Simplified Ray Tracing for the Millimeter Wave Channel: A Performance Evaluation

Mattia Lecci*, Paolo Testolina*, Marco Giordani*, Michele Polese*,
Tanguy Ropitault†, Camillo Gentile†, Neeraj Varshney†, Michele Zorzi*

*Department of Information Engineering, University of Padova, Italy, email:{name.surname}@dei.unipd.it
†National Institute of Standards and Technology (NIST), Gaithersburg, MD, 20899 USA, email:{name.surname}@nist.gov

***Abstract*—Millimeter-wave (mmWave) communication is one of the cornerstone innovations of fifth-generation (5G) wireless networks, thanks to the massive bandwidth available in these frequency bands. To correctly assess the performance of such systems, however, it is fundamental to have reliable channel models, based on a deep understanding of the propagation characteristics of the mmWave signal. In this respect, ray tracers can provide high accuracy, at the expense of a significant computational complexity, which limits the scalability of simulations. To address this issue, in this paper we present possible simplifications that can reduce the complexity of ray tracing in the mmWave environment, without significantly affecting the accuracy of the model. We evaluate the effect of such simplifications on link-level metrics, testing different configuration parameters and propagation scenarios.**

I. Introduction

The next generation of Cellular and Wireless Local Area Network (WLAN) will be the first to exploit millimeter wave (mmWave) frequencies to provide connectivity in the access network, i.e., in the links between base stations and mobile users. In particular, 3rd Generation Partnership Project (3GPP) NR has been designed to support a carrier frequency up to 52.6 GHz in Release 15 [1], and future Releases will consider an extension to 71 GHz and the sub-THz band [2]. Similarly, IEEE 802.11ad and 802.11ay exploit the unlicensed bands at 60 GHz [3]. The mmWave frequencies, indeed, feature large chunks of untapped bandwidth that can increase the data rate provided to the end users, making it possible to target the 5th generation (5G) requirements of ultra-high peak throughput (20 Gbps) and average user experienced rate (50-100 Mbps) [4]. Moreover, the small wavelength enables the design of antenna arrays with tens of elements in a small form factor, which could fit even smartphones or VR headsets.

The propagation characteristics of the Radio Frequency (RF) signals in these frequency bands, however, complicate the design of reliable communication systems [5]. First, the high propagation loss, which is proportional to the square of the carrier frequency, limits the coverage region of the mmWave base stations. This can be compensated by using beamforming with large antenna arrays, which could concentrate the power in narrow, directional beams and increase the link budget. Additionally, mmWave signals can be easily blocked by obstacles (e.g., vehicles, buildings, human bodies), preventing direct Line-of-Sight (LoS) communications. Furthermore, at mmWave frequencies, the increased diffraction results in deep shadow regions, thus further degrading propagation performance [6]. By considering the combination of these phenomena, the mmWave channel appears extremely volatile to mobile users, whose quality of experience may be poor unless a proper network design is adopted.

This work was partially supported by NIST under Award No. 70NANB18H273. Mattia Lecci and Paolo Testolina's activities were supported by *Fondazione CaRiPaRo* under the grants "Dottorati di Ricerca" 2018 and 2019, respectively.

A. Channel Modeling for mmWaves

As experimental platforms and testbeds at mmWaves are still at an early development stage [7], [8], analysis and simulation play a fundamental role for the performance evaluation of novel solutions for mmWave networks. Given the aforementioned behavior of the channel at such high frequencies, and the interplay with network deployment choices and beamforming design, the accuracy of analysis and simulation depends on that of the channel model even more than at conventional sub-6 GHz frequencies. Therefore, the research community has developed a number of channel modeling tools for mmWaves, with a varying degree of complexity and accuracy. Stochastic and analytical models are based on the combination of random variables fitted on traces and measurements, and are widely used for analysis [9], [10] and system-level simulations [11], [12]. However, the generality of these models and their stochastic nature fits poorly with the need to accurately characterize specific scenarios. Additionally, most stochastic models may not properly characterize the specific features of the mmWave channel that may affect the overall system performance, such as the temporally- and spatially-consistent updates of the LoS condition and the evolution of each single Multi Path Component (MPC).

These modeling challenges are instead addressed by Ray Tracers (RTs), which have been used to precisely characterize the propagation of RF signals in specific scenarios [13], [14], [15]. With ray tracing, the channel is modeled in terms of MPCs that generate from a certain location and angle of departure, are reflected (and, in complete models, diffused) on the scattering surfaces of the scenarios, and reach the position of the receiver with a certain angle of arrival, delay and power [16]. As the generation of MPCs is purely based on the geometry of the scenario, the channel is as accurate as the description of the scenario, and the MPCs are consistent with the mobility model of the communication endpoints. Additionally, ray tracers can be easily integrated into system-level simulators, by computing the channel matrix $\mathbf{H}$ that combines the different multipath components and the antenna arrays of the network nodes.

With currently available channel modeling tools, however, a higher accuracy translates into a higher computational complexity for the MPC generation and the simulations. As we discuss in [17], the complexity is proportional to two elements, i.e., the number of MPCs which need to be combined to generate the channel matrix $\mathbf{H}$, and the number of antennas at the two endpoints of the communication link (which represents the number of columns and rows of $\mathbf{H}$).

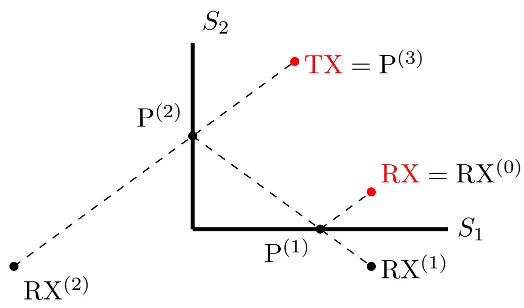

Fig. 1: Visualization for the Method of Images algorithm for a second-order reflection ($N = 2$).

B. Contributions

Based on the above introduction, in this paper we investigate whether it is possible to improve the trade off between accuracy and complexity in mmWave simulations, by studying simplification techniques for ray tracers that speed up the simulations and the ray tracer itself. Specifically, we consider processing only MPCs whose received power is above a certain threshold, which is relative to the strongest MPC, and limit the maximum number of reflections for each MPC. Our results show that it is possible to decrease the complexity of the simulations with a minimal reduction in accuracy, with respect to the baseline ray tracer implementation (i.e., without simplifications).

These promising results are a first step towards understanding and isolating which are the most fundamental channel modeling components at mmWave frequencies, and could stimulate further investigations into whether it is possible to develop simplified channel models (e.g., to be used also for mathematical analysis) that are more representative of the mmWave propagation than the widely-used Nakagami-m or Rayleigh fading models [18].

The rest of the paper is organized as follows. In Section II we provide details on the ray tracer we consider as baseline. We then introduce possible simplifications in Section III, while performance results are discussed in Section IV. Finally, we conclude the paper in Section V.

II. The Millimeter Wave Ray Tracer

To simulate a realistic channel, an open-source MATLAB ray tracer was used[1]. The ray tracer was built with mmWave propagation in mind and for this reason, given the deep shadow effect that diffraction yields at such high frequencies [6], only specular reflections are considered. In this work, diffuse reflections are ignored, but their importance is undoubted and will thus be part of our future analysis. Currently, the ray tracer accepts Computer-aided Design (CAD) files in AMF format with scenarios described by triangles of different materials.

Specular reflections are computed using the Method of Images (MoI), a basic principle from antenna theory [19]. Given two points in 3D space, i.e., the Transmitter (TX) and the Receiver (RX), and a surface S, the MoI defines the virtual image of the RX ($\mathrm{RX}^{(1)}$) to be the reflection of the RX ($\mathrm{RX}^{(0)}$) across the given surface S. By joining the TX with the $\mathrm{RX}^{(1)}$, it is possible to easily compute the point of specular reflection $\mathrm{P}^{(1)}$ as the intersection of the segment with S. It is necessary, though, to check if the reflection point is inside the bounded surface, otherwise the reflection will be discarded. Furthermore, it is also necessary to check that every other surface of the scenario does not intersect the two segments at any point, otherwise the whole ray will be considered obstructed and thus discarded.

When multiple reflections are considered, the MoI applies recursively. Specifically, given an array of reflecting surfaces $\mathcal{S} = (S_1, \dots, S_N)$, $\mathrm{RX}^{(n)}$ is computed as the virtual image of $\mathrm{RX}^{(n-1)}$ for surface $S_n,\ n = 1, \dots, N$. Then, defining $\mathrm{P}^{(\mathrm{N}+1)} = \mathrm{TX}$, the reflection point $\mathrm{P}^{(\mathrm{n})}$ is computed as the intersection between the surface S_n and the segment joining $\mathrm{P}^{(n+1)}$ and $\mathrm{RX}^{(n)}$. Finally, a check is needed on every path segment $\left(\mathrm{P}^{(n)}, \mathrm{P}^{(n+1)}\right)$ to asses whether it is obstructed by any triangle of the environment. Fig. 1 shows a visual example of the MoI algorithm.

To compute all possible reflections between the RX and TX, a *reflection tree* is created, based on the geometric information extracted from the CAD file. In a reflection tree, all nodes except the root (the TX) correspond to triangles of the environment. For each node, its children coincide to all the visible triangles of the environment with respect to that node. Thus, the depth of the tree corresponds to the maximum reflection order $\eta_{\max}$ (given as an input configuration parameter), i.e., the maximum number of reflections per MPC that the RT computes, and each path from the root to a node at depth d corresponds to an ordered array of d reflecting surfaces. By following all the paths for each tree depth d, all possible array of triangles are tested and thus every reflected ray is computed.

An accurate profiling of the software shows that the most demanding parts of the RT operations are the geometric computations (i.e., computing the position of virtual RXs, computing the point of specular reflection, check if the point is inside the bounded surface) and the obstruction checks. The complexity of the geometric computations is proportional to $\eta_{\max}$, while obstruction checks scale both with $\eta_{\max}$ (every segment has to be checked) and the environment complexity (any triangle can potentially block the propagation of the ray).

Finally, if the ray reflections are valid and it is not ob-

[1]Ray tracer implementation: https://github.com/wigig-tools/qd-realization

structed, path gain is computed as

$$PG = \left(\frac{\lambda}{4\pi d}\right)^2 - \sum_{n=1}^{N} RL_n, \tag{1}$$

where λ is the wavelength (that is a function of the carrier frequency f_c), d is the total distance traveled by the ray, and RL_n is the reflection loss of the material associated to the n-th reflecting surface [16]. Together with the delay, phase, angle of departure, and angle of arrival, the path gain is returned as an output by the ray tracer and written to a file in a specified format, which can be fed as input to other simulators (e.g., link-level or system-level simulators) to compute the channel between the two nodes.

To further simplify the software, no polarization is considered, and rays reflected by a surface experience a $180°$ phase rotation and a reflection loss of $RL_n = 7$–25 dB, depending on the material but irrespective of the angle of incidence.

III. How to Simplify the Millimeter Wave Channel

As introduced in Sec. II, the number of multipath components of the channel between two endpoints heavily affects both the RT computational complexity and the performance of network-level simulators [17]. For this reason, in this work we propose two different strategies to reduce the total number of MPCs, and analyze the effects that this simplification yields on the system behavior. The two strategies we introduce are based on considerations related to the power of each single MPC, based on the idea that weak MPCs do not significantly contribute to the overall signal at the receiver.

For the first simplification approach, we reduce the number of reflections $\eta_{\max}$ that the RT computes. Indeed, when considering an increasing number of reflections, the MPC experiences a decreasing path gain PG, as the absorption on the reflecting surfaces severely degrades the power of the ray and the length of the ray increases. Therefore, MPCs that bounce across multiple scattering surfaces do not impact very much the power at the receiver, and can be omitted from the RT computations. Limiting the maximum reflection order corresponds to setting a bound to the depth of the reflection tree, whose size, as mentioned in the previous section, is exponential on the maximum reflections order $\eta_{\max}$. Therefore, reducing $\eta_{\max}$ reduces the complexity of both the RT and of the network simulators that use it to model the channel.

Following similar considerations, the second strategy aims at discarding the weakest MPCs based on how low their path gain is, regardlessly of how many reflections they actually experience. Using this method, the path gain PG for a ray still needs to be computed, therefore the geometric operations will not be spared. However, the obstruction check may not need to be performed in the case PG is below a certain threshold, and the ray is not accounted for when computing the channel matrix $\mathbf{H}$. The method we propose is based on a threshold γ_{th} which is relative to the $PG_{\max}$ of the strongest MPC for a given channel realization. Notably, the rays with path gain $PG/PG_{\max} < \gamma_{\text{th}}$ are discarded. The path gain associated with the strongest ray $PG_{\max}$ is updated on-line.

TABLE I: Parameters used for simulations.

P_{TX}	30 dBm	TX Array Config.	8×8
f_c	60 GHz	RX Array Config.	4×4
Noise Figure (F)	5 dB	Element pattern	Omni-directional
Bandwidth	400 MHz	Element spacing	$\lambda/2$

Selecting the MPCs with a relative, rather than an absolute threshold, makes it possible to dynamically adapt the simplification to the actual quality of the channel. For example, in Non-Line-of-Sight (NLoS) conditions, the strongest MPC will be given by a reflected ray. Therefore, its path gain $PG_{\max}$ will be comparable to that of a higher number of MPCs (given by other reflections) than in a LoS scenario, where the strongest ray is the direct path, with a much higher power than the reflections. If the reflections have a PG similar to $PG_{\max}$, then the receiver experiences a strong fading. In this case, a *relative* threshold combines accuracy (in NLoS, significant MPCs are still computed) and reduction in complexity (in LoS, several negligible reflected MPCs are not accounted for), more than an absolute threshold, which would apply the same cut in both cases.

We implemented the proposed simplification techniques on top of the open-source ray tracer described in Sec. II. Although such methods are beneficial from a computational point of view, they may have some drawbacks, depending on the reliability requirements of the application for which the ray tracer is used. The most evident downside is that the power spatial distribution can be affected in complex and non-foreseeable ways, as we will show in Sec. IV.

IV. Performance Results

In this section we evaluate through simulations the effects of the ray tracer simplifications we presented in Sec. III. In particular, we investigate the impact of (i) the maximum number of reflections $\eta_{\max}$ per MPC, and (ii) the received power threshold γ_{th} (relative to the strongest path) below which MPCs are are discarded by the ray tracer.

Three scenarios are defined as follows:

1) *Indoor1*: A simple scenario of a box-like room (Fig. 2a) of size $10 \times 19 \times 3$ m. A TX is positioned close to the ceiling, half-way through the wall at $(5, 0.1, 2.9)$ m while the RX, 1.5 m tall, moves inside the room at a speed of 1.2 m/s in spiral-like motion;
2) *L-Room*: Details of this scenario are shown in Fig. 2b. Similarly to the *Indoor1* scenario, the RX of the same height as the previous one moves at a speed of 1.2 m/s;
3) *Parking-Lot*: Outdoor scenario with building around a parking area of about 120×70 m. A TX is positioned on a building 3 m high and a RX is moving at a speed of 4.17 m/s (15 km/h) around the parking lot (see Fig. 2c).

All scenarios have been sampled every 5 ms, thus a total of approximately $9\,000$, $12\,500$, and $15\,000$ time-steps respectively. A list of parameters used in our simulations is shown in Table I. Optimal single-stream SVD-based beamforming is used.

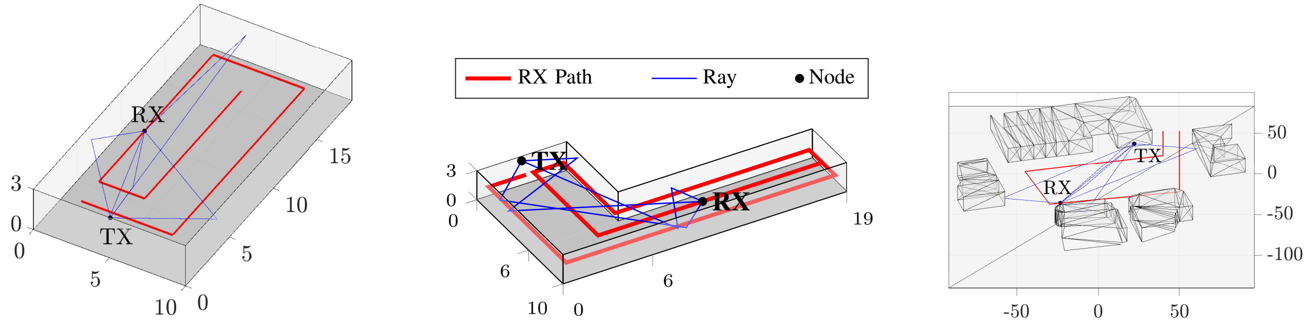

(a) *Indoor1* (b) *L-room* (c) *Parking Lot*

Fig. 2: Visual representations of the proposed scenarios

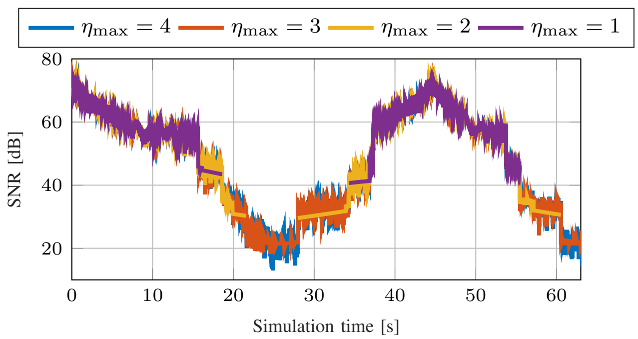

Fig. 3: Temporal evolution of the SNR experienced when a test RX moves in the *L-room* scenario vs. $\eta_{\max}$, fixing $\gamma_{\text{th}} = -\infty$.

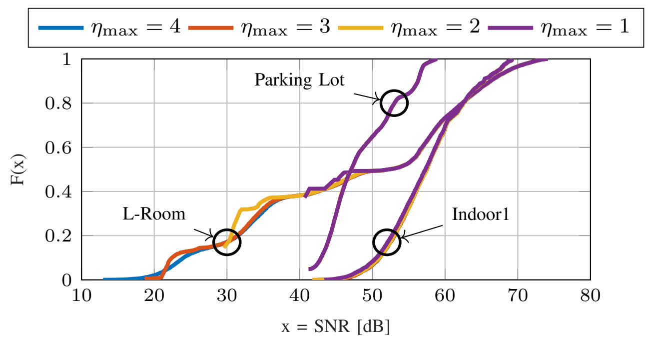

Fig. 4: Cumulative Distribution Function of the SNR when the test RX moves in different simulation scenarios vs. $\eta_{\max}$, with $\gamma_{\text{th}} = -\infty$.

The following performance metrics are considered:[2]

- The RT simulation time T_{RT} [s], i.e., the time taken by the RT software to compute the channel between each pair of nodes at each time-step;
- The MATLAB Net. Sim. Time T_{run} [s], i.e., the time taken by our custom MATLAB simulator to compute the relevant metrics starting from the output of the RT software;
- The Normalized Root Mean Square Error (NRMSE) of the Signal-to-Noise Ratio (SNR), an accuracy indicator that compares the SNR Γ_t experienced when the most accurate RT settings (e.g., with $\eta_{\max} = 4$ and $\gamma_{\text{th}} = -\infty$ for the *L-room* scenario) are considered and the SNR $\hat{\Gamma}$ experienced when different combinations of RT simplifications are applied. Formally, if σ_Γ represents the standard deviation of the baseline SNR Γ, we have

$$\text{NRMSE} = \frac{\text{RMSE}}{\sigma_\Gamma} = \frac{\sqrt{\mathbb{E}\left[\left(\Gamma - \hat{\Gamma}\right)^2\right]}}{\sigma_\Gamma}. \quad (2)$$

In Fig. 3 we consider two nodes and measure the SNR that the probing RX experiences when moving along the path shown in Fig. 2b vs. the maximum number of reflections per MPC $\eta_{\max}$. Rapid variations in the SNR are due to MPCs interfering constructively and destructively, since they observe slightly different path lengths and at least 5 of them have similar path gain in the LoS regions, specifically, the LoS ray and the first-order reflections from ceiling, floor and the two side walls, showing the strong fading in the order of $\lambda = 5$ mm (at 60 GHz). First, we notice that the SNR evolves consistently with the mobility of the RX: the SNR suddenly degrades when the RX enters a NLoS condition and is maximized when it is in LoS with its serving TX, i.e., around time $t = 0$ s and $t = 45$ s. At first glance, it appears that the effect of the RT simplifications is not negligible. In particular, considering the lowest possible value of the relative threshold, i.e., $\gamma_{\text{th}} = -\infty$, the trend of the SNR visibly changes when progressively limiting the maximum number of reflections for each MPC. The impact of those simplifications is particularly evident when the RX operates in NLoS, i.e., when the number of MPCs as little as one (when fading stops) or even none (when no power is received). Despite the above considerations, in the following results we will show more explicitly the accuracy vs. speed trade-off of these parameters in the different scenarios. Furthermore, we will suggest working points for which computation time is significantly reduced with only minor effects on the accuracy of the model.

In Fig. 4 we plot the Cumulative Distribution Function (CDF) of the SNR experienced in different scenarios as a function of the parameter $\eta_{\max}$. The abrupt termination of the CDFs for the *L-Room* and *Parking Lot* scenarios is due to positions of RX/TX for which no ray was able to reach

[2]In this paper, we focus on low-layer performance metrics. In turn, investigating the impact of the proposed RT simplifications on higher-layer performance metrics represents a very interesting research topic that will be part of our future work.

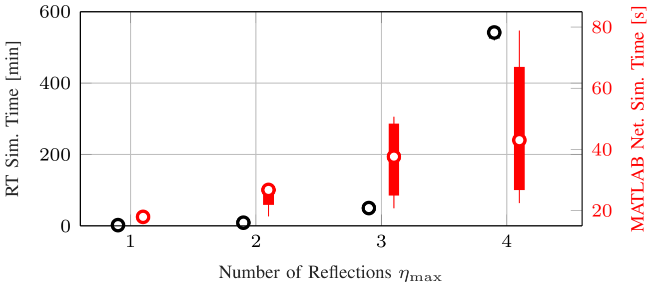

Fig. 5: Box-plots representing the computation time vs. $\eta_{\max}$ when a test RX moves in the *L-room* scenario for all values of γ_{th}. Each box is delimited by the first and the third quartiles of the simulation time, the box's center dot represents the median of the simulation time, and the lines extending from the box (*whiskers*) indicate variability outside the upper and lower quartiles. Each box includes every combination of γ_{th}.

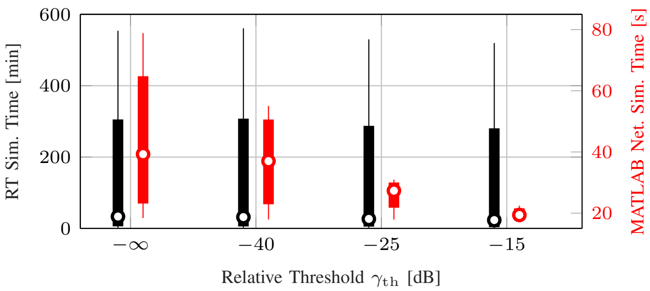

Fig. 6: Box-plots of the computation time vs. γ_{th} when a test RX moves in the *L-room* scenario for all values of $\eta_{\max}$. Each box is delimited by the first and the third quartiles of the simulation time, the box's center dot represents the median of the simulation time, and the lines extending from the box (*whiskers*) indicate variability outside the upper and lower quartiles. Each box includes every combination of $\eta_{\max}$.

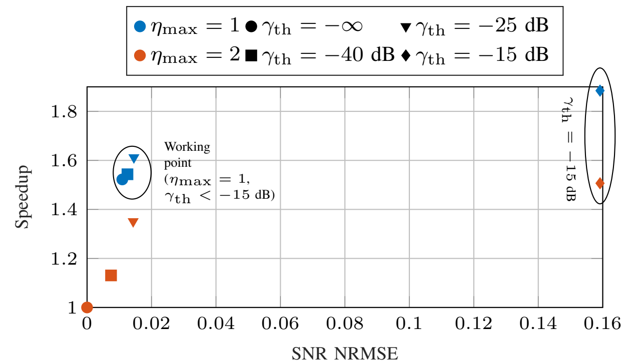

Fig. 7: Speedup vs. SNR NRMSE for different combinations of the RT simplifications when a test RX moves in the *Parking Lot* scenario. $N_{\text{run}} = 1000$ simulations are considered.

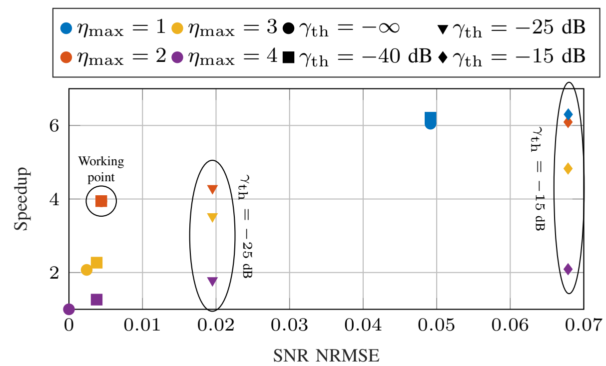

Fig. 8: Speedup vs. SNR NRMSE for different combinations of the RT simplifications when a test RX moves in the *L-room* scenario. $N_{\text{run}} = 1000$ simulations are considered.

with the given $\eta_{\max}$, thus resulting in a complete outage. We observe that, unlike in the *L-room* scenario, in the *Parking Lot* and *Indoor1* scenarios the RX preserves the LoS with its serving TX for the whole duration of the simulation, thereby maintaining very high values of SNR, i.e., above 40 dB. Moreover, Fig. 4 shows that, while in the LoS regime it is possible to reduce the number of reflections $\eta_{\max}$ for each MPC with a minor impact on the accuracy, in the NLoS regime of the *L-room* scenario (i.e., the leftmost part of the figure) the same operation significantly reshapes the CDF of the SNR, thereby confirming the results we obtained in Fig. 3.

On the other hand, decreasing $\eta_{\max}$ speeds up the simulations by several orders of magnitude, as exemplified by the boxplots in Figs. 5 and 6. We can see that the MATLAB simulation time can be reduced by a factor up to $2.4\times$ going from $\eta_{\max} = 4$ to $\eta_{\max} = 1$. The improvement is even more remarkable considering the RT simulation time: the speedup is as significant as $25\times$ considering $\eta_{\max} = 3$, and even $275\times$ for $\eta_{\max} = 4$. Fig. 5 also shows that the configurations with $\eta_{\max} = 3, 4$ exhibit very diverse simulation run time, which is an indication of the increased variability of the channel due to scattering and reflection of the MPCs from nearby surfaces. Similarly, the box-plot in Fig. 6 illustrates that the speedup factor is inversely proportional to the relative threshold γ_{th}, since higher values of γ_{th} makes it possible to reduce the number of MPCs to be processed by the ray tracer as described in Section III, which represents one of the most computationally intensive steps of a simulation and which directly impacts on the channel generation process.

Overall, it is possible to identify which level of simplification is most adequate, i.e., which one provides accurate results while minimizing the overall simulation time. To this aim, in Figs. 7 and 8 we plot the trade-off between simulation speedup and NRMSE of the SNR for the outdoor *Parking Lot* and *L-room* scenarios, respectively. Since, typically, simulations are used to evaluate how changing a set of parameters affects the network performance, and should be repeated with several random seeds to increase the robustness of the obtained results, simulation campaigns would reuse the same RT channel traces for hundreds or thousands of simulations. For this reason, we consider the speedup relative to the total campaign time, which is roughly equal to $T_{\text{TOT}} = T_{\text{RT}} + N_{\text{runs}} T_{\text{run}}$, where T_{RT} is the RT computation time, N_{runs} is the number of independent simulations that are run, and T_{run} is the network simulation run time.

For the *Parking Lot* case (Fig. 7), we can see that all the investigated combinations of simplifications with $\gamma_{\text{th}} < -15$ dB deliver similar values of SNR NRMSE, while reducing the computational complexity with respect to the baseline implementation (i.e., with $\eta_{\max} = 2$ and $\gamma_{\text{th}} = -\infty$). In this scenario, the measured power has limited contribution from the reflected rays, and the optimal approach would be to opt for the configuration with $\eta_{\max} = 1$ and $\gamma_{\rm th} = -25$ dB: the corresponding speedup is around 60% compared to the baseline ray tracing model.

For the *L-room* case (Fig. 8), instead, it is possible to identify two operational regimes. On the one hand, very high (low) values of $\gamma_{\rm th}$ ($\eta_{\max}$) would inevitably lead to a performance degradation in terms of SNR NRMSE, due to the dominant contribution of the reflected signals to the overall received power. On the other hand, reflected rays of order higher than the second have a negligible impact in terms of SNR NRMSE (the gap between the $\eta_{\max} = 2$ and $\eta_{\max} = 3$ configurations, with $\gamma_{\rm th} = -40$ dB, is just 0.001) in the face of a speedup improvement of around 100%. In this scenario, further reducing $\gamma_{\rm th}$ would result in a considerable increase of the system complexity while leading to negligible accuracy gain, and the optimal approach would be to select $\eta_{\max} = 2$ with $\gamma_{\rm th} = -40$ dB.

Finally, we highlight that, while limiting the number of MPCs reduces the ray tracer complexity, it may preclude the implementation of beamforming techniques that exploit the sparsity property of the channel to realize simultaneous beams in independent angular direction (e.g., MIMO techniques exploiting spatial multiplexing). Additionally, while simplifications might have minimal implication on low-layer performance metrics, e.g., the SNR, their effect on higher-layer metrics, e.g., end-to-end throughput and latency, is still unknown and deserves further investigation.

V. Conclusions and Future Work

In this paper, we presented possible simplifications to reduce the complexity of channel modeling through ray tracing. Notably, after an introduction on how a RT works, and on which are the sources of computational complexity for the RT, we discussed two strategies which aim at avoiding computations for MPCs which do not contribute significantly to the overall received power. The first limits the maximum reflection order, while the second removes MPCs with a path gain which is much smaller than that of the strongest ray. We then evaluated the impact of these simplifications on the SNR, in three different scenarios, and on the run time of the RT and of a network simulator. We highlighted that, for each scenario, there exists an optimal working point which minimizes the accuracy loss with respect to the baseline, but reduces the channel generation and modeling time by up to 4 times.

As future works, we will consider a more complex RT, which also includes diffuse components, according to a quasi-deterministic model [20]. Moreover, we will study the impact of the simplifications on the higher layers of the protocol stack, by using the ns-3 802.11ad module [21] (which already integrates the RT) and by extending the ns-3 mmWave module [12] to use RT traces.

References

[1] 3GPP, "NR and NG-RAN Overall Description - Rel. 15," TS 38.300, 2018.

[2] 3GPP, "New WID on extending current NR operation to 71 GHz," Qualcomm, RP-193229 - 3GPP TSG RAN Meeting 86, December 2019.

[3] IEEE, "IEEE Standard for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," *IEEE Std 802.11ad-2012 (Amendment to IEEE Std 802.11-2012, as amended by IEEE Std 802.11ae-2012 and IEEE Std 802.11aa-2012)*, pp. 1–628, Dec 2012.

[4] ITU-R, "IMT Vision - Framework and overall objectives of the future development of IMT for 2020 and beyond," Recommendation ITU-R M.2083, September 2015.

[5] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, March 2014.

[6] S. Deng, G. R. MacCartney, and T. S. Rappaport, "Indoor and Outdoor 5G Diffraction Measurements and Models at 10, 20, and 26 GHz," in *IEEE Global Communications Conference (GLOBECOM)*, Dec 2016.

[7] M. Polese, F. Restuccia, A. Gosain, and J. e. a. Jornet, "MillimeTera: Toward A Large-Scale Open-Source MmWave and Terahertz Experimental Testbed," in *Proceedings of the 3rd ACM Workshop on Millimeter-Wave Networks and Sensing Systems*, ser. mmNets'19. Los Cabos, Mexico: ACM, 2019, p. 27–32.

[8] S. K. Saha, Y. Ghasempour, M. K. Haider, T. Siddiqui, P. D. Melo, N. Somanchi, L. Zakrajsek, A. Singh, R. Shyamsunder, O. Torres *et al.*, "X60: A programmable testbed for wideband 60 GHz WLANs with phased arrays," *Computer Communications*, vol. 133, pp. 77–88, Jan. 2019.

[9] T. Bai and R. W. Heath, "Coverage and rate analysis for millimeter-wave cellular networks," *IEEE Transactions on Wireless Communications*, vol. 14, no. 2, pp. 1100–1114, February 2015.

[10] J. G. Andrews, T. Bai, M. N. Kulkarni, A. Alkhateeb, A. K. Gupta, and R. W. Heath, "Modeling and analyzing millimeter wave cellular systems," *IEEE Transactions on Communications*, vol. 65, no. 1, pp. 403–430, Jan 2017.

[11] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," 3rd Generation Partnership Project (3GPP), Technical Report (TR) 38.901, Jun 2018, version 15.0.0.

[12] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-End Simulation of 5G mmWave Networks," *IEEE Commun. Surveys Tuts*, vol. 20, no. 3, pp. 2237–2263, Third Quarter 2018.

[13] V. Degli-Esposti, F. Fuschini, E. M. Vitucci, M. Barbiroli, M. Zoli, L. Tian, X. Yin, D. A. Dupleich, R. Müller, C. Schneider, and R. S. Thomä, "Ray-Tracing-Based mm-Wave Beamforming Assessment," *IEEE Access*, vol. 2, pp. 1314–1325, 2014.

[14] S. G. Larew, T. A. Thomas, M. Cudak, and A. Ghosh, "Air interface design and ray tracing study for 5G millimeter wave communications," in *IEEE Globecom Workshops (GC Wkshps)*, Dec 2013, pp. 117–122.

[15] A. Maltsev, A. Pudeyev, A. Lomayev, and I. Bolotin, "Channel modeling in the next generation mmWave Wi-Fi: IEEE 802.11ay standard," in *22th European Wireless Conference*, May 2016.

[16] C. Lai, R. Sun, C. Gentile, P. B. Papazian, J. Wang, and J. Senic, "Methodology for Multipath-Component Tracking in Millimeter-Wave Channel Modeling," *IEEE Transactions on Antennas and Propagation*, vol. 67, no. 3, pp. 1826–1836, March 2019.

[17] P. Testolina, M. Lecci, M. Polese, M. Giordani, and M. Zorzi, "Scalable and Accurate Modeling of the Millimeter Wave Channel," in *International Conference on Computing, Networking and Communications (ICNC)*, 2019.

[18] M. Polese and M. Zorzi, "Impact of Channel Models on the End-to-End Performance of mmWave Cellular Networks," in *IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, June 2018.

[19] Z. Yun and M. F. Iskander, "Ray Tracing for Radio Propagation Modeling: Principles and Applications," *IEEE Access*, vol. 3, pp. 1089–1100, 2015.

[20] 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, May 2018.

[21] H. Assasa, J. Widmer, T. Ropitault, and N. Golmie, "Enhancing the Ns-3 IEEE 802.11ad Model Fidelity: Beam Codebooks, Multi-Antenna Beamforming Training, and Quasi-Deterministic MmWave Channel," in *Proceedings of the 2019 Workshop on Ns-3*, ser. WNS3 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 33–40. [Online]. Available: https://doi.org/10.1145/3321349.3321354