustive RT, which is what our tool performs by default, will process up to $M = N(N - 1)^{K-1}$ path candidates. Hence, each multipath cell, C_i , can be identified by a vector, $v_i \in \{0, 1\}^M$, of boolean values, with each entry indicating whether a path candidate is valid in that cell.

To reduce the memory footprint of our tool, we process the possible path candidates in smaller chunks, compressing the partial vectors into unique scalar values that may vary between snapshots. To ensure consistent coloring of multipath cells—which is necessary only for the visual output, but not for the metrics—we compress each multipath cell vector into a fixed-size space with the 256-bit Secure Hash Algorithm (SHA256). If no hash collisions occur, as observed in our experiments, this enables us to map each cell to a constant unique color at a much lower memory cost.

C. Related Metrics

While a visual tool has definitively some advantages, especially for debugging purposes, it does not convey any concrete value that one can later use to make a decision.

For a given cell C_i , we introduce the following metrics:

- the area covered by each multipath cell, $S_i = \text{area}(C_i)$;
- and the average minimal inter-cell distance, $\bar{d}_i$;

where i indicates the index of the multipath cell, $\bar{\cdot}$ is the ensemble average over a given cell, and the minimal inter-cell distance of an object $x \in C_i$ is:

$$d_i(x) = \min_{y \notin C_i} \text{dist}(x, y), \quad (1)$$




Fig. 2. Example of multipath cells for individual—and combined—types of interaction, for a fixed TX position (circle node) and a moving RX. Each color indicates a unique multipath cell, i.e., regions reached by the same type of interaction. When all types are considered simultaneously, more cells are created, each having a different multipath structure. Blank regions are part of the no-multipath cell. Fig. 2d exhibits 7 unique cells, split across 11 continuous regions.




Fig. 3. Street canyon scene from Sionna [3].

that is, the minimum distance the object x has to travel to leave the cell C_i .

The area per multipath cell (in m^2) provides insight into the size of the cells. However, as the same cell can be split among multiple spatial regions, as seen in Fig. 2, the average minimal inter-cell distance is also important; it provides an estimate of the minimal lifetime of a multipath structure before an object transitions to another cell. Furthermore, a transition from one cell to another implies exactly one different path in the multipath structure, allowing for the reuse of all other paths, as they remain valid. Nevertheless, it may be computationally challenging to determine which path changes. In other words, the Hamming distance between two neighboring cell vectors, v_i and v_j , is always 1.

D. Urban Street Canyon

For this example, we used the “*simple street canyon*” scene from Sionna [3], performing the simulation with our own open-source DiffRT toolbox, named DiffRT¹. A comprehensive tutorial² is also available to guide readers through the simulation procedure. This scene is spread over a 120 m by 185 m rectangular area, with building heights ranging from 20 m to 50 m. Each building in the scene has a square footprint of 30 m by 30 m and consist of triangular facets; we first merged these facets into quadrilaterals to ensure logical coherence in ray paths and align with the principles of DynRT. This adjustment prevents path candidates from varying across individual triangular facets of the same building face. Furthermore, this 6-Building (6B) scenario includes narrow streets

TABLE I
MEAN AND MEDIAN VALUES OVER ALL 50 SIMULATION SNAPSHOTS AND CELLS.

	6B		2B	
	S (m^2)	$\bar{d}$ (m)	S (m^2)	$\bar{d}$ (m)
Mean	225.62	1.40	840.56	3.08
Median	86.43	1.00	371.38	2.30

(15 m wide) that intersect the main street canyon (20 m wide) at perpendicular angles. These narrow openings are frequently omitted, simplifying the model to represent the canyon as a long 2-Building (2B) scenario, where all building faces along each side of the street are treated as co-planar. Accordingly, we conducted an additional simulation with a simplified 2B scene, grouping buildings on each side of the main street canyon into single structures.

a) Simulation Parameters: MLM and metrics were derived from 50 RT snapshots, each with TX positioned at linearly spaced intervals between the scene’s endpoints (see Fig. 4). TX was placed at a constant altitude of 32 m, above all buildings, with a grid of 500-by-500 RXs uniformly distributed across the scene at an altitude of 1.5 m. Only LOS and first-order reflection paths were included in the simulation.

b) Results: Fig. 4 illustrates MLM at three different RX positions. Using data from the 50 snapshots, we compiled a histogram of S_i and $\bar{d}_i$ shown in Fig. 5. Tab. I summarizes the mean and median values of all measured S_i and $\bar{d}_i$.

c) Comments: Even in this simple scenario with limited paths, Fig. 4 already exhibits the complex spatial distribution of multipath cells. Increasing the reflection order would only introduce more cells and further reduce the mean values of S_i and $\bar{d}_i$. However, diffraction would not affect these cells, as a diffraction path of any order exists wherever a corresponding reflection-only path is present. Assuming disk-shaped cells, the 6B scenario yields a cell radius from 5.25 m to 8.47 m (median to mean), aligning with spatial consistency values of 5 m to 10 m in local areas [13]. However, cells are rarely disk-shaped, as shown by the relatively small $\bar{d}$ and observed in Fig. 4; they tend to extend along a primary direction with the longest lifetime. Ultimately, as demonstrated by the 2B scenario, simplifying the scene can extend the lifetime

¹Repository: <https://github.com/jeertmans/DiffRT>.

²Tutorial: <https://differt.rtd.io/eucap2025/notebooks/multipath.html>.




Fig. 4. MLMs for fixed TX positions and a moving RX. Each unique MC_i is assigned a unique color. E.g., the purple color indicates LOS-only visibility. The same color is reused in other snapshots to indicate the same multipath cell. The no-multipath cell is transparent, thus leaving the ground apparent.




Fig. 5. Histograms of (top) area occupied by each multipath cell and (bottom) average distance a RX has to travel to leave its current cell.

of a given multipath structure, further enhancing DynRT's suitability.

V. CONCLUSION

In this work, we have presented a comparative analysis of DynRT and DiffRT methods for RT, using 3DSCAT and Sionna as respective examples. While DynRT offers manual differentiation, providing interpretability through symbolic derivatives, it is constrained by the tedious derivation process and limited variable applicability. Conversely, DiffRT leverages AD to calculate derivatives for nearly any scene parameter with ease and flexibility, though often at the cost of restricted operational primitives and interpretability.

Despite the unique benefits of both approaches, the availability of tools implementing DynRT or DiffRT remains limited. The development of our open-source framework, DiffeRT, aims at offering a RT-focused alternative to Sionna. Ultimately, expanding the range of open-source DynRT and DiffRT tools will be crucial for advancing the field of RT.

Then, we have introduced a novel simulation-based method, via MLM and two metrics, to investigate the scope of appli-

cations of extrapolating RT snapshots, as used in DynRT. The method shows that scene simplifications can increase the scope where DynRT is best suited, but also highlight the relative complex shapes of multipath cells even in simple scenarios. In future research, we intend to expand the scope of the study to encompass larger scenes and additional categories of dynamic objects, including multiple moving reflectors.

REFERENCES

- [1] F. Fuschini, E. M. Vitucci, M. Barbiroli, G. Falciascecca, and V. Degli-Esposti, "Ray tracing propagation modeling for future small-cell and indoor applications: A review of current techniques," *Radio Science*, vol. 50, no. 6, pp. 469–485, 2015.
- [2] D. Bilibashi, E. M. Vitucci, and V. Degli-Esposti, "Dynamic ray tracing: Introduction and concept," in *2020 14th European Conference on Antennas and Propagation (EuCAP)*, 2020, pp. 1–5.
- [3] J. Hoydis, S. Cammerer, F. Ait Aoudia, A. Vem, N. Binder, G. Marcus, and A. Keller, "Sionna: An open-source library for next-generation physical layer research," *arXiv preprint*, Mar. 2022.
- [4] 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.
- [5] 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.
- [6] V. Červený and F. Hron, "The ray series method and dynamic ray tracing system for three-dimensional inhomogeneous media," *Bulletin of the Seismological Society of America*, vol. 70, no. 1, pp. 47–77, 02 1980.
- [7] F. Quatresouz, S. Demey, and C. Oestges, "Tracking of interaction points for improved dynamic ray tracing," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 7, pp. 6291–6301, 2021.
- [8] P.-S. Kildal, "Analysis of numerically specified multireflector antennas by kinematic and dynamic ray tracing," *IEEE Transactions on Antennas and Propagation*, vol. 38, no. 10, pp. 1600–1606, 1990.
- [9] E. Reinhard, B. Smits, and C. Hansen, "Dynamic acceleration structures for interactive ray tracing," in *Rendering Techniques 2000*, B. Péroche and H. Rushmeier, Eds. Vienna: Springer Vienna, 2000, pp. 299–306.
- [10] G. Kedem, "Automatic differentiation of computer programs," *ACM Trans. Math. Softw.*, vol. 6, no. 2, p. 150–165, Jun. 1980.
- [11] J. Hoydis, F. A. Aoudia, S. Cammerer, F. Euchner, M. Nimier-David, S. ten Brink, and A. Keller, "Learning radio environments by differentiable ray tracing," 2023.
- [12] J. Eertmans, L. Jacques, and C. Oestges, "Fully differentiable ray tracing via discontinuity smoothing for radio network optimization," in *2024 18th European Conference on Antennas and Propagation (EuCAP)*, 2024, pp. 1–5.
- [13] S. Ju and T. S. Rappaport, "Millimeter-wave extended nysim channel model for spatial consistency," in *2018 IEEE Global Communications Conference (GLOBECOM)*, 2018, pp. 1–6.