

# Enabling Site-Specific Cellular Network Simulation Through Ray-Tracing-Driven ns-3

Tanguy Ropitault<sup>\*†</sup>, Matteo Bordin<sup>§</sup>, Paolo Testolina<sup>§</sup>,

Michele Polese<sup>§</sup>, Pedram Johari<sup>§</sup>, Nada Golmie<sup>‡</sup>, Tommaso Melodia<sup>§</sup>

<sup>\*</sup>Associate, CTL, National Institute of Standards and Technology, Gaithersburg, MD, USA

<sup>†</sup>Prometheus Computing LLC, Bethesda, MD, USA

<sup>§</sup>Institute for the Wireless Internet of Things, Northeastern University, Boston, MA, USA

<sup>‡</sup>CTL, National Institute of Standards and Technology, Gaithersburg, MD, USA

{tanguy.ropitault,nada.golmie}@nist.gov; {bordin.m, p.testolina, m.polese, p.johari, t.melodia}@northeastern.edu

**Abstract**—Evaluating cellular systems, from 5th generation (5G) New Radio (NR) and 5G-Advanced to 6th generation (6G), is challenging because the performance emerges from the tight coupling of propagation, beam management, scheduling, and higher-layer interactions. System-level simulation is therefore indispensable, yet the vast majority of studies rely on the statistical 3rd Generation Partnership Project (3GPP) channel models. These are well suited to capture average behavior across many statistical realizations, but cannot reproduce site-specific phenomena such as corner diffraction, street-canyon blockage, or deterministic line-of-sight conditions and angle-of-departure/arrival relationships that drive directional links.

This paper extends 5G-LENA, an NR module for the system-level Network Simulator 3 (ns-3), with a trace-based channel model that processes the Multipath Components (MPCs) obtained from external ray-tracers (e.g., Sionna Ray Tracer (RT)) or measurement campaigns. Our module constructs frequency-domain channel matrices, and feeds them to the existing Physical (PHY)/Medium Access Control (MAC) stack without any further modifications. The result is a geometry-based channel model that remains fully compatible with the standard 3GPP implementation in 5G-LENA, while delivering site-specific geometric fidelity. This new module provides a key building block toward Digital Twin (DT) capabilities by offering realistic site-specific channel modeling, unlocking studies that require site awareness, including beam management, blockage mitigation, and environment-aware sensing. We demonstrate its capabilities for precise beam-steering validation and end-to-end metric analysis. In both cases, the trace-driven engine exposes performance inflections that the statistical model does not exhibit, confirming its value for high-fidelity system-level cellular networks research and as a step toward DT applications.

**Index Terms**—Ray tracing, cellular networks, digital twin

## I. INTRODUCTION

System-level simulation is an indispensable tool for the design, evaluation, and optimization of 5th generation (5G) New Radio (NR) networks. Among the available simulators, the Network Simulator 3 (ns-3) 5G-LENA module [1] has established itself as a widely adopted open-source reference platform due to its compliance with 3rd Generation Partnership

This work was partially supported by the National Telecommunications and Information Administration (NTIA)’s Public Wireless Supply Chain Innovation Fund (PWSCIF) under Award No. 25-60-IF011, and by the National Science Foundation (NSF) under Award No. CNS 2112471.

The mention of commercial products, their sources, or their use in connection with material reported herein is not to be construed as either an actual or implied endorsement of such products by the Department of Commerce.

Project (3GPP) procedures and its support for configurable numerologies, bandwidth parts, and beamforming strategies. In its default configuration, however, 5G-LENA relies on the statistical channel model of 3GPP Technical Report (TR) 38.901 [2], [3] (in the remainder of the paper, we refer to 3GPP TR 38.901 as TR 38.901 for brevity), in which large-scale parameters (e.g., path-loss, shadowing) and small-scale clusters (powers, delays, angles) are generated according to scenario-dependent probability distributions. While this approach is well-suited for extensive Monte-Carlo studies and standardization activities, it cannot capture site-specific propagation features, e.g., deterministic blockage or dominant paths that drive the performance of highly directional links. In particular, statistical clusters blur the geometric relationship between the transmitter and the receiver, masking the effects of Line-of-Sight (LoS) occlusion and limiting the realism of beam-management or sensing algorithms that depend on accurate Angle of Departure (AoD) and Angle of Arrival (AoA) information.

To close this gap, we introduce a 5G-LENA extension that processes Multipath Component (MPC) traces, allowing ns-3 to compute a site-specific, physically realistic channel matrix through its native pipeline. The proposed *trace-based channel model* can operate with traces from ray-tracers (e.g., Sionna Ray Tracer (RT) [4], Remcom Wireless Insite, etc.) or measurement campaigns. The new module replaces the stochastic cluster synthesis of TR 38.901 with the delays, powers, phases, and angular information contained in the traces, while preserving 5G-LENA’s existing beamforming and scheduling block. Consequently, users can perform site-specific evaluations of link adaptation, beam management and design procedures under the same Medium Access Control (MAC) and Physical (PHY) implementations already available in the simulator.

The main contributions of this work are threefold. First, we design and release<sup>1</sup> SiOLENA, an ns-3 module that builds frequency-domain channel matrices from ray traces, modeling location-specific multipath and Doppler effects. Second, we maintain full compatibility with 5G-LENA release 4 without modifying its core scheduler, physical layer abstraction, or

<sup>1</sup><https://github.com/usnistgov/siolena>

antenna models, thereby enabling a high-fidelity alternative to the native TR 38.901 channel model. Third, we validate the module in two realistic three dimensional (3D) outdoor scenarios, showing that the proposed channel model reproduces the expected angular behavior of LoS links and captures the pronounced end-to-end metrics degradation that occurs in Non-Line-of-Sight (NLoS), effects that the statistical model largely smooths out.

Our ns-3 extension, by integrating detailed propagation traces with a 3GPP-compliant network stack, enables reproducible research in advanced beamforming, positioning, and sensing in complex environments, and supports Artificial Intelligence (AI)/Machine Learning (ML) data collection for wireless networks. Because realistic propagation modeling is crucial for Digital Twin (DT) use cases, integrating it with ns-3's system-level simulator, which offers 5G core capabilities, marks a substantial stride toward practical DT realization.

The remainder of the paper is organized as follows. Section II reviews the state of the art and related works, highlighting our contributions. In Section III, we offer a brief overview of the 5G-LENA NR module and describe our trace-based channel model extension. Section IV presents the beamforming validation results in the Place de l'Étoile scenario, analyzes end-to-end metrics in a Boston street canyon scenario, and discusses the implications of trace-based versus statistical modeling. Finally, Section V concludes the paper and outlines directions for future work.

## II. RELATED WORKS

Stochastic channel models [2], [3], [5] are extremely effective in capturing general channel properties, but fail to reliably capture the propagation characteristics of specific environments. In this context, RTs enable physically consistent, site-specific full-stack simulations [6], at the cost of a higher computational complexity. This shift toward realistic modeling aligns with the DT paradigm, where virtual representations aim to accurately mirror the physical network's behavior and dynamics. Among early efforts, [7] integrated a Quasi-Deterministic (Q-D) channel model [8] into the IEEE 802.11ad/ay ns-3 implementation, providing for the first time trace-based channel generation capability to ns-3. The Q-D methodology, specifically designed to accurately model IEEE 802.11ad/ay channels, generates MPCs based on the propagation conditions in a specific 3D environment, combining a deterministic ray-tracing component (large-scale fading) with a stochastic process (small-scale fading) whose distribution was extracted from an extensive measurement campaign [8]. The authors of [6] added similar capabilities to the 5G millimeter wave (mmWave) module.

More recently, RTs that exploit the parallelization capabilities of Graphics Processing Units (GPUs), e.g., Sionna RT [4], dramatically reduced the channel computation time. VaN3Twin framework [9] embeds Sionna RT into the ms-van3t vehicular simulator via a modular UDP client-server architecture. Unlike prior ns-3 integrations, VaN3Twin supports multi-RAT simulations across Institute of Electrical and Electronics

Engineers (IEEE) 802.11p, Long Term Evolution (LTE)-V2X, and NR-V2X, while enabling high-fidelity modeling of static and dynamic objects through 3D meshes. Ns3Sionna [10] introduces a client-server architecture based on ZeroMQ and protocol buffers to interface ns-3 with Sionna RT.

The integration with RT allows reproducing a space and time-consistent, site-specific communication channel. Namely, (i) LoS/NLoS conditions are deterministically set by the node positions in the 3D environment; (ii) the overall channel gain is computed based on the interactions of the signal with the 3D environment model across its path(s) from the transmitter to the receiver; and (iii) correspondingly, the AoA and AoD of the MPCs are consistently determined by the aforementioned interactions. We highlight that *the angular information plays a key role particularly in directional communication scenarios*, i.e., with directional antennas or multi-element antenna arrays, as it is needed to determine the MPC attenuation due to the antenna pattern and the beamforming vectors. However, [9], [10] reduce the channel response to an overall, pre-combined scalar gain, delay, and channel condition (LoS/NLoS), discarding the AoA and AoD information. Customized antenna patterns and beamforming strategies can still be used, but they have to be implemented in Sionna, which disconnects their implementation from the ns-3 antenna and beam management stack. In practice, the standard 5G NR beam management implementation, e.g., in the LENA module, is not compatible with these frameworks, and has to be re-implemented to interface with Sionna directly; or the simulation of directional antennas can not be controlled from ns-3, but has to be disabled in ns-3 and externally managed in Sionna. We believe that these elements limit the applicability of the aforementioned approaches, particularly for research in Multiple Input, Multiple Output (MIMO) and Frequency Range 2 (FR2), Frequency Range 3 (FR3), and sub-THz frequency bands, where directionality is necessary to overcome the path loss.

In contrast, our architecture generates the **full spatial channel matrix** between each antenna element at runtime directly in ns-3, supporting configurable Phased-array Antenna (PAA), based on the complete MPC information obtained from a RT or a channel sounding measurement campaign, and feeding it into the ns-3 pipeline. This enables directional transmission to be modeled and beamforming vectors to be generated at runtime by ns-3, taking full advantage of its capabilities for antenna and beam management, unlocking its potential for beamforming and MIMO research.

## III. RAYTRACING-BASED CHANNEL MODEL

### A. ns-3 LENA NR Module and 3GPP Stochastic Model

5G-LENA, the ns-3 NR module [1], supports 3GPP Release-15 end-to-end system-level simulations with support for flexible frame structures via multiple numerologies, Bandwidth Parts (BWPs), and both time- and frequency-domain scheduling with dynamic Time Division Duplexing (TDD), Orthogonal Frequency-Division Multiple Access (OFDMA), Hybrid Automatic Repeat reQuest (HARQ), and Modulation and Coding Scheme (MCS) adaptation. The PHY is abstracted through

Exponential Effective Signal to Interference plus Noise Ratio (SINR) Mapping (EESM) with lookup tables derived from a NR-compliant link-level simulator. Additionally, the NR module offers both ideal and realistic beamforming functionalities. Ideal methods rely on perfect channel or positional knowledge (e.g., cell-scan or direction-of-arrival), while realistic methods rely on Sounding Reference Signal (SRS) SINR/Signal-to-Noise-Ratio (SNR) measurements to determine BF vectors. The 5G-LENA module in ns-3 provides a TR 38.901 compliant channel modeling pipeline [3] that captures large- and small-scale propagation effects for 3GPP Urban Micro (UMi), Urban Macro (UMa), and Rural Macro (RMa) scenarios. It is based on four classes as depicted in Fig. 1, top:

- `ChannelConditionModel` determines the probabilistic channel condition, LoS, NLoS, or Outdoor to Indoor (O2I), based on the statistical models in TR 38.901. The scenario parameters directly influence the LoS probability, while the determined condition subsequently drives all path-loss and fading computations.
- `ThreeGppPropagationLossModel` implements large-scale fading mechanisms including path loss and shadow fading. Based on scenario configuration and LoS condition, it selects and applies the appropriate scenario-specific path-loss formula from Section 7 of TR 38.901, incorporating frequency-dependent corrections, antenna height effects, and blockage losses to generate realistic link budgets.
- `ThreeGppChannelModel` generates the complex channel matrix using the cluster-based MPC modeling according to [2]. Based on the scenario and channel condition, it determines cluster characteristics (number, powers, delays, angular spreads), generates MPC parameters, and constructs the spatial channel matrix accounting for antenna array geometry and spatial correlation.
- `ThreeGppSpectrumPropagationLossModel` applies frequency-selective small-scale fading and beamforming processing to the transmitted signal spectrum. This class computes and applies the Doppler frequency shifts and the beamforming gain, and integrates the complete channel response with the signal spectrum to produce realistic instantaneous SINR variations for link adaptation and scheduling algorithms.

### B. SioLENA Trace-based Channel Model

To complement the statistical 3GPP approach, we extended the LENA pipeline to utilize trace-based channel modeling. This approach enables to use MPCs obtained from ray-tracers or measurement campaigns to enhance the channel model realism. Two specialized classes (Fig. 1, bottom) operate in place of `ThreeGppChannelModel` and `ThreeGppSpectrumPropagationLossModel`:

- `TracesChannelModel` processes trace-based pre-computed MPC from time-indexed Comma-Separated Values (CSV) files generated by RT software such as Sionna, or from measurement campaigns. Each trace entry specifies the propagation delay, the amplitude, phase,

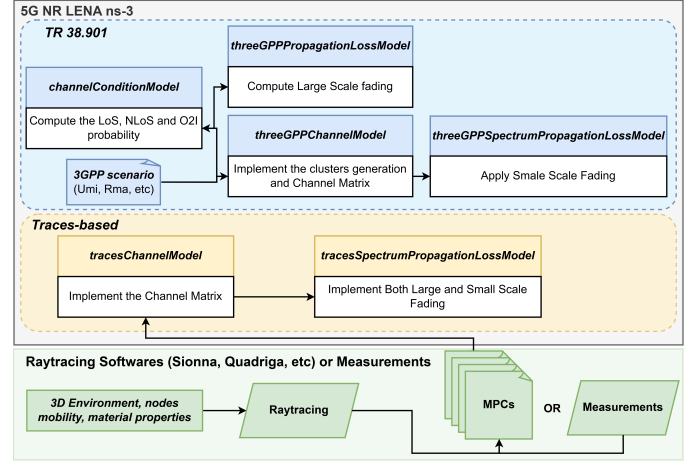


Fig. 1: Overall channel model pipeline in 5G-LENA ns-3: (top) TR 38.901-based classes, (bottom) newly added trace-based classes.

and both departure and arrival angles of an MPC, for every transmitter–receiver link pair. For each MPC, the antenna array response vectors are computed based on the provided angles, with appropriate phase shifts applied according to element positions and carrier wavelength. This process yields a frequency-domain channel matrix that faithfully represents the ray-traced or measured propagation environment.

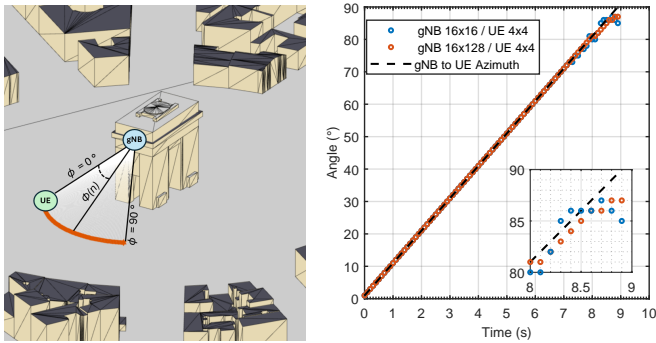
- `TracesSpectrumPropagationLossModel` applies these deterministic channel matrices to the signal spectrum. Unlike the statistical approach, the Doppler shifts are computed directly from node velocities projected onto the MPC propagation directions, providing physically accurate frequency shifts. Beamforming gains are calculated by multiplying the channel matrix by the transmit and receive steering vectors.

Since the RT-derived or measurement campaigns MPCs inherently encode the actual propagation conditions including LoS/NLoS states, path delays, and received power levels, the `TracesChannelModel` eliminates the dependency on probabilistic modeling, capturing site-specific correlation and multipath structures.

### C. Sionna RT Channel-Trace Generation for ns-3

While our trace-based channel model can use data from any RT or measurement campaign that follows the required CSV format, in this article we showcase the workflow using *Sionna RT* [4]. Sionna is an open-source Python library (TensorFlow-based) for link-level wireless simulation. Sionna RT extends it with GPU-accelerated 3D ray tracing that models reflection, diffraction and scattering under both LOS and NLOS conditions. To feed ns-3, we built a lightweight pipeline<sup>2</sup> that (i) defines device trajectories (static, linear, circular, or user-provided), (ii) calls Sionna RT to compute the MPCs (delays, complex gains, AoA, AoD, Doppler) for each Tx–Rx link, and (iii) exports these MPCs in the exact CSV schema expected by `TracesChannelModel`

<sup>2</sup><https://github.com/wineslab/sionna-channel-generator>



(a) Visualization of the considered (b) Time evolution of the steering azimuth  $\acute{E}$ toile scenario. The angle  $\phi$  denotes the selected for two gNBs PAAs. An inset around the azimuth AOD, with  $\phi(\tau)$  representing 8-9 s highlights the increased deviation near the AOD at time step  $\tau$ . The UE trajectory is illustrated using orange markers.

Fig. 2: Place de l'Étoile scenario.

and `TracesSpectrumPropagationLossModel`. This exporter bridges Sionna RT's detailed propagation output with our `SioLENA ns-3` extension, enabling site-specific channel simulation without further manual conversion.

#### IV. EVALUATION

The goal of this section is twofold. First, we verify that the trace-based channel faithfully preserves the geometric information needed for beamforming, i.e., that the steering angles selected by 5G-LENA closely match the true LoS directions embedded in the RT trace. Second, we show that this geometric fidelity propagates to end-to-end metrics: SINR, MCS, and user throughput. In both case studies, we use the lightweight pipeline from Sec. III-C to bridge Sionna RT and ns-3. Sionna RT generates the MPCs, configured to include diffraction and up to four orders of specular reflection, from the 3D environment and node trajectories. Our pipeline then exports these MPCs into the CSV schema expected by the `SioLENA` extension, enabling seamless import into ns-3.

The two scenarios intentionally involve a single Next Generation Node Base (gNB)–User Equipment (UE) pair and operate in the mmWave band FR2 so that the results remain easily interpretable. The proposed module, however, is agnostic to the number of nodes and to the carrier frequency: it can import traces for multiple gNBs and UEs simultaneously and works unchanged in FR1 or the emerging FR3 spectrum.

##### A. Place de l'Étoile: Beamforming validation

We evaluate our trace-based channel model extension in a LoS urban scenario at Place de l'Étoile, Paris. As depicted in Fig. 2a, a single gNB, operating at 28 GHz with 100 MHz of bandwidth, is mounted atop the Arc de Triomphe (10 m height), while a UE at 1.5 m height moves along a circular trajectory, advancing by  $1^\circ$  every 100 ms (i.e., the gNB's azimuth AoD sweeps from  $0^\circ$  to  $90^\circ$  over 9 s).

Beamforming training is performed every 100 ms using the `idealBeamforming` procedure from LENA 5G-NR, configured to scan the azimuth plane in  $1^\circ$  steps and the elevation plane in  $10^\circ$  steps. At each beam-training interval, the gNB selects the steering vector that yields the highest received

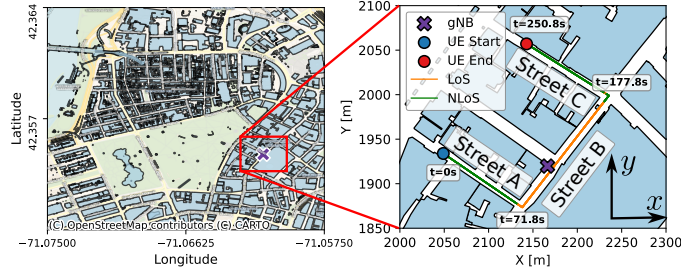


Fig. 3: Boston simulation scenario.

power at the UE. We emphasize that `idealBeamforming` is idealized: beam-training frames are not transmitted over the medium, so no interference arises during beam training. Two gNB array configurations are studied:  $16 \times 16$  PAA and  $16 \times 128$  PAA while the UE uses a  $4 \times 4$  PAA. The PAA elements are positioned in the  $y$ - $z$  plane and the antenna elements are isotropic.

Fig. 2b compares the azimuth of the beamforming vector selected in ns-3 (markers) with the true UE azimuth (dashed line). The two patterns are indistinguishable for most of the 9 s trajectory, confirming the correctness of the model. The matching persists until the LoS angle approaches about  $70^\circ$ , beyond which the steering angle approaches the end-fire region, as the PAA elements lie in the  $y$ - $z$  plane, and thus exhibit limited azimuthal resolution. The inset highlights the increased beam spread between 8 s and 9 s, with the three largest errors occurring within this period:  $5^\circ$  for the  $16 \times 16/4 \times 4$  array and  $3^\circ$  for the  $16 \times 128/4 \times 4$  array.

The beam-steering for the  $16 \times 16/4 \times 4$  configuration has a mean error of  $-0.17^\circ$  and a Root Mean Squared Error (RMSE) of  $0.74^\circ$ , while the  $16 \times 128/4 \times 4$  configuration achieves a mean error of  $-0.14^\circ$  and an RMSE of  $0.51^\circ$ . The lower RMSE of the latter array reflects its larger element count: more elements yield a narrower beamwidth and finer angular discrimination in azimuth, reducing quantization error and the magnitude of end-fire outliers.

These results confirm that (i) the trace-based channel model preserves directional information from the imported MPCs, and (ii) `idealBeamforming` coupled with the trace-based channel model recovers the true LoS direction in most cases. The expected degradation near end-fire, and the reduced error when moving from  $16 \times 16$  to  $16 \times 128$ , demonstrate that our framework can be used not only to study PAA design trade-offs, to evaluate beamforming and beam-training algorithms in realistic scenarios, but also to generate realistic, labeled datasets for ML-based beam-steering or channel estimation algorithms. Such ML-driven methods require training data with geometrically consistent steering directions, which cannot be obtained using the 3GPP statistical channel model.

##### B. Boston Street Canyon: End-to-End Analysis

This validation quantifies the end-to-end impact of our trace-based channel model. Leveraging the Boston digital-twin street-canyon model of [11], we deploy a single gNB at 10 m height and let a pedestrian UE at 1.5 m height walk 375 m in 250 s at 1.5 m/s. The route crosses three consecutive streets,


A, B, and C, and alternates between LoS and NLoS visibility: Streets A and C are NLoS, whereas Street B is predominantly LoS with a brief blockage near its end (see Fig. 3).

We compare the results obtained using our trace-based channel model, based on the MPCs generated in Sionna, against those obtained with the UMi street-canyon 3GPP channel model. The gNB employs a  $16 \times 16$  PAA, whereas the UE is equipped with a  $4 \times 4$  PAA. Downlink traffic is generated at a constant User Datagram Protocol (UDP) rate of 10 000 packets/s with 1500 byte packets, corresponding to 122 Mb/s; the gNB buffer stores up to 10 000 packets. Beam training is performed every 100 ms using a  $10^\circ$  scan grid, and the 3GPP LoS/NLoS condition is re-evaluated at the same 100 ms interval.

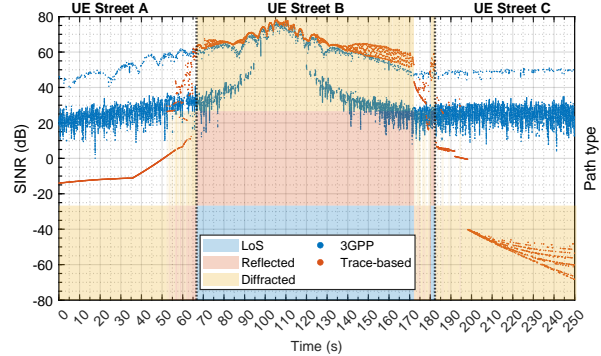
**Channel Characterization.** Fig. 4a shows the LoS agreement between 3GPP and trace-based channel model, and Fig. 4b shows both the MPCs (LoS, reflected, diffracted) in the Sionna trace and the SINR evolution for 3GPP and trace-based channel model. We can identify three distinct regions: (1) *Early NLoS* (0 s to 67 s): The UE is in Street A (physically NLoS). The trace-based model correctly remains NLoS, whereas TR 38.901 erroneously reports LoS (red bars). For the trace-based channel, we can observe that diffraction dominates, with specular reflections appearing only as the UE rounds the corner into Street B (around 53 s). (2) *Predominantly LoS* (67 s to 182 s): In Street B the link is LoS except during a blockage from 172 s to 180 s (see Fig. 3). Both models agree most of the time (green bars), though TR 38.901 misclassifies  $\sim 32\%$  of samples as NLoS (orange bars). The trace-based channel is driven by a strong LoS MPC (except during the blockage), with reflected and diffracted components. (3) *Late NLoS* (182 s to 250 s): Upon entering Street C, the physical link returns to NLoS. The trace-based model remains strictly NLoS, while TR 38.901 occasionally indicates LoS. Again, for trace-based channel model, brief specular reflections punctuate an otherwise diffraction-dominated channel.

**SINR.** Fig. 4b shows the instantaneous downlink SINR at the UE for both channel models. This SINR is governed by a tight interplay between the time-varying channel and the beamforming. As the UE moves, some MPCs appear while others fade; their relative phases determine whether they add constructively or destructively in the PAA plane. Periodic beam training then computes the weights that maximize the combined power at the array ports. Thus, the received power, and therefore the SINR, is the joint outcome of the instantaneous MPC set and the current beam-steering weights. Two global trends emerge. (i) A switch from NLoS to LoS visibility yields an immediate  $\sim 20$  dB jump in SINR. (ii) For both models the SINR rises from 0 s to 106 s, while the UE and gNB approach their closest separation, and then falls as they move apart.

A closer inspection of the trace-based MPCs in Fig. 4b shows that the channel is dominated by diffracted paths in the intervals 0 s to 67 s and again after 182 s. In the former interval, a single diffracted ray keeps the SINR stable around




(a) Comparison of LoS status between the trace-based channel model and TR 38.901 (3GPP).




(b) Instantaneous downlink SINR for the two channel models (red markers trace-based and blue markers 3GPP). Coloured backgrounds mark the type of MPCs generated by Sionna (diffracted, reflected, or LoS).


Fig. 4: LoS agreement and SINR evolution versus time.



(a) MCS at the gNB.



(b) UDP throughput measured at the UE.



(c) Packet delay experienced at the UE.

Fig. 5: Link adaptation, throughput, and delay under the two channel models.

$-15$  dB, increasing slowly. After 35 s, as the gNB nears the corner of Street A, the diffraction angle becomes more favorable: the AoA of the diffraction path aligns more closely with the boresight of the UE's PAA, resulting in a sharper SINR increase. On the contrary, after 197 s, the SINR sharply declines as the diffraction path's AoA steers the beam into the array's end-fire region, reducing significantly the azimuthal gain. Between 53 s and 67 s, and 182 s and 195 s, intermittent specular reflections raise the SINR by  $\sim 20$  dB. Once the UE enters Street B at 67 s and the LoS is established, the SINR climbs above 60 dB, peaks near 80 dB at minimum gNB-UE separation, then declines as the nodes move apart. A brief blockage from 172 s to 180 s removes  $\sim 20$  dB, with a short rebound from 180 s to 182 s before the UE turns into Street C.

The TR 38.901 statistical model smooths these fluctuations: the SINR stays mostly in the 10 dB to 45 dB range even during

the initial NLoS interval, with frequent LoS appearances, and rarely falls below 10 dB, due to distance-based cluster power distributions that do not reflect the site geometry.

**End-to-end Metrics.** Fig. 5a shows how Adaptive Modulation and Coding (AMC) adapts to the time-varying SINR of Fig. 4b. In 5G-LENA, AMC selects each transport-block's MCS to maximize spectral efficiency under a 10% BLER constraint. During early NLoS when the UE is in Street A (0 s to 45 s), AMC picks MCS 0 (SINR  $\sim -15$  dB), whereas TR 38.901 reports (SINR  $\sim 20$  dB) and thus MCS 8 to 28 but mostly larger than 24. As the UE nears the Street A–B corner (45 s to 50 s), the trace-based SINR rises and MCS climbs to 8 to 28 (depending on reflections), while the 3GPP model remains at high indices. When the UE enters Street B (67 s to 182 s), LoS emerges and both models reach MCS 28, though 3GPP intermittently dips to 16 due to the NLoS condition; trace-based stays at 28 except during blockage (down to 22). After the UE entering Street C (182 s to 250 s), trace-based MCS gradually falls to 0 as reflections vanish, whereas the statistical model selects MCS 12 to 28.

The UE's throughput (Fig. 5b) and delay (Fig. 5c) mirror the MCS: lower indices shrink transport-block size and spectral efficiency, requiring more radio link control segments, extra HARQ rounds, and fewer bits per transmission time interval, thereby reducing throughput and raising delay. In the trace-based model, while the UE traverses Street A under NLoS (0 s to 45 s), MCS remains at 0, capping throughput at  $\sim 10$  Mb/s and inflating delay to  $\sim 8$  ms. Between 45 s and 67 s, as the UE rounds the Street A–B corner, emerging specular reflections and a diffraction angle that aligns more closely with the UE's PAA boresight, drive MCS up to 28, elevating throughput to 122 Mb/s and cutting delay below 0.5 ms. While traversing Street B under sustained LoS (67 s to 172 s), throughput remains at 122 Mb/s with minimal delay; when the obstacle blocks LoS (172 s to 180 s), MCS dips to 22 yet maintains  $\sim 100$  Mb/s and low delay. In Street C, from 182 s to 197 s, throughput drops to  $\sim 50$  Mb/s and delay peaks at  $\sim 2$  ms, except when specular reflection punctually appears (186 s to 195 s), restoring full rate and low delay. Finally, at 197 s, the UE is steering into its PAA end-fire region, forcing MCS back to 0, reducing throughput to  $\sim 10$  Mb/s and returning delay to  $\sim 8$  ms. These results clearly capture the site-specific performance variations induced by channel evolution.

By contrast, the TR 38.901 statistical street-canyon model smooths localized SINR fluctuations, sustaining throughput near 122 Mb/s (mean 121.92 Mb/s) and limiting delay to 0.5–1 ms (mean 0.54 ms), and thus does not reflect these site-specific effects.

**Discussion.** Deterministic MPC traces are indispensable whenever the *geometric* components of the channel, and not merely its average gain, drive system behavior (e.g. joint communication-sensing, and blockage prediction). Moreover, statistical and trace-based models are complementary: the TR 38.901 offers rapid Monte-Carlo sweeps for large network layouts, whereas the trace pinpoints worst-case corners, indoor–outdoor transitions, or moving blockers. Additionally,

the runtime exposure of full channel matrices supports hybrid or learning-based beam selection *without* re-running the ray-tracer. Finally, as the trace parser is agnostic to the ray-tracing tool and to the carrier frequency, the approach extends to sub-THz industrial sensing and satellite backhaul where LoS and specular reflections dominate.

## V. CONCLUSIONS

This work integrates site-specific ray-tracing into the open-source 5G-LENA simulator while preserving its 3GPP-compliant beamforming, scheduling, and link-adaptation stack. The trace-based channel model reproduces deterministic angular dynamics and, when compared with the standard TR 38.901 implementation, reveals significant differences in beam-steering accuracy, SINR, MCS selection, and user throughput under blockage and corner diffraction.

Ongoing work targets more challenging scenarios: (i) multi-site deployments with inter-cell interference, (ii) spectrum sharing and coexistence, (iii) mobility with multiple gNBs and UEs executing concurrent beam training and handover, and (iv) dynamic blocker models and traffic heterogeneity. Future directions also include exploring how this realistic propagation modeling could contribute to DT applications for network optimization and AI-driven network management.

## REFERENCES

- [1] N. Patriciello, S. Lagen, B. Bojovic, and L. Giupponi, "An E2E simulator for 5G NR networks," p. 101933, 2019.
- [2] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," Technical Report (TR) 38.901, 2024.
- [3] T. Zugno, M. Polese, N. Patriciello, B. Bojović, S. Lagen, and M. Zorzi, "Implementation of a Spatial Channel Model for ns-3," in *Proceedings of the 2020 Workshop on Ns-3*, ser. WNS3 '20, 2020, p. 49–56.
- [4] J. Hoydis, F. Ait Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna RT: Differentiable ray tracing for radio propagation modeling," in *2023 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2023, pp. 317–321.
- [5] T. S. Rappaport, S. Sun, and M. Shafi, "Investigation and Comparison of 3GPP and NYUSIM Channel Models for 5G Wireless Communications," in *IEEE 86th Vehicular Technology Conference (VTC-Fall)*, 2017.
- [6] M. Lecci, P. Testolina, M. Polese, M. Giordani, and M. Zorzi, "Accuracy Versus Complexity for mmWave Ray-Tracing: A Full Stack Perspective," *IEEE Transactions on Wireless Communications*, vol. 20, no. 12, pp. 7826–7841, 2021.
- [7] H. Assasa, J. Widmer, T. Ropitault, A. Bodi, and N. Golmie, "High Fidelity Simulation of IEEE 802.11ad in ns-3 Using a Quasi-deterministic Channel Model," in *Proceedings of the 2019 Workshop on Next-Generation Wireless with Ns-3*, ser. WNGW 2019, 2019, p. 22–25.
- [8] C. Gentile, P. B. Papazian, R. Sun, J. Senic, and J. Wang, "Quasi-deterministic channel model parameters for a data center at 60 GHz," *IEEE Antennas and Wireless Propagation Letters*, vol. 17, no. 5, pp. 808–812, 2018.
- [9] R. Pegurri, D. Gasco, F. Linsalata, M. Rapelli, E. Moro, F. Raviglione, and C. Casetti, "VaN3Twin: the Multi-Technology V2X Digital Twin with Ray-Tracing in the Loop," 2025. [Online]. Available: <https://arxiv.org/abs/2505.14184>
- [10] A. Zubow, Y. Pilz, S. Rösler, and F. Dressler, "Ns3 meets Sionna: Using Realistic Channels in Network Simulation," 2024. [Online]. Available: <https://arxiv.org/abs/2412.20524>
- [11] P. Testolina, M. Polese, P. Johari, and T. Melodia, "Boston Twin: the Boston Digital Twin for Ray-Tracing in 6G Networks," in *Proceedings of the ACM Multimedia Systems Conference 2024*, ser. MMSys '24. ACM, Apr. 2024, p. 441–447.