




POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE



EXECUTIVE SUMMARY OF THE THESIS

Rain Attenuation Mitigation for EHF NGSO Systems Using Site and/or Orbital Diversity

TESI MAGISTRALE IN TELECOMMUNICATION ENGINEERING – INGEGNERIA DELLE
TELECOMUNICAZIONI

Author: SOFIA LECCI

Advisor: LORENZO LUINI

Co-Advisor: ENRICO POLO

Academic year: 2024-2025

1. Introduction

The current evolution of satellite communications is characterized by two concurrent trends: the large-scale adoption of Non-Geostationary Orbit (NGSO) constellations and the progressive migration toward higher frequency bands (Ka and, increasingly, Q/V). While these bands enable wider bandwidth and higher throughput, they also expose the Earth-space links to a propagation environment where tropospheric effects become a dominant design constraint. In particular, rain attenuation is the key impairment above approximately 10 GHz, and it can reach tens of dB during intense events, directly threatening link availability if not properly mitigated.

NGSO systems amplify this challenge compared to traditional GEO links because satellite motion causes continuously varying elevation angles and slant-path lengths, with long time spent at low elevations. Static fade margins become inefficient or impractical, and system design must leverage dedicated Fade Mitigation Techniques (FMTs). This thesis focuses on spatial diversity strategies.

Specifically, Site Diversity (SD) mitigates fades by using multiple geographically separated gateways, whereas Orbital Diversity (OD) exploits the simultaneous visibility of multiple satellites from a single site and selects, at each time, the least attenuated link. In realistic architectures, SD and OD can also be combined, potentially delivering the largest improvements in the high-attenuation tail that drives high-availability design. The main objective of this thesis is to quantify and compare the performance of SD, OD, and joint SD+OD for EHF NGSO systems under realistic propagation conditions. In addition, a statistical model for OD is proposed linking attenuation differences to the angular separation between simultaneous links, enabling faster performance prediction without time-domain simulation.

2. Simulation Framework

2.1 Scenario Definition

The simulation framework combines realistic space-time rain fields generated through ST-MultiEXCELL with detailed NGSO orbital

geometry. Precipitation is modeled over a $200 \text{ km} \times 200 \text{ km}$ domain with 1 km spatial resolution, allowing pixel-level integration of specific attenuation along dynamically varying slant paths.

The NGSO constellation considered is the SpaceX VLEO sub-constellation operating at 40 GHz, composed of 7,518 satellites distributed over three orbital shells at altitudes around 340 km. Only satellites with elevation angle $\geq 5^\circ$ are considered visible, and the simulator enforces that all ground stations have access to the same number of simultaneously visible satellites at each epoch, ensuring not trivial differences in diversity order. Ground-station geometries, shown in Figure 1, are parameterized as line, triangle, and square configurations with variable inter-site separation S . The geometry is defined in a local East-North-Up (ENU) frame and mapped consistently onto the rain-field grid. For each station and each visible satellite, attenuation time series are produced and stored in a structured way to allow long-term post-processing. The performance assessment is then carried out through Complementary Cumulative Distribution Functions (CCDFs) of attenuation and through diversity gains, which quantify how much the effective attenuation is reduced at a given exceedance probability compared to a single-link reference.




Figure 1: Ground station configurations.

2.2 Orbital-Propagation Module

An auxiliary orbital-propagation module is developed as a standalone component to preserve the modularity and extensibility of the overall framework. Although it is not directly used in the simulations (which rely on SES-provided NGSO ephemerides), the module enables future studies based on fully synthetic or hybrid constellations by allowing the generation and propagation of customizable LEO orbits starting from user-defined Two-Line Element (TLE) files. In the presented example, the TLE is configured to emulate a nearly circular LEO satellite (eccentricity 0.0005) with inclination 50° , RAAN 100° , and mean motion 15 rev/day (≈ 96 min period), representative

of broadband constellations at $\sim 550\text{--}600$ km altitude. From these orbital elements, the satellite state vector is computed at the epoch and propagated over one year with 60 s resolution under a simplified two-body gravitational model. The resulting trajectory is expressed in the local ENU frame of a fixed ground station to derive time series of azimuth and elevation angles. A 5° elevation mask is applied to identify visibility, and the collected visible samples are further used to estimate empirical PDFs of azimuth and elevation. These distributions highlight typical LEO visibility features—strong skew toward low elevations and bimodal azimuth patterns linked to ascending/descending passes—providing useful geometric insight and supporting future integration of synthetic orbital datasets into the rain-attenuation analysis chain

3. Performance Analysis and Results

3.1 Orbital Diversity

OD is evaluated by considering the set of satellites simultaneously visible from a single ground station and selecting, at each epoch, the satellite yielding the minimum attenuation. Results confirm that OD improves monotonically with increasing diversity order, i.e., with the number of simultaneously available satellites considered by the selection algorithm. Nonetheless, OD has intrinsic limitations: since all candidate links originate from the same site, they share local meteorological conditions and an initial common path segment, so full decorrelation is not achievable, especially for small angular separations. To quantify how performance scales with the number of available alternatives, the selection is repeated for different candidate-set sizes N_{vis} . At each epoch, the visible-satellite set is uniformly subsampled to retain exactly N_{vis} links and the minimum attenuation is selected within that subset. As N_{vis} increases, the gateway can exploit a richer set of link geometries; due to partial decorrelation of rain attenuation across paths, the probability that all candidates are simultaneously affected by deep fades decreases, and the OD-selected attenuation $A_{\text{OD}}(t)$ more frequently approaches clear-sky conditions, as visible in Figure 3.1. To mitigate sensitivity to the specific rain-field realization, the analysis is repeated for four different ground-station placements within

the rain map and the final statistics are obtained by averaging the corresponding CCDFs:

$$\hat{P}_j(A_{OD} > a) = \frac{1}{T} \sum_{n=1}^T \mathbf{1}\{A_{OD}^j(t_n) > a\}, \quad j \in \{1,2,3,4\} \quad (1)$$

where T is the number of epochs and $\mathbf{1}\{\cdot\}$ is the indicator function. From an operational standpoint, OD also translates into a significant reduction in outage probability: for an outage threshold of $A_{OD} = 30$ dB, the exceedance probability decreases from $\bar{P}(A_{OD} > 30 \text{ dB}) \approx 3.29\%$ in the single-link case to $\bar{P}(A_{OD} > 30 \text{ dB}) \approx 0.052\%$ when selecting over the full visible set (~ 170 satellites), i.e., an outage reduction by a factor of about 64x.




Figure 2.1: OD CCDFs for different candidate-set sizes.

3.2 Site Diversity

Site Diversity is assessed by deploying multiple gateways within the same precipitation field and selecting, at each epoch, the site with minimum attenuation. In this case, performance is governed primarily by the inter-site distance and by the number of sites. Increasing separation generally reduces the correlation between rain attenuation processes, since intense rain cells have finite spatial extent and can affect one site much more than another. Consequently, SD typically yields substantial gains, particularly for rare, high-attenuation events that are associated with localized convective structures.

Simulation results highlight a clear dependence of diversity gain on both the number of deployed sites and their spacing. When sites are too close, correlation remains high and the gain is limited; as separation increases, the probability of simultaneous deep fades decreases and the diversity system provides a sharper reduction of exceedance probability. With a dual-site configuration, increasing the inter-site separation from $S = 20 \text{ km}$ to $S = 80 \text{ km}$, the exceedance

probability drops from $\hat{P}_{\text{single-site}}(A > 30 \text{ dB}) \approx 0.635\%$ to $\hat{P}_{\text{SD}, 20\text{km}}(A > 30 \text{ dB}) \approx 0.191\%$ and $\hat{P}_{\text{SD}, 80\text{km}}(A > 30 \text{ dB}) \approx 0.0272\%$, corresponding to an outage reduction by factors of ~ 3.3 and ~ 23 , respectively.

Moreover, increasing the number of sites enhances the selection opportunity, further improving the tail of the attenuation CCDF: the diagonal configuration ($k = 2$ ground stations) provides the smallest improvement, the triangular geometry ($k = 3$) yields a further downward shift, and the square configuration ($k = 4$) provides the most significant reduction of exceedance probability. This trend is reported in Figure 2.2.

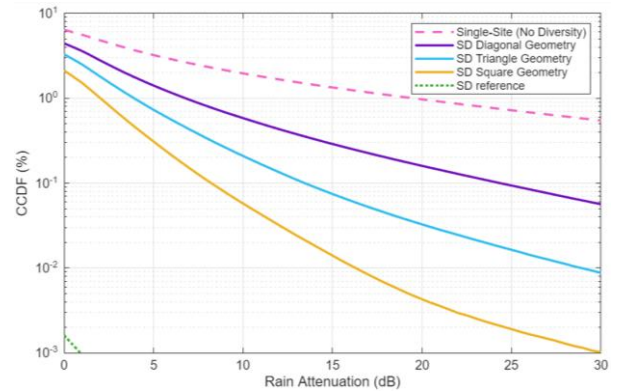


Figure 2.2: SD CCDFs for three geometry configurations with fixed site separation distance $S = 50 \text{ km}$.

3.3 Joint Site and Orbital Diversity

The joint SD+OD scheme applies a two-stage selection: first across sites, then across satellites, which is equivalent to selecting the global minimum attenuation across all site–satellite combinations:

$$A_{\text{joint}}(t_n) = \min_s \min_j A_{s,j}(t_n). \quad (2)$$

where s indexes the ground site and j indexes the common set of satellite links.

The joint strategy consistently provides the best performance among the tested configurations. Compared to SD-only, the joint results exhibit a markedly stronger reduction in the high-attenuation tail, which is the region most relevant for high-availability design. This indicates that the additional OD stage is particularly effective at mitigating deep fades. In the final analysis, the objective is to highlight the substantial benefits achieved by exploiting both the diversity order (i.e., the number of ground stations) and the deployment geometry, (i.e., inter-site separation). The simulation scenario, adopted for SD-only and SD+OD analyses, considers three ground-station

configurations - diagonal ($k = 2$), triangle ($k = 3$) and square ($k = 4$) – that are centered on the same rain-map region and two different inter-site separation distances are examined, $S = 30 \text{ km}$ and $S = 50 \text{ km}$. SD+OD reduces the exceedance probability by a factor of about 34x with respect to SD-only at $S = 30 \text{ km}$ and by about 670x with respect to SD-only at $S = 50 \text{ km}$. As shown in Figure 2.3, for higher attenuation thresholds, the joint SD+OD exceedance probability drops to the order of $10^{-4}\%$, rapidly approaching the statistical resolution limit of the dataset.

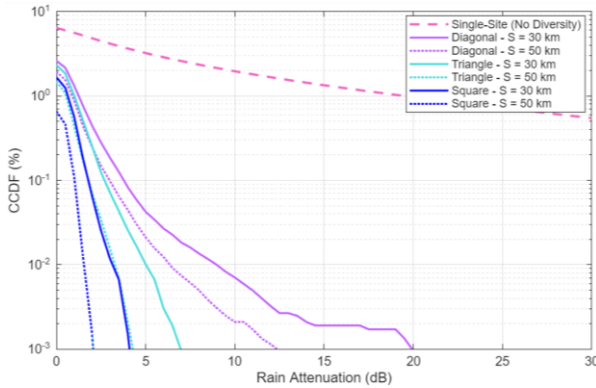


Figure 2.3: SD+OD CCDFs for three geometry configurations with $S = 30 \text{ km}$ and $S = 50 \text{ km}$.

4. Statistical Model for Orbital Diversity

While simulation-based studies provide high fidelity, they are computationally demanding because they require long time series, realistic rain-field evolution, and full constellation geometry. For early-phase design and dimensioning, it is therefore valuable to build reduced-complexity statistical models that preserve the key dependencies of OD performance on geometry and frequency. To this aim, this thesis introduces a statistical model centered on the attenuation difference between two simultaneously visible links. The analysis is performed on simulator outputs generated at three carrier frequencies (20, 30, and 40 GHz) to investigate possible frequency-dependent trends. For each pair of visible links (i, j), the absolute attenuation difference is computed as $|\Delta A| = |A_i - A_j|$ and associated with the angular separation $\theta \in [0, 180^\circ]$ between the two link directions. The resulting dataset is summarized through density plots: to improve robustness and generality, the analysis was not carried out on a single rain-map placement of the

gateway, but five independent simulation runs were performed, each using the same scenario configuration, but a different ground-station location. By conditioning the statistics of ΔA on this separation angle, the analysis shows that $|\Delta A|$ can be effectively modeled with a lognormal distribution, whose parameters depend smoothly on angle and frequency. Specifically, for each angular class c (centred at θ_c), the conditional distribution is modelled as:

$$f_{|\Delta A|}(|\Delta A| | \theta = \theta_c) \sim \text{Lognormal}(\mu(\theta_c), \sigma(\theta_c)). \quad (3)$$

where $\mu(\theta)$ and $\sigma(\theta)$ are the mean and standard deviation of $\ln|\Delta A|$. Repeating the estimation for all classes yields discrete estimates $\{\mu(\theta_c)\}$ and $\{\sigma(\theta_c)\}$ for each frequency. To enable analytic generation of the distribution at arbitrary θ , continuous parametric models for $\mu(\theta)$ and $\sigma(\theta)$ are introduced.

4.1 Model for $\mu(\theta)$

The selected model for $\mu(\theta)$ is:

$$\mu(\theta) = a_0 + a_1 \ln\left(\frac{\theta + \varepsilon}{180}\right) + a_2 \exp\left(\frac{\theta - 180}{\tau}\right), \varepsilon = 1^\circ \quad (4)$$

This structure captures a slow growth driven by the $\ln(\cdot)$ term, and a controlled “end-angle” behaviour through the exponential term centered at 180° . The fitted coefficients are reported in Table 1.1 for $\mu(\theta)$ at the three considered frequencies.

The frequency dependence of $\mu(\theta)$ is mainly captured by a vertical shift: the parameters a_1, a_2, τ remain approximately invariant, while the dominant variation occurs in a_0 . As a result, the curves preserve their shape across frequencies and therefore the final compact frequency-scalable expression is:

$$\mu(\theta, f) \approx \mu_{20}(\theta) + \beta \ln\left(\frac{f}{20}\right), \quad \beta = 1.94337. \quad (5)$$

Table 1.1: Coefficient values for $\mu(\theta)$ model.

Frequency	a_0	a_1	a_2	τ
20 GHz	0.05341	0.37	2.687	9.787
30 GHz	0.8849	0.3868	2.682	9.146
40 GHz	1.421	0.3868	2.623	8.788

4.2 Model for $\sigma(\theta)$

The modelling of $\sigma(\theta)$ is more complex. The behaviour from data estimate shows a sharp low-angle feature (rise and relaxation), a smoother mid-

angle trend and a tail toward large separations. For counting all these contributions, a two-component model is used:

- A smooth baseline (two exponential end-corrections):

$$\sigma_{base}(\theta) = \sigma_0 + c_1 \exp\left(-\frac{\theta}{\tau_1}\right) + c_2 \exp\left(-\frac{(180 - \theta)}{\tau_2}\right) \quad (6)$$

- A low-angle bump modelled as a lognormal-shaped term in θ :

$$\sigma_{bump}(\theta) = A \exp\left(-\frac{(\ln(\theta + \varepsilon) - \mu_B)^2}{2s_B^2}\right) \quad (7)$$

where $\varepsilon = 1^\circ$; the position of the maximum will be $\theta_{peak} \approx e^{\mu_B} - \varepsilon$. In this way, the complete model results as the sum of the two components. The fitted coefficients are reported in Table 1.2 at the three considered frequencies.

Table 1.2: Coefficient values for $\sigma(\theta)$ model.

Frequency	σ_0	c_1	τ_1	c_2	τ_2	A	μ_B	s_B
20 GHz	1.6685	-0.7107	7.98209	-0.23585	30.0623	0.4631	1.3065	0.9976
30 GHz	1.5805	-0.3195	25.9746	-0.27965	24.8418	0.4818	1.6145	1.3994
40 GHz	1.5653	0.2096	32.7017	-0.30595	29.8321	0.0805	1.7612	0.6282

At this step of study is not possible to find a relationship between those parameters and the working frequency, as the trend observing for $\sigma(\theta)$ does not show an approximately invariant shape across frequency over the whole angular range. This is most evident at low angular separations, where the relative ordering and curvature of the 20/30/40 GHz traces differ. This lack of shape invariance implies that a global frequency scaling based on a constant shift and/or scale factor cannot simultaneously fit the full θ -range. Figure 3.1 shows the fitted $\mu(\theta)$ (a) and $\sigma(\theta)$ (b) models versus per class estimates from simulated data.

4.3 Model-Based Reconstruction of $|\Delta A|$ Statistics

The statistical OD model is validated both locally and globally. Locally, for selected angular classes, the empirical PDFs of $|\Delta A|$ are compared with the corresponding lognormal PDFs defined by the fitted functions $\mu(\theta)$ and $\sigma(\theta)$. The model accurately reproduces not only the mean but also the overall distribution shape, showing excellent agreement for small and moderate angular separations. For large separations ($\theta \gtrsim 140^\circ$), minor discrepancies appear near the lowest attenuation bins, mainly due to reduced sample size and

increased histogram noise, although the tail behaviour remains well captured. Some PDFs examples are reported in Figure 3.2.

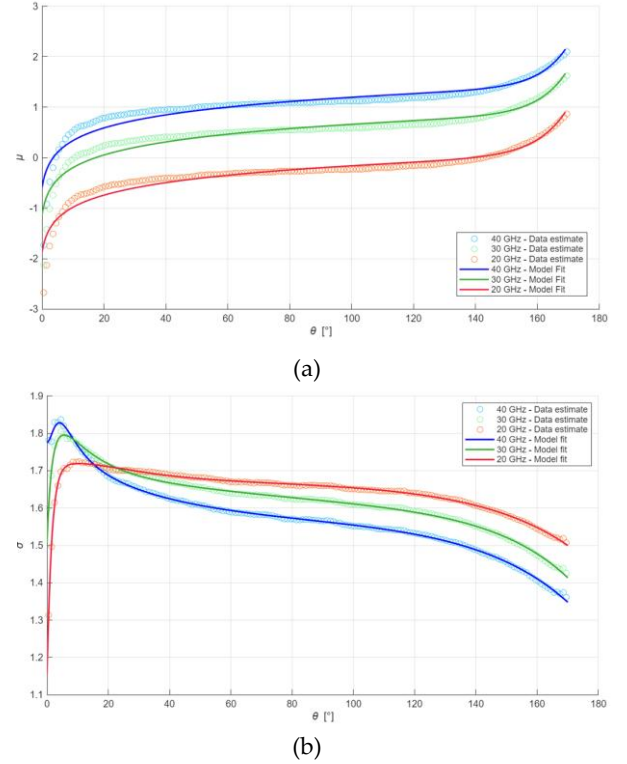


Figure 3.1: Fitted $\mu(\theta)$ and $\sigma(\theta)$ models versus per class estimates from simulated data.

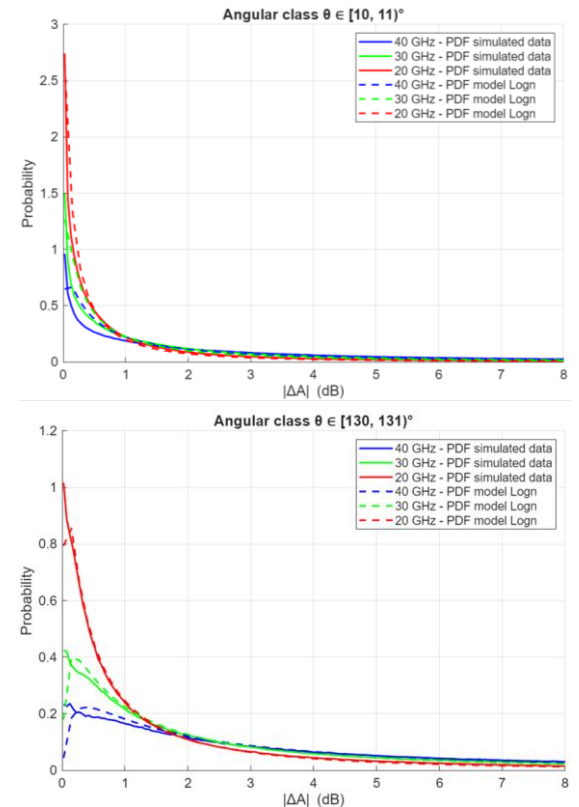


Figure 3.2: Empirical vs model-based PDFs for different angular classes.

Global validation is performed by reconstructing the full 2D density map through an analytic Monte Carlo procedure driven by the parametric model. Synthetic samples generated from the angle-dependent lognormal distributions are binned and normalized using the same procedure adopted for simulation results. The reconstructed maps closely match the empirical ones, with only a slight shift of the modal region toward smaller $|\Delta A|$. In Figure 3.3 is possible to see the comparison between empirical vs model-based density plots. Overall, while the lognormal assumption does not perfectly reproduce the fine structure around the mode, it provides the most consistent and accurate modelling framework across angular separations and frequencies among the tested distributions.




Figure 3.3: Empirical vs model-based density plots: (left) simulation results, (right) analytic Monte Carlo reconstruction using $\mu(\theta)$ and $\sigma(\theta)$ models.

5. Conclusions

As a whole, the work provides an engineering-oriented basis for designing spatial-diversity strategies in EHF NGSO systems, where availability is driven by rare, high-attenuation events. Beyond quantifying performance through CCDF analysis, a key outcome is a clearer interpretation of what limits each diversity mechanism: Orbital Diversity is primarily governed by the angular correlation of the rain field, while Site Diversity is governed by how fast spatial correlation decays with inter-site separation

and by the achievable gateway deployment density. The joint SD+OD scheme emerges as the most effective option because it leverages two complementary decorrelation mechanisms, increasing the probability that at least one site-satellite path avoids the most intense precipitation structures. In addition, the proposed OD statistical model represents a first step toward replacing heavy time-domain simulations with semi-analytical tools: by expressing $|\Delta A|$ statistics as a lognormal process conditioned on 3D angular separation and scalable with frequency, it enables fast exploration of OD design choices (e.g., candidate-set size and geometric separation) during early trade-off phases. Future developments should extend the framework from rain-only fading to full tropospheric impairment, assess more realistic switching/combining policies, and refine the OD model using more flexible distributions (e.g., mixtures) and broader parameterization across frequency, elevation angle, and climatic conditions.

References

- [1] ITU-R. *Propagation data and prediction methods required for the design of Earthspace telecommunication systems*. Recommendation P.618-14, 2023.
- [2] L. Luini, C. Capsoni, “MultiEXCELL: A new rain field model for propagation applications”, *IEEE Transactions on Antennas and Propagation*, vol. 59, no. 11, Page(s): 4286 – 4300, November 2011
- [3] Luini, L. & Wioland, Bastien & Riva, Carlo Giuseppe. (2022). *Orbital and Spatial Diversity for Next Generation Large NGSO Satellite Constellations*.1-6.10.1109/MMS55062.2022.9825604
- [4] ITU-R. *Rain attenuation model for rain for use in prediction methods*. Recommendation P.838-3, 2005.
- [5] T. Butash, P. Garland, and B. G. Evans, “Non-geostationary satellite orbit communications satellite constellations history,” *International Journal of Satellite Communications and Networking*, vol. 39, no. 1, pp. 1–5, 2021, doi: 10.1002/sat.1375.