engineering proceedings

MDPI

Proceeding Paper

Cross-Layer Evaluation of Generalized Fading Models in Simu5G-Based Vehicular Downlink Scenarios †

Ricardo Mena Villacis [1,*], Patricio Flor [1], Marco Serrano [1], Edison Tatayo [1] and Julian Simbaña [2]

[1] Departamento de Electrónica, Telecomunicaciones y Redes de Información (DETRI), Escuela Politécnica Nacional (EPN), Ladrón de Guevara E11-253, Quito 170517, Ecuador; patricio.flor@epn.edu.ec (P.F.); marco.serrano@epn.edu.ec (M.S.); edison.tatayo@epn.edu.ec (E.T.)

[2] Departamento de Formación Básica (DFB), Escuela Politécnica Nacional (EPN), Ladrón de Guevara E11-253, Quito 170517, Ecuador; julian.simbana@epn.edu.ec

* Correspondence: ricardo.menav@epn.edu.ec

† Presented at the XXXIV Conference on Electrical and Electronic Engineering (JIEE 2026), Quito, Ecuador, 26–30 October 2026.

Abstract

System-level fifth-generation (5G) vehicle-to-everything (V2X) studies can conceal physical-layer changes when application traffic does not saturate the radio link. This paper evaluates five small-scale fading profiles—no fading, Rayleigh, Rician, Nakagami-m, and κ-μ—for 5G New Radio (NR) in a Simu5G framework coupled with Simulation of Urban MObility (SUMO) and Vehicular Network Simulation (Veins). The 45-configuration matrix spans urban, rural, and highway environments and three traffic levels. The implementation follows Third Generation Partnership Project (3GPP)-style propagation settings and evaluates the signal-to-interference-plus-noise ratio (SINR), channel quality indicator (CQI), and delivered throughput. Rayleigh fading reduces median downlink SINR by 5.0–6.0 dB and CQI by 2–3 index points relative to the no-fading baseline in all nine environment–traffic conditions, while Rician, Nakagami-m, and κ-μ remain closer to the baseline. Median per-vehicle throughput remains within 0.046–0.056 Mbps, with a maximum within-condition spread of 0.0045 Mbps. A standardized five-feature clustering analysis identifies five operating regimes organized primarily by environment and traffic scale; the first two principal components retain 97.6% of the multimetric variance. The combined metrics achieve a clear cross-layer characterization: channel-quality indicators distinguish fading profiles even when the configured low-rate service preserves delivered application throughput.

Keywords: 5G NR; C-V2X; Simu5G; vehicular mobility; small-scale fading; κ-μ; Nakagami-m; SINR; CQI

1. Introduction

Vehicle-to-everything (V2X) communication combines rapidly changing geometry, multipath propagation, and heterogeneous traffic demand. System-level conclusions therefore depend on both the channel representation and the layer at which performance is measured. The Third Generation Partnership Project (3GPP) channel framework specifies large-scale and spatially consistent propagation models for frequencies from 0.5 to 100 GHz [1], while link-level fading is commonly represented through statistical models. Rayleigh fading describes a scattered channel without a dominant component [2]; Rician, Nakagami-m, and the generalized κ-μ distribution provide progressively different control of dominant power and cluster structure [3–5].

Check for updates

Published: 2 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Simu5G supports end-to-end fifth-generation (5G) New Radio (NR) evaluation in the Objective Modular Network Testbed in C++ (OMNeT++) [6]. Simulation of Urban MObility (SUMO) supplies microscopic road mobility [7], and Vehicular Network Simulation (Veins) connects the traffic and network simulators through the Traffic Control Interface (TraCI) [8]. This toolchain makes it possible to observe physical-layer indicators and delivered traffic under the same mobility trace. An application stream below radio capacity can maintain its offered rate while the channel margin changes, making the joint observation of radio and application metrics especially valuable.

This work contributes: (i) a configured Simu5G channel-model extension that evaluates Rician, Nakagami-m, and κ-μ profiles alongside Rayleigh and a no-fading reference; (ii) a 45-configuration comparison across three road environments and three traffic levels; and (iii) a cross-layer analysis that keeps median downlink signal-to-interference-plus-noise ratio (SINR), channel quality indicator (CQI), aggregate throughput, and per-vehicle throughput as distinct quantities. Building on these established tools and models, this evaluation-oriented contribution provides reproducible matched comparisons, distinguishes the cross-layer roles of the metrics, and traces reported claims to preserved comparative tables.

Research Gap and Evaluation Strategy

Vehicular network studies often compare propagation models through a single physical-layer indicator or report application performance without retaining the channel context. These views answer different questions. SINR represents the instantaneous link margin after propagation, interference, and noise; CQI translates radio conditions into a discrete indicator used by link adaptation; aggregate throughput reflects the traffic served by the cell; and per-vehicle throughput measures the service delivered to individual mobile users. A propagation profile can therefore produce a clear SINR and CQI separation while the offered traffic remains fully served. Treating these metrics as interchangeable can hide the operating mechanism that produced the observed application rate.

The present evaluation addresses this gap through two complementary analyses. First, it compares the five fading profiles within each matched environment–traffic condition, preserving the same road geometry, mobility trace, vehicle population, and radio configuration. This controlled comparison reveals how the selected small-scale distribution changes channel-quality and service-delivery metrics. Second, it applies a multimetric clustering procedure across all 45 configurations. The clustering jointly considers radio quality, cell-level load, individual service rate, and vehicle population, making it possible to identify operating regimes that are not defined in advance by the scenario labels. Principal component analysis (PCA) is then used only to display those groups in two dimensions.

This combined strategy extends the evaluation beyond a ranking of fading models. The within-condition analysis isolates profile-dependent behavior, whereas the unsupervised analysis reveals how mobility environment and offered load organize the complete performance space. Together, they provide a reproducible basis for selecting channel indicators, designing representative simulation campaigns, and interpreting application-layer performance in 5G vehicular studies.

2. Simulation Framework and Data Audit

2.1. Channel and Mobility Configuration

The evaluated downlink combines the extended Simu5G channel path with SUMO mobility and Veins synchronization. The large-scale component follows the configured 3GPP-style path-loss and shadowing settings, while the received power is multiplied by a stochastic small-scale gain. Five profiles are compared: a unit-gain no-fading reference,

https://doi.org/10.3390/engproc2026158004

Rayleigh, Rician with $K = 6$ dB, Nakagami-m with $m = 2$, and κ-μ with $\kappa = 2$ and $\mu = 1.5$. They are fixed operating points selected for reproducible matched comparisons, rather than universal fits; sensitivity across parameter ranges and independent mobility realizations is a constructive future improvement. To make the selection transparent, $K = 6$ dB introduces a dominant-component contrast to Rayleigh; for Nakagami-m, $m = 0.5$ is very severe, $m = 1$ recovers Rayleigh, $m = 2$ is moderate, and $m = 3$–5 is mild. For κ-μ, $\kappa \approx 0$ has no strong dominant component, $\kappa = 1$–3 is moderate, $\kappa > 5$ is strong, and $\mu = 1$–3 provides representative multicluster scenarios. Accordingly, $m = 2$, $\kappa = 2$, and $\mu = 1.5$ form a moderate, distinguishable comparison point.

The common radio configuration uses a 2.0 GHz carrier, 10 MHz bandwidth (50 resource blocks), 46 dBm next-generation NodeB (gNB) transmit power, and 100 s executions. Downlink G.711 periodic traffic offers approximately 54 kbps per vehicle. Urban mobility uses a grid, speeds of 30–50 km/h, and a maximum modeled separation of approximately 1.65 km. Rural and highway corridors extend to approximately 4 km, with speeds of 90–120 and 100–130 km/h, respectively. Shadowing is disabled so that the comparative matrix isolates the selected small-scale fading profile. The controlled configuration isolates profile-dependent comparisons; consequently, its quantitative findings apply to the evaluated 2.0 GHz, 10 MHz, low-rate traffic, 100 s, and shadowing-disabled settings, without representing all deployment conditions. Table 1 reports the vehicle counts recovered from the result files.

Table 1. Experimental matrix of environments, road representations, and vehicle populations.

Environment	Road Representation	Light	Medium	Dense
Urban	Grid, 1.65 km	30	60	100
Rural	Corridor, 4 km	20	50	80
Highway	Multi-lane corridor, 4 km	30	70	120

2.2. Metrics, Software, and Reproducibility

The scenario generator writes the mobility files, the Network Description (NED) language topology, and the initialization (INI) configuration blocks used by OMNeT++. The topology contains two gNBs, a user-plane function, a router, an application server, and vehicles created dynamically through Veins. Within each environment–traffic condition, the same SUMO mobility seed and trace are reused for all five fading profiles. OMNeT++ run index 0 identifies the analyzed execution. Table 2 consolidates the radio, traffic, software, and execution settings.

Four metrics are kept separate. Downlink SINR is the median of `measuredSinrDl` samples in dB; CQI is the median of `averageCqiDl`; aggregate throughput is the median of 1-s execution-level bins; and per-vehicle throughput is received bits divided by each vehicle's active interval, summarized by the median across valid vehicles. Simulation vectors were processed into graphical user interface (GUI) statistics and cross-profile comparative tables. The run identifiers agree across all 45 configurations, and the comparative tables provide one consistent canonical source for every numerical value used below.

https://doi.org/10.3390/engproc2026158004

Table 2. Radio, traffic, mobility, and software configuration.

Parameter	Configured Value
Radio link	2.0 GHz carrier; 10 MHz bandwidth; 50 resource blocks; 46 dBm gNB transmit power
Fading profiles	No fading; Rayleigh; Rician ($K = 6$ dB); Nakagami-m ($m = 2$); κ-μ ($\kappa = 2$, $\mu = 1.5$)
Traffic and scheduler	Downlink G.711 periodic stream, approximately 54 kbps per vehicle; proportional-fair scheduling
Mobility	SUMO-generated grid or bidirectional corridor; 30–50 km/h urban, 90–120 km/h rural, 100–130 km/h highway
Execution control	100 s; OMNeT++ run index 0; scenario-specific mobility seed shared across fading profiles
Simulation stack	OMNeT++ 6.0; INET 4.4; Veins 5.2; SUMO 1.11; MATLAB R2024b; Python 2.7.18
Simu5G source	Locally extended Simu5G checkout; source commit not recorded in the archived result package
Build platform	Linux workspace (/home/simu5g); exact C++ compiler identifier not recorded in the archived result package

2.3. Multimetric Analysis

The distributional throughput comparison uses the vehicle-level percentiles preserved for every configuration. Each box spans the 25th–75th percentiles, its center line is the median, and its whiskers mark the 5th and 95th percentiles. This construction preserves within-configuration variability without reconstructing or pooling unavailable individual observations. The same nine environment–traffic groups and five fading colors used in the channel comparison are retained.

For the unsupervised analysis, each configuration is represented by five features: median downlink SINR, median CQI, median aggregate downlink throughput, median per-vehicle downlink throughput, and the number of valid vehicles. Each feature is standardized to zero mean and unit variance before clustering, preventing units or numerical scale from controlling the distance calculation. Deterministic k-means uses seed 2026 and 100 initializations for every candidate k from 2 to 5. The number of groups is selected by the highest mean silhouette coefficient. PCA is calculated from the standardized feature matrix and is used only for the two-dimensional display; cluster membership is obtained in the full five-feature space.

3. Results and Discussion

3.1. Physical-Layer Separation

Figure 1 presents the configuration medians. Its labels combine environment (U: urban; R: rural; H: highway) and traffic level (L: light; M: medium; D: dense); each point is a configuration median and the lines support visual comparison. Urban median SINR lies between 19.4 and 26.1 dB, whereas rural and highway medians span 5.9–13.7 dB. The resulting separation characterizes the combined effect of the configured path loss, geometry, distance, and mobility.

The profile ordering is consistent across the nine environment–traffic conditions. Rayleigh median SINR is 5.0–6.0 dB below the no-fading reference (mean gap 5.3 dB across configurations), and its median CQI is 2–3 points lower. Rician is 0.4–1.2 dB below the reference, Nakagami-m is 0.8–1.5 dB below, and κ-μ differs by −0.1 to 1.0 dB. The matrix therefore establishes Rayleigh as the most severe evaluated profile, while Rician, Nakagami-m, and κ-μ form a closer intermediate group.

Traffic-level changes are smaller than the separation produced by the fading profiles. The design therefore achieves its main comparative objective: it reveals stable cross-profile differences across three distinct mobility environments while preserving the application service rate.

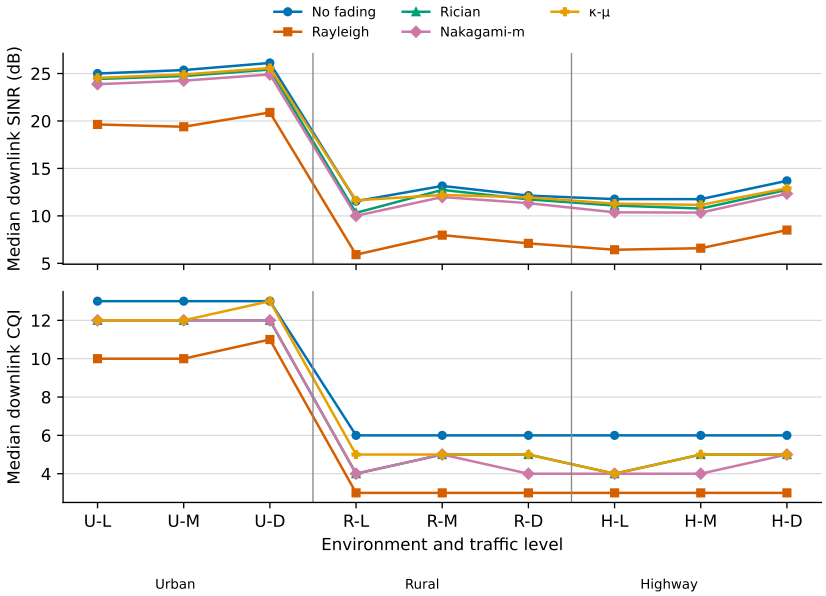


Figure 1. Median downlink SINR and CQI across environments, traffic levels, and fading profiles.

3.2. Application Throughput and Cross-Layer Interpretation

Figure 2 provides the application-layer distribution for every configuration. Its environment–traffic abbreviations follow Figure 1. Boxes represent the interquartile interval, center lines represent medians, and whiskers extend from the 5th to the 95th percentile of the valid vehicle rates. Across all 45 configurations, medians occupy the narrow 0.046–0.056 Mbps interval. The largest spread among fading-profile medians within the same environment–traffic condition is 0.0045 Mbps. Urban executions are centered higher (0.053–0.056 Mbps) than rural and highway executions (0.046–0.051 Mbps), and the five profiles remain closely grouped within each condition.

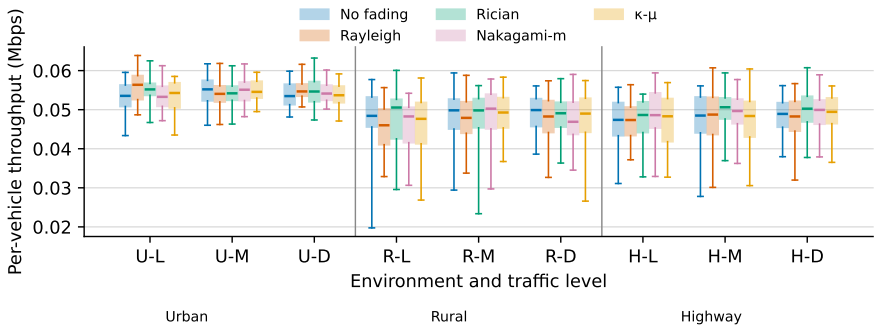


Figure 2. Per-vehicle downlink throughput distributions across environments, traffic levels, and fading profiles.

This narrow median range is consistent with an offered rate near 54 kbps per vehicle and a non-saturated radio link. The distributions add a second result: urban interquartile intervals are compact and concentrated near the offered rate, whereas rural and highway cases exhibit longer lower whiskers, reaching approximately 0.02–0.04 Mbps in several configurations. The scheduler therefore preserves the central application rate while a subset of vehicle observations experiences greater variability in the extended corridors. Aggregate throughput is a separate execution-level key performance indicator (KPI) that grows with the number of active flows, while the per-vehicle distribution measures service delivery

and fairness across mobile users. Their combination with SINR and CQI provides a more complete cross-layer assessment than either layer alone.

3.3. Multimetric Operating Regimes

The silhouette comparison selected five clusters, with a coefficient of 0.557 for $k = 5$, compared with 0.525, 0.493, and 0.522 for $k = 2$, 3, and 4, respectively. Figure 3 displays the resulting groups after projecting the standardized feature matrix onto its first two principal components. These two axes retain 97.6% of the standardized variance. Cluster colors encode the five operating groups, while circles, squares, and triangles identify urban, rural, and highway configurations.

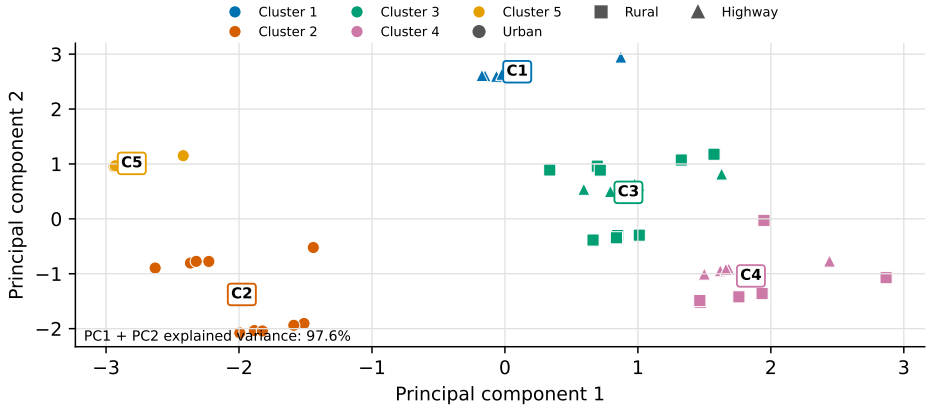


Figure 3. Principal component map of the five multimetric configuration clusters.

All five characteristics used to produce this result are reflected in Table 3: median SINR and CQI represent radio quality; median aggregate throughput captures cell-level delivered load; median per-vehicle throughput represents individual service; and valid-vehicle count represents traffic scale. Because clustering is performed after standardization, each characteristic contributes on a comparable numerical scale. The table reports cluster means in their original units to support direct engineering interpretation.

Table 3. Mean performance characteristics and composition of the five clusters.

Cluster	*n*	SINR (dB)	CQI	Aggregate (Mbps)	Per Vehicle (Mbps)	Vehicles	Configuration Composition
C1	5	12.0	4.8	2.740	0.0494	120.0	Highway–dense
C2	10	23.6	11.8	1.115	0.0546	45.0	Urban–light and urban–medium
C3	14	11.1	4.8	1.510	0.0492	67.9	Rural/highway, mainly medium and dense
C4	11	9.8	4.2	0.653	0.0481	27.3	Rural/highway, mainly light
C5	5	24.6	12.2	2.468	0.0541	100.0	Urban–dense

The clustering separates the configurations along two complementary performance dimensions. C2 and C5 contain the urban cases, combining the highest SINR, CQI, and per-vehicle throughput; their separation is produced primarily by vehicle population and aggregate traffic. C4 collects the lower-load rural and highway cases, while C3 captures their medium-to-dense operating range. C1 isolates highway–dense operation, where 120 valid vehicles raise aggregate throughput to 2.740 Mbps while the per-vehicle median remains near 0.049 Mbps. Each cluster contains multiple fading profiles, showing that environment and traffic scale define the global operating regimes, while the matched comparisons in Figure 1 reveal the systematic within-regime effect of the fading distribution.

3.4. Cross-Layer Design Implications

The combined results establish a metric-selection strategy for vehicular system-level studies. When the objective is to distinguish propagation profiles, median SINR and CQI provide the strongest separation: the Rayleigh gap is present in every matched condition, even where the delivered application rate changes only slightly. When the objective is service assessment, the throughput boxes add information that a single median cannot provide. Their interquartile widths and lower whiskers reveal how consistently the configured rate reaches individual vehicles across the mobility trace. Consequently, channel-model evaluation and service evaluation should be treated as complementary tasks with different primary indicators.

The load dimension requires a second distinction. Aggregate throughput increases with the number of simultaneous application flows and therefore identifies how much traffic the cell serves. Per-vehicle throughput instead remains close to the configured source rate and indicates whether that service is maintained for individual users. The clustering formalizes this relationship: C2 and C5 have similar radio-quality and per-vehicle values, but they separate because urban–dense operation has more vehicles and a larger aggregate rate. Likewise, C1 identifies the highway–dense regime through its 120 vehicles and 2.740 Mbps aggregate median, while its per-vehicle median remains comparable with the other non-urban groups.

Table 4 converts these observations into an evaluation guide. The mapping is useful for simulation planning because it links each research objective to the metric combination that responds directly to it. It also prevents a stable application rate from being interpreted as an unchanged radio channel and prevents aggregate cell traffic from being interpreted as an individual-user rate.

Table 4. Cross-layer metric selection for vehicular simulation objectives.

Evaluation Objective	Primary Evidence	Result-Based Interpretation
Differentiate fading profiles	Median SINR and CQI	Resolves the consistent Rayleigh separation and the closer Rician, Nakagami-m, and κ-μ group.
Assess individual service	Per-vehicle median, P25–P75 box, and P5–P95 whiskers	Quantifies the delivered rate and exposes the wider lower tails observed in rural and highway corridors.
Characterize cell load	Aggregate throughput and valid-vehicle count	Separates traffic scale from the individual service rate and distinguishes dense operating regimes.
Compare complete regimes	Standardized five-feature clustering	Integrates radio quality, cell load, user service, and vehicle population into reproducible operating groups.

This hierarchy supports efficient experiment design. A compact screening campaign can first use SINR and CQI to verify whether a channel extension produces the intended physical-layer separation. Throughput percentiles can then confirm service stability and user-level spread under the target application. Finally, aggregate traffic and multimetric grouping can be used to compare complete environment–load regimes. The approach preserves the physical meaning of each KPI while producing a unified interpretation of the simulated network.

3.5. Reproducible Analysis Workflow

The analysis pipeline imports the 45 comparative records, standardizes environment and profile names, checks the configuration matrix, and calculates all reported ranges and within-condition differences. It generates the channel comparison, percentile-based throughput boxes, and multimetric clustering map from the same canonical table. Median-based processing keeps the evaluated metrics aligned across configurations and provides

direct traceability from each manuscript claim to its source row. The scenario-specific mobility seed is held constant across fading profiles, so each five-profile comparison uses the same road geometry, vehicle population, and mobility realization. The pipeline exports the cluster assignment, original-unit cluster means, silhouette scores for every candidate k, and principal component coordinates, allowing every point and group in Figure 3 to be independently reconstructed.

4. Conclusions

Across the 45-configuration matrix, Rayleigh fading produces the lowest physical-layer medians: its SINR is 5.0–6.0 dB and its CQI 2–3 points below the no-fading reference in every environment–traffic condition. The other three stochastic profiles remain closer to the reference. Per-vehicle throughput medians remain between 0.046 and 0.056 Mbps, with at most 0.0045 Mbps separation among profiles within a condition, while the percentile boxes reveal wider lower tails in the extended rural and highway corridors. The five-feature analysis identifies five operating regimes and separates high-quality urban operation from the load-dependent rural and highway groups. The matched matrix demonstrates profile-dependent SINR and CQI separation, throughput distributions describe user-level service, and clustering relates both to environment and load. This evaluation-oriented evidence can guide reproducible simulation campaigns; broader system designs remain a constructive future extension. The resulting cross-layer comparison is specific to the configured fading parameters, mobility conditions, and offered traffic; alternative parameter choices can change the quantitative gaps, motivating future sensitivity campaigns. Future campaigns can incorporate spatial shadowing, additional bands and bandwidths, saturated and heterogeneous traffic, longer observation windows, and mobility-consistent Doppler–fading coupling.

Author Contributions: Conceptualization, R.M.V. and P.F.; methodology, P.F., E.T. and M.S.; software, P.F. and E.T.; validation, M.S., J.S. and R.M.V.; formal analysis, P.F. and R.M.V.; investigation, P.F. and E.T.; data curation, P.F. and R.M.V.; writing—original draft, P.F., E.T., J.S. and R.M.V.; writing—review and editing, J.S. and R.M.V.; supervision, R.M.V. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by project PISF-25-06, "MODELADO DINÁMICO DEL CANAL 5G Y POSTERIOR PARA LA EVALUACIÓN DEL DESEMPEÑO V2X EN ENTORNOS URBANOS".

Data Availability Statement: The processed data supporting the findings are available from the corresponding author upon reasonable request.

Acknowledgments: The authors acknowledge the support of the Departamento de Electrónica, Telecomunicaciones y Redes de Información at Escuela Politécnica Nacional.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. 3rd Generation Partnership Project. *Study on Channel Model for Frequencies from 0.5 to 100 GHz*; Technical Report TR 38.901; Version 18.1.0; Release 18; 3GPP: Sophia Antipolis, France, 2026.
2. Sklar, B. Rayleigh Fading Channels in Mobile Digital Communication Systems Part I: Characterization. *IEEE Commun. Mag.* **1997**, *35*, 90–100. [CrossRef]
3. Simon, M.K.; Alouini, M.S. *Digital Communication over Fading Channels*, 2nd ed.; John Wiley & Sons: Hoboken, NJ, USA, 2005.
4. Nakagami, M. The m-Distribution—A General Formula of Intensity Distribution of Rapid Fading. In *Statistical Methods in Radio Wave Propagation*; Pergamon Press: Elmsford, NY, USA, 1960; pp. 3–36.
5. Yacoub, M.D. The κ-μ Distribution and the η-μ Distribution. *IEEE Antennas Propag. Mag.* **2007**, *49*, 68–81. [CrossRef]
6. Nardini, G.; Sabella, D.; Stea, G.; Thakkar, P.; Virdis, A. Simu5G: An OMNeT++ Library for End-to-End Performance Evaluation of 5G Networks. *IEEE Access* **2020**, *8*, 181176–181191. [CrossRef]

7. Behrisch, M.; Bieker, L.; Erdmann, J.; Krajzewicz, D. SUMO—Simulation of Urban MObility: An Overview. In Proceedings of the SIMUL 2011, The Third International Conference on Advances in System Simulation, Barcelona, Spain, 23–29 October 2011; pp. 63–68.
8. Veins Developers. Veins: Vehicular Network Simulation Framework. 2024. Available online: https://veins.car2x.org/ (accessed on 26 September 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.