Improving the Efficiency of MIMO Simulations in ns-3

Matteo Pagin
University of Padova
Padova, Italy
paginmatte@dei.unipd.it

Sandra Lagén
Biljana Bojović
Centre Tecnològic de Telecomunicacions de Catalunya
Barcelona, Spain
{sandra.lagen,biljana.bojovic}@cttc.es

Michele Polese
Institute for the Wireless Internet of Things,
Northeastern University
Boston, MA, USA
m.polese@northeastern.edu

Michele Zorzi
University of Padova
Padova, Italy
zorzi@dei.unipd.it

ABSTRACT

Channel modeling is a fundamental task for the design and evaluation of wireless technologies and networks, before actual prototyping, commercial product development and real deployments. The recent trends of current and future mobile networks, which include large antenna systems, massive deployments, and high-frequency bands, require complex channel models for the accurate simulation of massive MIMO (m-MIMO) in millimeter wave (mmWave) and Terahertz (THz) bands. To address the complexity/accuracy trade-off, a spatial channel model has been defined by 3GPP (TR 38.901), which has been shown to be the main bottleneck of current system-level simulations in ns-3. In this paper, we focus on improving the channel modeling efficiency for large-scale MIMO system-level simulations. Extensions are developed in two directions. First, we improve the efficiency of the current 3GPP TR 38.901 implementation code in ns-3, by allowing the use of the Eigen library for more efficient matrix algebra operations, among other optimizations and a more modular code structure. Second, we propose a new performance-oriented MIMO channel model for reduced complexity, as an alternative model suitable for mmWave/THz bands, and calibrate it against the 3GPP TR 38.901 model. Simulation results demonstrate the proper calibration of the newly introduced model for various scenarios and channel conditions, and exhibit an effective reduction of the simulation time (up to 16 times compared to the previous baseline) thanks to the various proposed improvements.

CCS CONCEPTS

• **Networks** → **Network simulations**; **Network performance analysis**.

WNS3 2023, June 28–29, 2023, Arlington, VA, USA

ACM ISBN 979-8-4007-0747-6/23/06…$15.00
https://doi.org/10.1145/3592149.3592167

KEYWORDS

channel model, mmWave, Terahertz, network simulation, ns-3

ACM Reference Format:
Matteo Pagin, Sandra Lagén, Biljana Bojović, Michele Polese, and Michele Zorzi. 2023. Improving the Efficiency of MIMO Simulations in ns-3. In *2023 Workshop on ns-3 (WNS3 2023), June 28–29, 2023, Arlington, VA, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3592149.3592167

1 INTRODUCTION

Mobile networks play a key role in our society and are poised to become ever more important in the coming years. In fact, the International Telecommunications Union (ITU) foresees that in 2030 and beyond wireless broadband will be ubiquitous, and will be required to provide connectivity not only to humans, but also to a plethora of intelligent devices such as wearables, road vehicles, Unmanned Aircraft Systems (UASs) and robots [16]. Moreover, novel use cases such as holographic communications, Extended Reality (XR) and tactile applications will further exacerbate the throughput and latency requirements which were posed by enhanced Mobile Broadband (eMBB) and Ultra-Reliable Low-Latency Communications (URLLC) [15].

To meet these goals, future cellular systems will further evolve 5th generation (5G) networks, which have introduced a flexible, virtualized architecture, the support for mmWave communications and the use of m-MIMO technologies [9]. Notably, the research community is considering a more central role for mmWaves, a further expansion of the spectrum towards the THz band, and an Artificial Intelligence (AI)-native network design, with the goal of achieving autonomous data-centric orchestration and management of the network [35], possibly down to the air interface [14].

The THz and mmWave bands offer large chunks of untapped bandwidth which operators can leverage to meet the Tb/s peak rates that are envisioned by the ITU [16]. However, this portion of the spectrum is plagued by unfavorable propagation characteristics, comprising a marked free-space propagation loss and susceptibility to blockages [11, 20], which make it challenging to harvest its potential. Although the harsh propagation environment can be partially mitigated by using directional links and densifying network deployments [36], the support for mmWave and THz bands entails a major redesign not only of the physical layer, but of the whole cellular protocol stack [42]. For instance, the intrinsic directionality of

the communication requires ad hoc control procedures [13], while the frequent transitions between Line-of-Sight (LoS) and Non-Line-of-Sight (NLoS) conditions call for an ad hoc transport layer design, such as novel Transmission Control Protocol (TCP) algorithms [45].

In addition, as the network progressively becomes increasingly complex and heterogeneous, the push for spectrum expansion will be coupled with an AI-native design which, thanks to the ongoing virtualization, will not be limited to the radio link level, but will encompass the orchestration of large scale deployments as well [34]. Nevertheless, how to design, test and eventually deploy management and orchestration algorithms is an open research challenge [33]. First, the training data must accurately capture the interplay of the whole protocol stack with the wireless channel. Furthermore, optimization frameworks such as Deep Reinforcement Learning (DRL) also call for preliminary testing in isolated yet realistic environments, with the goal of minimizing the performance degradation to actual network deployments [5, 22].

In these regards, system-level network simulators have a central role to play. Indeed, an end-to-end evaluation of algorithms and protocols becomes paramount when considering frequencies above 6 GHz, given the impact of their peculiar propagation characteristics on the whole protocol stack. At the same time, end-to-end simulators can also serve as both sources of training data for AI models, and testing platforms for preliminary evaluation of Machine Learning (ML) algorithms prior to their deployment in commercial networks. However, the suitability of end-to-end network simulators to these tasks largely depends on the accuracy of the channel model [43] and on the scalability for realistically-sized deployments. In fact, system-level simulators generally abstract the actual link-level transmission via an error model, which maps the Signal-to-Interference-plus-Noise Ratio (SINR) of the wireless link to a packet error probability [23]. Eventually, the latter is used to determine whether the packet has been successfully decoded by the receiver. As a consequence, the accuracy of the simulator heavily depends on the reliability of the SINR estimation, especially when considering the mmWave and THz bands.

The well known ns-3 simulator features the implementation of the 3rd Generation Partnership Project (3GPP) channel model [1], which, according to the 3GPP, represents the state-of-the-art channel model for drop-based end-to-end simulations of devices operating at frequencies between 0.5 to 100 GHz. Despite its accuracy, the TR 38.901 channel model is particularly demanding from a computational point of view, and thus limits the scalability of the simulated scenarios. At the same time, the simpler channel models which are found in analytical studies fail to capture the peculiar characteristics of mmWave and THz links.

To fill this gap, in this paper we propose optimizations to the ns-3 implementation of the TR 38.901 channel model of [47], both at the codebase and at the design level, which aim to provide wireless researchers with the tools for simulating future dense wireless scenarios in a computationally efficient manner. Specifically, we significantly improve the runtime of simulations involving the 3GPP TR 38.901 channel model [1] by porting the intensive linear algebra operations to the open-source library `Eigen` [10]. To this end, we also design and implement a set of common linear algebra APIs, which increase the modularity of the `spectrum` module with respect to the underlying data structures and algorithms. Then, we propose a simplified channel model, based on [1], which aims to provide an additional order of magnitude of runtime reduction, at the cost of a slight accuracy penalty. Profiling results show that the support for `Eigen`, coupled with further TR 38.901 optimizations, leads to a decrease of up to 5 times in the simulation time of typical Multiple Input Multiple Output (MIMO) scenarios. Furthermore, the proposed performance-oriented channel model further improved the runtime of simulations, which now take as low as 6 % with respect to the full TR 38.901 channel model, with a negligible loss in accuracy.

The remainder of the paper is organized as follows. Section 2 reports the state of the art on channel models. Sections 3 and 4 describe the contributions of this work, in terms of optimizations to the ns-3 implementation of the TR 38.901 framework and the design of a performance oriented channel model, respectively. Finally, Section 5 presents benchmarks of the introduced optimizations and discusses the main use cases of these channel models, while Section 6 concludes the paper by mentioning possible future extensions of this work.

2 RELATED WORK

Channel modeling is a fundamental task for the design, simulation, and evaluation of current and future wireless networks. It is especially relevant to perform system-level simulations to test new algorithms, procedures, and architectures, before going into real deployment/device implementations. In the recent decades, the challenges for understanding the propagation at mmWave and THz frequencies with large antenna arrays and the use of MIMO have further motivated the channel modeling efforts in those frequency ranges [12, 25]. As a result, multiple channel measurement campaigns have been performed by the academic and industry communities [39], leading to different families of channel models. The various channel models differ in their degree of simplicity and accuracy. They range from simple models that just consider a propagation loss component combined with Nakagami-m or Rayleigh fading but fail to capture the spatial dimension of the channel and the interactions with beamforming [3], to deterministic models that are very accurate in specific scenarios but are much more complex and require a precise characterization of the environment [24]. To address the complexity-accuracy trade-off, the 3GPP has adopted a stochastic channel model for simulations of 5G and beyond networks [1]. Stochastic channel models are generic, thanks to their stochastic nature, but at the same time can model interactions with multiple-antenna arrays. Specifically, the 3GPP defined in TR 38.901 the spatial channel model for simulations that address frequency ranges from 0.5 GHz to 100 GHz [1], which is parameterized for various simulation scenarios, including indoor office, indoor factory, urban macro, urban micro, and rural macro.

However, for system-level simulations of large-scale systems including multiple nodes and large antenna arrays, the 3GPP spatial channel model still introduces a significant overhead in terms of computational complexity. In this line, in [37], a simplified channel model for the system-level simulations of MIMO wireless networks is proposed. Therein, the end-to-end channel gain is obtained as the sum of several loss and gain terms that account for large-scale phenomena such as path loss and shadowing, small-scale fading,

and antenna and beamforming gains. Notably, the latter terms represent a fundamental component for studies concerning modern wireless systems. In particular, an accurate characterization of the antenna radiation pattern and of the effect of the presence of multiple radiating elements becomes extremely important when studying mmWave and THz frequencies. Following the model in [37], the combined antenna and beamforming gain can be computed according to [40], the path loss and shadowing components can follow the 3GPP model in [1], and the small-scale fading can be sampled from various statistical distributions. For the small-scale fading, authors in [37] propose to use a Nakagami-m distribution, which has been shown to provide a good fit with the 3GPP model, provided that the m parameter is appropriately chosen. Another option for small-scale fading modeling is the so-called Fluctuating Two-Ray (FTR) fading model presented in [41], which models more accurately the fading that occurs at mmWaves.

The 3GPP TR 38.901 spatial channel model was included in ns-3 thanks to the efforts of Tommaso Zugno in 2019 Google Summer of Code [47], and later extended to address vehicular scenarios in [46] and industrial scenarios in [38]. The current spatial channel model implemented in ns-3 is very accurate for simulations in line with 3GPP specifications for a wide range of frequencies, but represents the main bottleneck in terms of computational complexity when considering large-scale simulations with many multi-antenna nodes, especially when equipped with large antenna arrays. This is because of the intrinsic complexity in the generation of the channel model according to 3GPP specifications and the need to deal with 3D channel matrix structures. The channel matrix in the ns-3 implementation of the 3GPP spatial channel model is implemented as a 3D structure whose dimensions depend on the number of the transmit antennas, receive antennas, and clusters. Currently, in ns-3, the 3GPP channel model uses a vector of vectors of vectors to represent 3D arrays, such as the channel matrix.

The design of computationally efficient yet accurate channel models has been a topic of interest also in the Wireless Local Area Network (WLAN) space. The authors of [18, 19] present a frequency-selective channel for WLANs, and use Exponential Effective SNR Mapping (EESM) Link-to-System (L2S) mapping to integrate their model with the ns-3 system-level Wi-Fi implementation. Moreover, they develop a framework which leverages cached statistical channel matrix realizations to directly estimate the effective Signal-to-Noise Ratio (SNR), thus further improving the computational efficiency of the model. Specifically, the latter is modeled as a parameterized log-SGN random variable. They extend their work in [17], by accounting for the channel correlation over time. Moreover, [26] compares statistical channel models for the 60 GHz band with the Quasi Deterministic (QD) Ray Tracer (RT) of [4].

In this paper, we summarize the efforts carried out by Matteo Pagin in 2022 Google Summer of Code to further optimize the code in ns-3 in two directions: 1) improving the efficiency of the code by allowing the use of `Eigen` library, and 2) proposing a new performance-oriented MIMO channel model for reduced complexity in ns-3 large-scale simulations. First, we have improved the efficiency of the 3GPP spatial channel model in ns-3 by allowing the usage of `Eigen` to represent matrices, so that when `Eigen` is available the 3GPP channel matrix is represented as an `std::vector` of `Eigen` matrices. This already improves the performance of current models. Second, we propose an alternative model, based on the FTR channel model [41], in which the channel is characterized by a single scalar instead of 3D matrices, and we have calibrated such model to align with the 3GPP TR 38.901 spatial channel model for various scenarios and channel conditions. This model is especially useful to speed up ns-3 large-scale simulations, when simplicity is prioritized.

3 EFFICIENT MIMO MODELING WITH THE EIGEN LIBRARY

The use of multiple antennas both at the transmitter and at the receiver, a fundamental feature of modern wireless systems, makes a scalar representation of the channel impulse response insufficient. Instead, MIMO channels are usually represented in the form of a complex matrix $\boldsymbol{H} \in \mathbb{C}^{U \times S}$, whose elements depict the channel impulse response between the U and S radiating elements of the transmitting and receiving antenna arrays, respectively [1]. This peculiarity significantly increases the computational complexity of MIMO channel models, compared to Single Input Single Output (SISO) ones, since the complex gain of the channel must be evaluated for each pair of transmit and receive antennas. Notably, previous analyses identified in statistical channel models the main bottleneck for system-level MIMO wireless simulations. In typical m-MIMO 5G scenarios, where the devices feature a high number of antennas, the channel matrix generation and the computation of the beamforming gain represent up to 90% of the simulation time [43].

In light of these limitations, as the first of our contributions, we optimized the implementation of the 3GPP TR 38.901 model in ns-3 introduced in [47]. First, we observed that, as of ns-3.37, part of the trigonometric operations of the `GetNewChannel` method of the `ThreeGppChannelModel` class are unnecessarily repeated for each pair of transmitting and receiving radiating elements. This represents a significant inefficiency, since the inputs of these functions, i.e., the angular parameters of the propagation clusters, depend on the cluster index only. Moreover, the standard library `sin` and `cos` functions are particularly demanding to evaluate. Therefore, we cached the trigonometric evaluations of these terms prior to the computation of $\boldsymbol{H}$'s coefficients, effectively reducing the complexity of the trigonometric operations from $O(U \times S \times N)$ to $O(N)$, where N is the number of propagation clusters.

Then, we focused on improving the algebra manipulations of the channel matrix performed in the `ThreeGppSpectrumPropagationLossModel` by introducing the support for the open-source library `Eigen` in ns-3. `Eigen` is a linear algebra C++ template library that offers fast routines for algebra primitives such as matrix multiplication, decomposition and space transformation [10], and is used by many open-source frameworks such as TensorFlow.

We set `Eigen` as an optional, external ns-3 dependency, with the goal of minimizing future code maintenance efforts, and thus mimicking the support for other third-party libraries. To get `Eigen`, ns-3 users can either rely on packet managers, i.e., install the package `libeigen3-dev` (`eigen`) for Linux (Mac) systems, or manually install the library by following the official instructions[1]. Then, `Eigen` can be enabled via a custom flag defined in the `macros-and--definitions.cmake` file, and its presence in the system is shown to

[1] https://gitlab.com/libeigen/eigen/-/blob/master/INSTALL

the user by exposing whether it has been found or not via the `ns3--config-table.cmake` file. The latter also defines the preprocessor definition `HAVE_EIGEN3`, which is used in the ns-3 source files to discern `Eigen`'s availability. Finally, the linking of `Eigen` with the ns-3 source files is taken care of by the `CMake` configuration file provided by the library itself, as suggested in the related ns-3 guide.

To prevent the need for `Eigen` to be installed in the host system, we developed a common set of APIs between the `Eigen`- and the Standard Template Library (STL)-based data structures and primitives. Thanks to this choice, the remainder of the `spectrum` code is completely abstracted with respect to the presence of the library. Given that most of the needed operators can not be overloaded for STL C++ vectors (for instance, `operator()`), the common interface for both Eigen and STL's based vectors and matrices has been implemented by defining ad hoc structs with custom operators. In particular, we defined:

- The complex vector type `PhasedArrayModel::ComplexVector`. This data-structure is defined as an `std::vector` of `std::complex<double>` whenever `Eigen` is not installed, and as an `Eigen` vector of `std::complex<double>` otherwise. The set of APIs includes operators `[]` and `!=`, which can be used to access the vector entries and to compare pairs of vectors, respectively. Additionally, we defined the STL-like methods `size`, `norm` and `resize`, which return the vector size, its $\mathcal{L}^2$-norm, and allow the user to resize the underlying container, respectively. These definitions follow the typical STL notation, as it is supported by `Eigen` as well.
- The complex matrix type `MatrixBasedChannelModel::Complex2DVector`. In this case, the underlying type is a nested `std::vector` of `std::complex<double>` for when `Eigen` is disabled, and an `Eigen` matrix whose entries are of type `std::complex<double>` otherwise.
  In this case, we aligned the notation to the APIs provided by `Eigen`. Specifically, the matrix elements can be accessed via the operator `()`, which takes as arguments the row and column indices of the entry, while the method `resize` allows users to resize matrices by specifying the number of rows and columns. In turn, these can be accessed via the `rows` and `columns` methods, respectively.
- The 3D matrix `MatrixBasedChannelModel::Complex3DVector`. This data structure is defined, regardless of `Eigen`'s availability, as an `std::vector` of `MatrixBasedChannelModel::Complex2DVector`. In this case, the only method provided is `MultiplyMatByLeftAndRightVec`, which computes a product of the type $\boldsymbol{w}_T \boldsymbol{H} \boldsymbol{w}_R^T$, where $\boldsymbol{H} \in \mathbb{C}^{U \times S}$, $\boldsymbol{w}_T \in \mathbb{C}^{1 \times U}$ and $\boldsymbol{w}_R \in \mathbb{C}^{1 \times S}$. Notably, this computationally demanding evaluation, which is required for computing the beamforming gain in `ThreeGppSpectrumPropagationLossModel`, leverages `Eigen`'s optimized algorithms whenever the library is installed in the host system.

Finally, we remark that the support for `Eigen` in the ns-3 codebase can possibly be further extended to improve the efficiency of other linear algebra operations, such as the Singular Value Decomposition (SVD) which is used in the `mmwave` and `nr` modules to compute optimal beamformers, and the matrix-by-matrix multiplications needed for relayed channels [30].

4 A PERFORMANCE-ORIENTED MIMO STATISTICAL CHANNEL MODEL

The second approach to reduce computational complexity we propose in this paper is a MIMO channel model for simulating large m-MIMO scenarios, implemented in the class `TwoRaySpectrumPropagationLossModel`. The goal of this auxiliary model is to offer a faster, albeit slightly less accurate, statistical channel model than the 3GPP TR 38.901 framework of [47] by preventing the need for the computation of the complete channel matrix. In line with [1], the frequency range of applicability of this model is $0.5 - 100$ GHz, although the framework can be possibly extended to support higher frequencies as well.

The overall channel model design follows the approach of [37], i.e., the end-to-end channel gain is computed by combining several loss and gain terms which account for both large- and small-scale propagation phenomena, and the antenna and beamforming gains. In particular, let T be a device transmitting a signal x with power P_T^x, and R be another device in the simulation (which may or may not be the intended destination of x). The proposed model implements the `PhasedArraySpectrumPropagationLossModel` interface by estimating P_R^x, i.e., the power of x received at R, as follows:

$$\begin{aligned}\mathrm{P}_R^x[dBm] = \mathrm{P}_T^x[dBm] - \mathrm{PL}_{T,R}[dB] \\ + \mathrm{S}_{T,R}[dB] + G_{T,R}[dB] + F_{T,R}[dB],\end{aligned} \tag{1}$$

where the terms $\mathrm{PL}_{T,R}$ and $\mathrm{S}_{T,R}$ represent the path loss and the shadowing, respectively, while $G_{T,R}$ and $F_{T,R}$ denote the antenna and beamforming gain and the small-scale fading, respectively. The remainder of this section describes in detail how each of these terms is computed.

4.1 Path loss, Shadowing, and LoS Condition

The large-scale propagation phenomena are modeled according to the 3GPP TR 38.901 model [1], since its implementation of [47] is not computationally demanding. Nevertheless, the channel model can in principle be coupled with arbitrary classes extending the `ChannelConditionModel` interface.

Specifically, we first determine the 3GPP scenario. Then, for each link we set the LoS condition in a stochastic manner, using the class extending `ThreeGppChannelConditionModel` which corresponds to the chosen scenario.

Then, we compute the path loss using the 3GPP TR 38.901 formula

$$PL_{T,R} = A \log_{10}(d) + B + C \log_{10}(f_C)[dB], \tag{2}$$

where d is the 3D distance between the transmitter and the receiver, f_C is the carrier frequency, and A, B and C are model parameters which depend on the specific scenario and the LoS condition.

To account for the presence of blockages, an optional log-normal shadowing component $S_{T,R}$ and an outdoor-to-indoor penetration loss term are added to $PL_{T,R}$.

4.2 Antenna and Beamforming Gain

The combined array and beamforming gain is computed using the approach of [40]. The proposed model supports the presence of multiple antenna elements at the transmitter and at the receiver, and arbitrary analog beamforming vectors and antenna radiation patterns. Therefore, ns-3 users can use this model in conjunction

with any class that implements the `AntennaModel` interface. In this implementation, we focus on Uniform Planar Arrays (UPAs), although the methodology is general and can be applied to arbitrary antenna arrays.

Let θ and φ be the relative zenith and azimuth angles between transmitter and receiver, respectively, and let $\boldsymbol{w}(\theta_0, \varphi_0)$ denote the beamforming vector pointing towards the steering direction (θ_0, φ_0). We denote with $U = U_h U_v$ the total, horizontal, and vertical number of antenna elements, respectively, and with d_h, d_v their spacing in the horizontal and vertical domains of the array, respectively.

Considering first isotropic antennas, the gain pattern of a UPA, in terms of received power relative to a single radiating element, can be expressed as [6]

$$G^{iso}_{T,R}(\theta, \varphi) = \left| \boldsymbol{a_i}^{\mathrm{T}}(\theta, \varphi) \boldsymbol{w}(\theta_0, \varphi_0) \right|^2, \tag{3}$$

where $\boldsymbol{a_i}(\theta, \varphi)$ is the array response vector, whose generic entry m, n with $m \in \{0, \ldots, U_v - 1\}, n \in \{0, \ldots, U_h - 1\}$ reads

$$a_i(\theta, \varphi)_{m,n} = \exp\left(j \frac{2\pi}{\lambda} m d_v \cos(\theta) \right) \times \exp\left(j \frac{2\pi}{\lambda} n d_h \sin(\theta) \sin(\varphi) \right). \tag{4}$$

In this work, which supports arbitrary antennas, each antenna element (m, n) actually exhibits a generic radiation pattern $g(\theta, \varphi)_{m,n}$ towards direction (θ, φ). In particular, we assume that $g(\theta, \varphi)_{m,n}$ is constant for all the elements of the array, i.e., $g(\theta, \varphi)_{m,n} \equiv g(\theta, \varphi)$. Accordingly, we compute $G_{T,R}(\theta, \varphi)$ in the `ComputeBeamformingGain` function of the `TwoRaySpectrumPropagationLossModel` class as

$$G_{T,R}(\theta, \varphi) = G^{iso}_{T,R}(\theta, \varphi) \left| g(\theta, \varphi) \right|^2. \tag{5}$$

Figures 1a and 1b report $G_{T,R}(\theta, \varphi)$ for both the isotropic (`IsotropicAntennaModel`) and the 3GPP (`ThreeGppAntennaModel`) radiation patterns, respectively.

It can be noted that our model abstracts the computation of the received signal power as a SISO keyhole channel [7], which is then combined with the spatial antenna gain patterns at the transmitter/receiver to obtain the received power. This approximation is possibly imprecise when considering NLoS links, due to the lack of a dominant multipath component. To account for this limitation, we introduce a multiplicative correction factor η which scales the beamforming gain as $G'_{T,R}(\theta, \varphi) \equiv \eta G_{T,R}(\theta, \varphi)$. In line with [21], we set $\eta = 1/19$.

4.3 Fast Fading

The widely used Rayleigh and Rician distributions fail, even in their generalized forms, to capture the intrinsic bimodality exhibited by mmWave scenarios [8, 28, 44]. Therefore, in our implementation we model fast fading using the more general FTR model of [41]. This fading model assumes that the received signal comprises two dominant specular components and a mixture of scattered paths, thus modeling the amplitude of the received signal V_r as

$$V_r = \sqrt{\xi} \exp(j\phi_1) + \sqrt{\xi} \exp(j\phi_2) + X + jY, \tag{6}$$

where ϕ_1, ϕ_2 are statistically independent random phases, distributed as $\phi_i \sim \mathcal{U}[0, 2\pi]$. X and Y are independent Gaussian

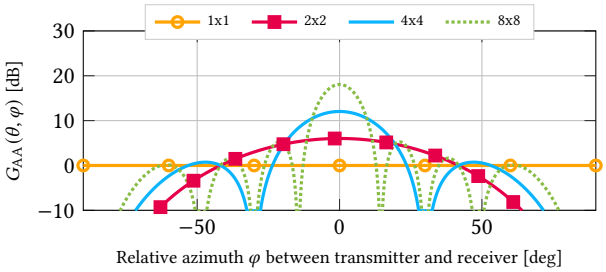

(a) Isotropic Radiating Elements

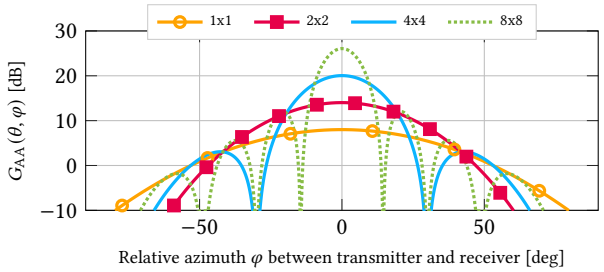

(b) Directional Antenna Radiation Pattern of [1, Section 7.3]

Figure 1: Overall Array and Beamforming Gain of a Uniform Planar Array, for Isotropic and 3GPP [1, Section 7.3] Radiating Elements and {1x1, 2x2, 4x4, 8x8} UPAs. The Steering Direction is Fixed to $(\theta_0, \varphi_0) = (0°, 0°)$, and $\theta \equiv 0°$

random variables, i.e., $X, Y \sim \mathcal{N}(0, \sigma^2)$, which represent the diffuse component of the received signal, which is assumed to be the superposition of multiple weak scattered waves with independent phase. Finally, ξ is a unit-mean Gamma distributed random variable with rate m and Probability Density Function (PDF)

$$f_\xi(u) = \frac{m^m u^{m-1}}{\Gamma(m)} exp(-mu). \tag{7}$$

In our implementation, $F_{T,R} = |V_r|^2$ is sampled via the `GetFtrFastFading` function of the `TwoRaySpectrumPropagationLossModel` class.

The FTR fading model is usually expressed as a function of the Gamma rate m and the auxiliary parameters

$$K \doteq \frac{V_1^2 + V_2^2}{2\sigma^2} \tag{8}$$

$$\Delta \doteq \frac{2V_1 V_2}{V_1^2 + V_2^2} \in [0, 1], \tag{9}$$

where K represents the ratio of the power of the specular components with respect to the diffuse ones, while Δ denotes how similar the received powers of the specular components are. By tuning these parameters, a high degree of flexibility can be achieved. Notably, a choice of $\Delta = 0$ effectively yields a Rician-distributed signal amplitude [41].

4.3.1 Calibration. In our work, we calibrated the V_1, V_2 and m parameters of the FTR fading model using the full 3GPP TR 38.901 channel model as a reference. In particular, we first obtained the statistics of the small-scale fading of the 3GPP model, using an ad hoc

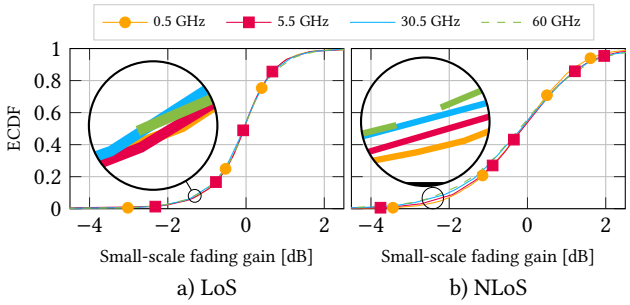

Figure 2: Small-scale Fading Gain Statistics for the UMi Propagation Scenario Versus the Carrier Frequency f_C, for both LoS and NLoS Channel Conditions

calibration script (`three-gpp-two-ray-channel-calibration.cc`). The script produces a collection of channel gain samples obtained by using the `ThreeGppSpectrumPropagationLossModel` and the `ThreeGppChannelModel` classes, and neglecting the beamforming gain, path-loss, shadowing and blockages. Accordingly, we isolate the variation around the mean received power caused by the small-scale fading only. A separate set of these samples has been retrieved for both LoS and NLoS channel conditions, the different propagation scenarios of [1], and a set of carrier frequencies ranging from 0.5 to 100 GHz. However, a preliminary evaluation of the obtained data showed a negligible dependence of the small-scale fading with respect to the carrier frequency, as can be observed in Figure 2. Therefore, we calibrated the FTR parameters considering only the channel condition and the propagation scenario.

The small-scale fading samples have been used to estimate the Δ, K and m FTR parameters, and then derive analytically the values of V_1 and V_2 yielding the fading realizations that are the closest (in a goodness-of-fit sense) to the TR 38.901 model. To this end, we defined a discrete grid of FTR parameters, spanning their whole domain, and considered the corresponding set of parameterized FTR distributions. To find the best matching one, we measured the distance between each of these distributions and the 3GPP reference curves by using the Anderson-Darling goodness-of-fit test [2]. This test is used to discern whether a sorted collection of n samples $\{Y_1, \ldots, Y_n\}$ originates from a specific distribution, by evaluating the test statistic [2]

$$A^2 = -n - S(\mathcal{F}), \tag{10}$$

where

$$S(\mathcal{F}) = \sum_{i=1}^{n} \frac{2i-1}{n} \left[\ln(\mathcal{F}(Y_i)) + \ln\left(1 - (\mathcal{F}(Y_{n+1-i}))\right)\right], \tag{11}$$

and $\mathcal{F}$ is the Cumulative Distribution Function (CDF) of the target distribution. In the standard Anderson-Darling test, A^2 is then compared to a pre-defined critical value to validate the hypothesis. Instead, in our work we find the FTR distribution $\mathcal{F}_{m,K,\Delta}$ which yields the lowest S. Specifically, for each combination of propagation scenario, LoS condition and corresponding samples $\{Y_1, \ldots, Y_n\}$ we find

$$\mathcal{F}_{m^*,K^*,\Delta^*} \doteq \underset{m,K,\Delta}{\operatorname{argmin}} S(\mathcal{F}_{m,K,\Delta}). \tag{12}$$

Finally, we exported the calibrated FTR parameters into ns-3, by storing them in `SIM_PARAMS_TO_FTR_PARAMS_TABLE`, i.e., an `std::map` which associates the propagation scenario and condition to the corresponding best fitting FTR parameters. We remark that this calibration process represents a pre-computation step which needs to be done only once. Indeed, when running a simulation with this channel model, the FTR parameters get simply retrieved from the pre-computed lookup table by the `GetFtrParameters` function. Nevertheless, for the sake of reproducibility and maintainability of the code, we provide this functionality in the Python script `two-ray-to-three-gpp-ch-calibration.py`.

5 BENCHMARKS, EXAMPLES AND USE CASES

In this section, we provide an example on how to use the performance-oriented channel model presented above, in conjunction with the New Radio (NR) [32] module, to simulate 5G MIMO networks. Moreover, we present benchmarks which quantify the simulation time reduction achieved with this work, and we outline some possible use cases.

5.1 Examples and Benchmarks

We demonstrate how to use the performance-oriented channel model in the `cttc-nr-demo-two-ray` script, i.e., a custom version of the `cttc-nr-demo` example which is included in the NR module. The script deploys N_{gNB} 5G NR base stations, along with N_{UE} users in each cell. Each User Equipment (UE) uploads data using two Bandwidth Parts (BWPs) operating at 28 and 30 GHz, respectively. Both base stations and user terminals feature UPAs with multiple radiating elements.

Most simulation parameters can be tuned by ns-3 users. Notably, the script provides the possibility to choose whether to use the 3GPP TR 38.901 channel model of [47] or the FTR-based channel model proposed in this work. In such regard, the use of the `TwoRaySpectrumPropagationLossModel`, instead of the TR 38.901 one, is achieved by:

(1) Setting the `TypeId` of the `SpectrumPropagationLossModel` factory to `TwoRaySpectrumPropagationLossModel`;
(2) Creating an instance of the `TwoRaySpectrumPropagationLossModel` class using the above factory, and setting the corresponding pointer as the `SpectrumPropagationLossModel` of both BWPs;
(3) Setting the attribute `Frequency` of the `TwoRaySpectrumPropagationLossModel` instance as the BWP carrier frequency;
(4) Specifying the 3GPP propagation scenario by setting the attribute `Scenario`; and
(5) Creating and setting the `ChannelConditionModel` by using the `TwoRaySpectrumPropagationLossModel` class `ChannelConditionModel` attribute.

On the other hand, the `Eigen` optimizations simply require users to have the corresponding library installed in their system, and to enable `Eigen` when configuring ns-3, using the flag `enable-eigen`.

We validated our contributions by benchmarking the simulation times exhibited by the above simulation script, which depicts a typical MIMO 5G NR scenario. To such end, we varied the number

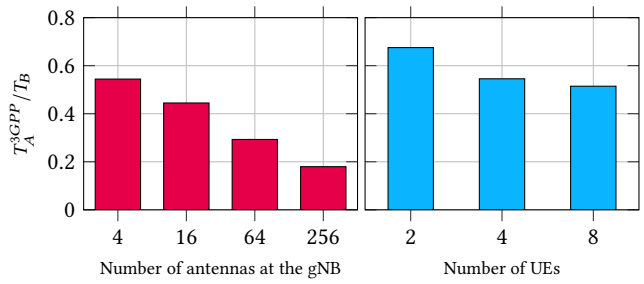

Figure 3: Ratio of the Median Simulation Times After the Merge of this Work with the Eigen Integration (T_A^{3GPP}) and as per ns-3.37 (T_B), when Using the 3GPP Channel Model of [1]

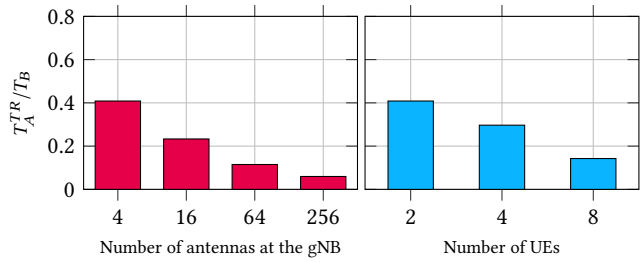

Figure 4: Ratio of the Median Simulation Times Using the Performance-Oriented Channel Model Presented in this Work (T_A^{TR}) and the 3GPP Channel Model of [1] After the Merge of this Work. In this Case, `Eigen` is Disabled

of Next Generation Node Base (gNB) antennas and UEs deployed, and we timed 100 simulation runs for each parameter combination. Figure 3 reports the ratio of the median simulation time achieved when using the `Eigen`-based optimizations, and of the same metric obtained using the vanilla ns-3.37. It can be seen that the matrix multiplication routines offered by `Eigen` can significantly reduce simulation times. For instance, a reduction of 5 times in the simulation time is achieved when equipping gNBs with 256 radiating elements. Similarly, Figure 4 depicts the ratio of the median simulation time obtained by using the FTR-based channel model, and the 3GPP TR 38.901 with `Eigen` disabled. In this case the computational complexity improvement is even more dramatic, with simulations taking as low as 6 % of the time to complete, with respect to the 3GPP model implementation of [47]. As a reference, the median simulation time obtained on an Intel© i5-6700 processor system, before the merge of this work and for $\{2, 4, 8\}$ users is $\{64.7, 210.5, 666.6\}$ [s], respectively.

Finally, we also computed (using the same simulation script, i.e., `cttc-nr-demo-two-ray`) the SINR statistics achieved by the proposed FTR-based model, and compared them to those obtained using the model of [47]. As can be seen in Figure 5, the two models provide similar results. Indeed, a non-negligible difference can be found only in the case of the `InH-OfficeMixed` propagation scenario.

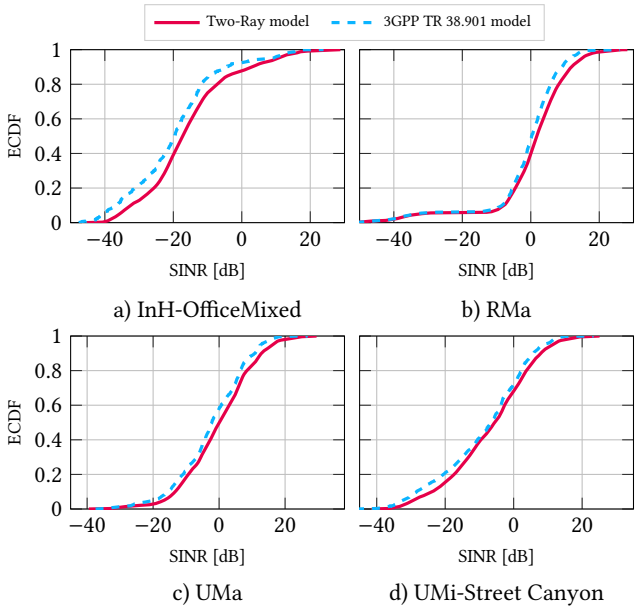

Figure 5: ECDF of the SINR Obtained Using the 3GPP Channel Model of [1], and the Performance-Oriented Channel Model Presented in this Work, for Different Propagation Scenarios

We remark that all the results presented in this section can be reproduced by using the SEM [27] scripts which we provide[2].

5.2 Use Cases

The main goal of both the performance oriented channel model and the optimizations to the 3GPP TR 38.901 model is to enable system-level simulations of large-scale MIMO scenarios for which the implementation of [47] exhibits prohibitive computational complexity. Specifically, our contributions allow ns-3 users to simulate wireless deployments where the devices feature antenna arrays with more than hundreds of radiating elements, and/or the number of communication endpoints is particularly high. For example, the modifications presented in this work can be used in the NR and `mmwave` [29] modules (which both already support the proposed channel models) to simulate massive MIMO 5G NR networks. Notably, a preliminary version of the `Eigen` port has been used in conjunction with the `mmwave` [29] module to simulate 5G networks aided by Intelligent Reflective Surfaces (IRSs), i.e., devices which feature up to 100×100 reflecting elements [31].

Moreover, since the supported frequency range is $0.5 - 100$ GHz, this encompasses not only terrestrial 5G and Long Term Evolution (LTE) deployments, but also most non-terrestrial networks and IEEE Radio Access Technologies (RATs). Finally, the proposed `TwoRaySpectrumPropagationLossModel` can be further extended to support frequencies above 100 GHz using reference fading and path loss statistics.

[2] https://gitlab.com/pagmatt/ns-3-dev/-/tree/gsoc-wns3

6 CONCLUSIONS AND FUTURE WORK

In this paper, we presented a set of optimizations concerning the simulation of MIMO wireless channels in ns-3. First, we introduced the support for the linear algebra library `Eigen` in ns-3, and reduced the computational complexity of the channel matrix generation procedure by avoiding the unnecessary repetition of trigonometric evaluations. Then, we designed and implemented in ns-3 a performance-oriented statistical channel model based on the FTR fading model, which further reduces the simulation time of MIMO scenarios.

Profiling results showed that, thanks to this work, the simulation of MIMO deployments in ns-3 using the 3GPP TR 38.901 channel model takes as little as 20% of the original time. Furthermore, whenever the complexity of the simulations represents a major bottleneck, ns-3 users are now given the possibility of using an additional auxiliary channel model, which achieves a further reduction in simulation time, at the cost of a negligible accuracy penalty with respect to the full 3GPP TR 38.901 model.

As part of our future work, we plan to study more refined beamforming gain correction factors, using the 3GPP statistical channel model as a reference, and possibly making the estimation of such term scenario-dependent. Moreover, we envision to design more efficient storage/access data structures and linear algebra operations for 3D matrices, by better leveraging `Eigen` also in this context. Finally, we will consider using Single instruction, multiple data (SIMD) for speeding up the evaluation of trigonometric functions, and caching the beamforming gain in the `TwoRaySpectrumPropagationLossModel` class to further reduce the simulation time of MIMO scenarios in ns-3.

ACKNOWLEDGMENTS

The work of Matteo Pagin was partially funded by the Google Summer of Code 2022 program. The work of Michele Polese is partially supported by the U.S. NSF Grant CNS 2225590. CTTC authors have received funding from Grant PID2021-126431OB-I00 funded by MCIN/AEI/10.13039/501100011033 and "ERDF A way of making Europe", and TSI-063000-2021-56/57 6G-BLUR project by the Spanish Government. This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on "Telecommunications of the Future" (PE0000001 - program "RESTART"). Moreover, the authors would like to thank Tom Henderson, Eduardo Almeida, and Gabriel Ferreira for their useful suggestions and support during this work.

REFERENCES

[1] 3GPP. 2019. Study on Channel Model for Frequencies from 0.5 to 100 GHz. TR 38.901 (Rel. 15), V15.0.0.
[2] Theodore W Anderson and Donald A Darling. 1954. A Test of Goodness of Fit. *Journal of the American statistical association* 49, 268 (1954), 765–769.
[3] Jeffrey G. Andrews, Tianyang Bai, Mandar N. Kulkarni, Ahmed Alkhateeb, Abhishek K. Gupta, and Robert W. Heath. 2017. Modeling and Analyzing Millimeter Wave Cellular Systems. *IEEE Transactions on Communications* 65, 1 (Jan. 2017), 403–430.
[4] Anuraag Bodi, Steve Blandino, Neeraj Varshney, Jiayi Zhang, Tanguy Ropitault, Mattia Lecci, Paolo Testolina, Jian Wang, Chiehping Lai, and Camillo Gentile. 2021. The NIST Q-D Channel Realization Software.
[5] Amir Ashtari Gargari, Andrea Ortiz, Matteo Pagin, Anja Klein, Matthias Hollick, Michele Zorzi, and Arash Asadi. 2023. Safehaul: Risk-Averse Learning for Reliable mmWave Self-Backhauling in 6G Networks. In *IEEE Conference on Computer Communications (INFOCOM)*. New York, USA.
[6] Henrik Asplund, David Astely, Peter von Butovitsch, Thomas Chapman, Mattias Frenne, Farshid Ghasemzadeh, Måns Hagström, Billy Hogan, George Jöngren, Jonas Karlsson, Fredric Kronestedt, and Erik Larsson. 2020. Chapter 4 - Antenna Arrays and Classical Beamforming. In *Advanced Antenna Systems for 5G Network Deployments*. Academic Press, 89–132.
[7] Dmitry Chizhik, Gerald J. Foschini, and Reinaldo A. Valenzuela. 2000. Capacities of Multi-element Transmit and Receive Antennas: Correlations and Keyholes. *Electronics Letters* 36, 13 (Jun. 2000), 1.
[8] Simon L Cotton. 2014. Human Body Shadowing in Cellular Device-to-Device Communications: Channel Modeling Using the Shadowed κ-μ Fading Model. *IEEE Journal on Selected areas in Communications* 33, 1 (Nov. 2014), 111–119.
[9] Amitabha Ghosh, Andreas Maeder, Matthew Baker, and Devaki Chandramouli. 2019. 5G Evolution: A View on 5G Cellular Technology Beyond 3GPP Release 15. *IEEE Access* 7 (Sep. 2019), 127639–127651.
[10] Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. http://eigen.tuxfamily.org.
[11] Chong Han and Yi Chen. 2018. Propagation Modeling for Wireless Communications in the Terahertz Band. *IEEE Communications Magazine* 56, 6 (Jun. 2018), 96–101.
[12] Ibrahim A. Hemadeh, Katla Satyanarayana, Mohammed El-Hajjar, and Lajos Hanzo. 2017. Millimeter-Wave Communications: Physical Channel Models, Design Considerations, Antenna Constructions and Link-Budget. *IEEE Communications Surveys Tutorials* 20, 99 (Dec. 2017), 870–913.
[13] Yuqiang Heng, Jeffrey G Andrews, Jianhua Mo, Vutha Va, Anum Ali, Boon Loong Ng, and Jianzhong Charlie Zhang. 2021. Six Key Challenges for Beam Management in 5.5G and 6G Systems. *IEEE Communications Magazine* 59, 7 (Jul. 2021), 74–79.
[14] Jakob Hoydis, Fayçal Ait Aoudia, Alvaro Valcarce, and Harish Viswanathan. 2021. Toward a 6G AI-Native Air Interface. *IEEE Communications Magazine* 59, 5 (May 2021), 76–81.
[15] ITU-R. 2015. *IMT Vision - Framework and Overall Objectives of the Future Development of IMT for 2020 and Beyond*. Report M.2083.
[16] ITU-R. 2022. *Future Technology Trends of Terrestrial International Mobile Telecommunications Systems Towards 2030 and Beyond*. Report M.2516-0.
[17] Sian Jin, Sumit Roy, and Thomas R Henderson. 2021. EESM-log-AR: an Efficient Error Model for OFDM MIMO Systems Over Time-Varying Channels. In *Proceedings of the 2021 Workshop on ns-3*. Association for Computing Machinery, Virtual Event, USA.
[18] Sian Jin, Sumit Roy, and Thomas R Henderson. 2021. Efficient PHY Layer Abstraction for Fast Simulations in Complex System Environments. *IEEE Transactions on Communications* 69, 8 (May 2021), 5649–5660.
[19] Sian Jin, Sumit Roy, Weihua Jiang, and Thomas R Henderson. 2020. Efficient Abstractions for Implementing TGn Channel and OFDM-MIMO Links in ns-3. In *Proceedings of the 2020 Workshop on ns-3*. Association for Computing Machinery, Gaithersburg, MD, USA.
[20] Josep Miquel Jornet and Ian F Akyildiz. 2011. Channel Modeling and Capacity Analysis for Electromagnetic Wireless Nanonetworks in the Terahertz Band. *IEEE Transactions on Wireless Communications* 10, 10 (Aug. 2011), 3211–3221.
[21] Mandar N Kulkarni, Eugene Visotsky, and Jeffrey G Andrews. 2018. Correction Factor for Analysis of MIMO Wireless Networks with Highly Directional Beamforming. *IEEE Wireless Communications Letters* 7, 5 (Mar. 2018), 756–759.
[22] Andrea Lacava, Michele Polese, Rajarajan Sivaraj, Rahul Soundrarajan, Bhawani Shanker Bhati, Tarunjeet Singh, Tommaso Zugno, Francesca Cuomo, and Tommaso Melodia. 2022. Programmable and Customized Intelligence for Traffic Steering in 5G Networks using Open RAN Architectures. *arXiv preprint arXiv:2209.14171* (2022).
[23] Sandra Lagen, Kevin Wanuga, Hussain Elkotby, Sanjay Goyal, Natale Patriciello, and Lorenza Giupponi. 2020. New Radio Physical Layer Abstraction for System-Level Simulations of 5G Networks. In *International Conference on Communications (ICC)*. IEEE, Virtual Event.
[24] Mattia Lecci, Paolo Testolina, Marco Giordani, Michele Polese, Tanguy Ropitault, Camillo Gentile, Neeraj Varshney, Anuraag Bodi, and Michele Zorzi. 2020. Simplified Ray Tracing for the Millimeter Wave Channel: A Performance Evaluation. In *Information Theory and Applications Workshop (ITA)*. IEEE, San Diego, CA USA.
[25] Shanyun Liu, Xianbin Yu, Rongbin Guo, Yajie Tang, and Zhifeng Zhao. 2021. THz Channel Modeling: Consolidating the Road to THz Communications. *China Communications* 18, 5 (May 2021), 33–49.
[26] Yuchen Liu, Shelby K Crisp, and Douglas M Blough. 2021. Performance Study of Statistical and Deterministic Channel Models for mmWave Wi-Fi Networks in ns-3. In *Proceedings of the 2021 Workshop on ns-3*. Association for Computing Machinery, Virtual Event, USA.
[27] Davide Magrin, Dizhi Zhou, and Michele Zorzi. 2019. A Simulation Execution Manager for ns-3: Encouraging Reproducibility and Simplifying Statistical Analysis of ns-3 Simulations. In *Proceedings of the 22nd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*. Miami Beach, FL USA.

[28] Theodoros Mavridis, Luca Petrillo, Julien Sarrazin, Aziz Benlarbi-Delai, and Philippe De Doncker. 2015. Near-Body Shadowing Analysis at 60 GHz. *IEEE Transactions on Antennas and Propagation* 63, 10 (Jul. 2015), 4505–4511.

[29] Marco Mezzavilla, Menglei Zhang, Michele Polese, Russell Ford, Sourjya Dutta, Sundeep Rangan, and Michele Zorzi. 2018. End-to-End Simulation of 5G mmWave Networks. *IEEE Communications Surveys & Tutorials* 20, 3 (Apr. 2018), 2237–2263.

[30] Matteo Pagin, Marco Giordani, Amir Ashtari Gargari, Alberto Rech, Federico Moretto, Stefano Tomasin, Jonathan Gambini, and Michele Zorzi. 2022. End-to-End Simulation of 5G Networks Assisted by IRS and AF Relays. In *Proc. IEEE MedComNet*. Paphos, Cyprus.

[31] Matteo Pagin, Tommaso Zugno, Michele Polese, and Michele Zorzi. 2022. Resource Management for 5G NR Integrated Access and Backhaul: A Semi-Centralized Approach. *IEEE Transactions on Wireless Communications* 21, 2 (Jul. 2022), 753–767.

[32] Natale Patriciello, Sandra Lagen, Biljana Bojovic, and Lorenza Giupponi. 2019. An E2E Simulator for 5G NR Networks. *Simulation Modelling Practice and Theory* 96 (Nov. 2019), 101933.

[33] Michele Polese, Leonardo Bonati, Salvatore D'Oro, Stefano Basagni, and Tommaso Melodia. 2022. ColO-RAN: Developing Machine Learning-Based xApps for Open RAN Closed-Loop Control on Programmable Experimental Platforms. *IEEE Transactions on Mobile Computing* (Jul. 2022), 1–14.

[34] Michele Polese, Leonardo Bonati, Salvatore D'Oro, Stefano Basagni, and Tommaso Melodia. 2023. Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges. *IEEE Communications Surveys & Tutorials* (Jan. 2023), Early Access.

[35] Michele Polese, Marco Giordani, Marco Mezzavilla, Sundeep Rangan, and Michele Zorzi. 2021. 6G Enabling Technologies. In *6G mobile wireless networks*. Springer, 25–41.

[36] Michele Polese, Josep Miquel Jornet, Tommaso Melodia, and Michele Zorzi. 2020. Toward End-to-End, Full-Stack 6G Terahertz Networks. *IEEE Communications Magazine* 58, 11 (Nov. 2020), 48–54.

[37] Michele Polese and Michele Zorzi. 2018. 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)*. Kalamata, Greece.

[38] Andrea Ramos, Yanet Estrada, Miguel Cantero, Jaime Romero, David Martín-Sacristán, Saúl Inca, Manuel Fuentes, and José Monserrat. 2022. Implementation and Calibration of the 3GPP Industrial Channel Model for ns-3. In *Proceedings of the 2022 Workshop on ns-3*. Association for Computing Machinery, Virtual Event, USA.

[39] Theodore S. Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N. Wong, Jocelyn K. Schulz, Mathew Samimi, and Felix Gutierrez. 2013. Millimeter Wave Mobile Communications for 5G Cellular: It Will Work! *IEEE Access* 1 (May 2013), 335–349.

[40] Mattia Rebato, Laura Resteghini, Christian Mazzucco, and Michele Zorzi. 2018. Study of Realistic Antenna Patterns in 5G mmWave Cellular Scenarios. In *Proc. IEEE ICC*. Kansas City, MO, USA.

[41] Juan M. Romero-Jerez, F. Javier Lopez-Martinez, José F. Paris, and Andrea J. Goldsmith. 2017. The Fluctuating Two-Ray Fading Model: Statistical Characterization and Performance Analysis. *IEEE Transactions on Wireless Communications* 16, 7 (May 2017), 4420–4432.

[42] Mansoor Shafi, Jianhua Zhang, Harsh Tataria, Andreas F Molisch, Shu Sun, Theodore S Rappaport, Fredrik Tufvesson, Shangbin Wu, and Koshiro Kitao. 2018. Microwave vs. Millimeter-Wave Propagation Channels: Key Differences and Impact on 5G Cellular Systems. *IEEE Communications Magazine* 56, 12 (Dec. 2018), 14–20.

[43] Paolo Testolina, Mattia Lecci, Michele Polese, Marco Giordani, and Michele Zorzi. 2020. Scalable and Accurate Modeling of the Millimeter Wave Channel. In *International Conference on Computing, Networking and Communications (ICNC)*. IEEE, Big Island, Hawaii, USA.

[44] Michel Daoud Yacoub. 2007. The κ-μ Distribution and the η-μ Distribution. *IEEE Antennas and Propagation Magazine* 49, 1 (Feb. 2007), 68–81.

[45] Menglei Zhang, Michele Polese, Marco Mezzavilla, Jing Zhu, Sundeep Rangan, Shivendra Panwar, and Michele Zorzi. 2019. Will TCP Work in mmWave 5G Cellular Networks? *IEEE Communications Magazine* 57, 1 (Jan. 2019), 65–71.

[46] Tommaso Zugno, Matteo Drago, Sandra Lagén, Zoraze Ali, and Michele Zorzi. 2021. Extending the ns-3 Spatial Channel Model for Vehicular Scenarios. In *Proceedings of the 2021 Workshop on ns-3*. Association for Computing Machinery, Virtual Event, USA.

[47] Tommaso Zugno, Michele Polese, Natale Patriciello, Biljana Bojović, Sandra Lagén, and Michele Zorzi. 2020. Implementation of a Spatial Channel Model for ns-3. In *Proceedings of the 2020 Workshop on ns-3*. Association for Computing Machinery, Gaithersburg, MD, USA.