ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DEPARTMENT of Electrical, Electronic, and Information Engineering

SECOND CYCLE DEGREE

Telecommunication Engineering

Ka-Band Propagation Channel Study based on In-Orbit NTN LEO Link Measurements

Dissertation in Satellite Communication and Navigation Systems

Supervisor:
Chiar.mo Prof. Alessandro Vanelli Coralli

Defended by:
Giorgio Cascelli

Co-Supervisor:
Stefano Cioni

Graduation Session July 2026

Academic Year 2025/2026

Abstract

This thesis develops and applies a complete signal processing chain for the study of Ka-band propagation from a Low Earth Orbit (LEO) satellite, a geometry whose fast, continuously varying elevation sets it apart from the well studied geostationary case. A 19.2 GHz continuous wave beacon transmitted by the LIDE cubesat was recorded at the ESTEC site of the European Space Agency by two co-located terminals: a high-gain Cobham SeaTel 1500 on a building roof and a low-profile Kebni vehicular terminal on a van. The chain first compensates the predictable terms of the link budget, namely the satellite transmit-power variation, the free space path loss and the gaseous attenuation, each reconstructed from independent ancillary data. A dedicated segmentation step, robust to the deep multipath fades of the van terminal, then isolates the interval in which the beacon is genuinely received. On the conditioned signal the excess attenuation and the tropospheric scintillation are characterised and compared with predictions derived from radiosonde data through the Rytov model. The measured scintillation agrees with the prediction, in the median, above about 25° of elevation, whereas its temporal spectrum cannot be resolved on the short, noise-limited LEO records. By exploiting the co-location of the two antennas, the dominant residual impairments are shown to be instrumental rather than atmospheric: the periodic pointing cycle of the Cobham and the ground multipath of the Kebni are identified and separated from the propagation effects. A sun-outage screen confirms that the Sun only grazed the antenna beams, raising the noise floor by less than a decibel without ever blocking the link. The results, and an honest account of their limitations, motivate a set of improvements for future campaigns and the development of a machine learning model for the prediction of link outages.

Contents

List of Abbreviations

AOS	Acquisition of Signal
CCDF	Complementary Cumulative Distribution Function
CSV	Comma-Separated Values
CW	Continuous Wave
ESA	European Space Agency
ESTEC	European Space Research and Technology Centre
FFT	Fast Fourier Transform
FSPL	Free Space Path Loss
GEO	Geostationary Earth Orbit
G/T	Antenna gain-to-noise-temperature ratio
IF	Intermediate Frequency
I/Q	In-phase and Quadrature
ITU-R	International Telecommunication Union, Radiocommunication Sector
KNMI	Royal Netherlands Meteorological Institute
LEO	Low Earth Orbit
LNB	Low Noise Block
LO	Local Oscillator
LoS	Line of Sight
LOS	Loss of Signal
MEO	Medium Earth Orbit
PA	Power Amplifier
PSD	Power Spectral Density
RAAN	Right Ascension of the Ascending Node
RAOB	Radiosonde Observation
RF	Radio Frequency
SDR	Software Defined Radio
SGP4	Simplified General Perturbations 4
SNR	Signal-to-Noise Ratio
SOTM	Satcom-On-The-Move
TCA	Time of Closest Approach
TLE	Two-Line Element
USRP	Universal Software Radio Peripheral
UTC	Coordinated Universal Time
VSAT	Very Small Aperture Terminal

Chapter 1

Introduction

The demand for reliable, high-capacity satellite communications has grown steadily over the past three decades, driven by the proliferation of broadband services, real-time multimedia applications, and the increasing need for global connectivity in sectors ranging from maritime navigation to emergency response. Satellite systems operating at higher microwave frequencies offer access to wider bandwidths and the possibility of using smaller ground terminals, both of which are fundamental prerequisites for the next generation of communication infrastructures. In this context, the Ka band, occupying the 17 to 40 GHz frequency range, has progressively emerged as a central candidate for commercial and institutional satellite services, attracting significant research interest to characterise the propagation environment and design robust link budgets [1].

At the same time, the orbital paradigm of the space segment is shifting. The traditional dominance of Geostationary Earth Orbit (GEO) satellites is being complemented, and in some application domains increasingly replaced, by constellations of Low Earth Orbit (LEO) satellites [2]. The substantially reduced altitude of LEO systems, typically between 300 and 2000 km, translates into lower propagation latency, reduced Free Space Path Loss (FSPL), and the possibility of polar coverage, making them particularly attractive for broadband internet access, Earth observation, and time-sensitive communication services. The combination of LEO geometry with Ka-band frequencies, however, introduces propagation challenges that differ significantly from the well-studied GEO case, and the experimental characterisation of these effects remains an open research area [3].

This thesis presents the work carried out in the framework of a measurement campaign conducted at the European Space Research and Technology Centre (ESTEC) of the European Space Agency (ESA), focused on the reception and processing of a Ka-band Continuous Wave (CW) beacon signal transmitted by the LIDE LEO cubesat [4]. The campaign employed two distinct ground terminals installed at the same site, allowing a direct comparison of propagation effects under different hardware and environmental conditions. The analysis developed in this work addresses the complete signal processing chain, from raw power samples to characterised propagation impairments, with particular attention to the specific challenges posed by the LEO geometry.

1.1 The Role of Low Earth Orbit Satellite Communications

LEO satellites occupy altitudes ranging from approximately 300 to 2000 km, completing a full orbit in roughly 90 to 130 minutes and achieving orbital velocities of the order of 7.5 km/s [5]. Their proximity to Earth results in an FSPL that is roughly 20 to 40 dB lower than that of a GEO satellite at 35 786 km, depending on the slant range, and their round-trip signal latency, typically below 15 ms, is more than an order of magnitude smaller than the roughly 250 ms round-trip delay of a GEO link.

The commercial landscape for LEO communications has evolved considerably. Early ventures such as the Teledesic T1 demonstration satellite, launched in 1998 into a 565 km polar orbit to evaluate Ka-band link performance for a planned 288-satellite constellation, demonstrated the technical feasibility of broadband LEO services but did not reach commercial deployment [6]. The Iridium constellation, operating at L band for subscriber handsets but using Ka-band inter-satellite and gateway links, demonstrated the viability of large-scale LEO networks and remained in continuous operation after a financial restructuring [6]. More recently, constellations such as Starlink and OneWeb have deployed hundreds to thousands of satellites in Ku and Ka band with the objective of providing global broadband coverage to fixed and mobile terminals [2].

Despite this commercial momentum, the propagation models used in the design of LEO Ka-band links remain largely adapted from GEO-derived empirical datasets, collected from fixed-elevation beacon experiments on satellites such as OLYMPUS, ITALSAT, and ACTS [6, 3]. A critical difference between GEO and LEO measurements is that the elevation angle of a LEO satellite changes rapidly and continuously throughout a pass. The speed at which the propagation path sweeps through the turbulent troposphere is then set by the satellite motion rather than by the wind. This alters the scintillation spectrum and makes the direct application of GEO-derived models unreliable [7, 8]. The LIDE campaign at ESTEC is motivated by the need to collect controlled propagation data under this dynamic geometry, contributing experimental evidence to support the adaptation and validation of propagation models for the LEO case.

1.2 Characteristics and Advantages of the Ka Band

The Ka band spans the frequency range from 17 to 40 GHz, with commercial satellite downlinks predominantly allocated in the 17.7 to 21.2 GHz portion and uplinks in the 27.5 to 30 GHz region, depending on regional regulatory frameworks. The primary motivation for operating in this band is the availability of several gigahertz of spectrum, enabling data rates that are unattainable in the congested C and Ku bands [1, 2]. For a given antenna aperture, antenna gain scales as the square of frequency, meaning that Ka-band terminals can achieve higher directive gain with smaller physical dimensions, a property of direct relevance for compact mobile and vehicular terminals.

The critical disadvantage of Ka-band operation is the sensitivity of the propagation channel to atmospheric impairments. Rain attenuation, the dominant impairment above 10 GHz, scales strongly with frequency: at 20 GHz, rain rates of the order of 10 to 15 mm/h can produce attenuation exceeding 7 to 8 dB at low elevation angles, and values well above 20 dB have been observed under heavy precipitation in many regions of the world [6]. Water vapour absorption reaches a local maximum near 22.235 GHz, adding a spectrally selective contribution to the total path attenuation that must be accounted for in the link budget [9]. Cloud liquid water content, though less severe than rain, contributes a diffuse attenuation component that is relevant for availability calculations at high elevation angles [10]. Tropospheric scintillation, caused by refractive index fluctuations in the lower troposphere induced by turbulent mixing of temperature and humidity, produces rapid amplitude and phase fluctuations that cannot be mitigated by simple power control [11].

The strongly stochastic nature of these impairments has driven decades of experimental beacon campaigns over GEO satellites, producing the semi-empirical models codified in recommendations of the International Telecommunication Union, Radiocommunication Sector (ITU-R), such as ITU-R P.618, P.676 and P.840 [11, 9, 10]. Building on earlier European campaigns conducted at lower frequencies (the Italian SIRIO and the ESA OTS, operating in the 11–18 GHz range), the OLYMPUS satellite, launched by ESA in 1989, provided the first large-scale European dataset at 20 and 30 GHz from a wide range of climatic sites [12], and the subsequent ACTS campaign, operated by NASA from 1993 onwards, extended these results across North America and firmly established the current set of rain attenuation and scintillation statistics used for GEO link design [13]. These models, however, were derived under the assumption of a fixed elevation angle, a condition that holds for GEO links but not for LEO passes, as the effective scintillation bandwidth is shifted to higher frequencies by the rapid orbital motion [8].

1.3 The ESTEC Measurement Context

ESTEC, located in Noordwijk, the Netherlands, is the main technical establishment of ESA and hosts a broad range of space engineering research facilities. The measurement campaign described in this thesis was carried out in the context of an ongoing effort at ESTEC to characterise the Ka-band LEO propagation channel using the LIDE cubesat as the signal source. LIDE transmits a CW tone at 19.2 GHz on its downlink, providing a spectrally clean, narrowband signal well suited to power-level measurements and Doppler analysis, following a similar approach to the beacon experiments performed over ROCSAT-1 at 19.5 GHz [7] and over FedSat at 20.13 GHz [6].

The experimental setup deployed two receiving terminals at the ESTEC site in a fixed configuration, designed to allow simultaneous acquisition from both stations during each satellite pass. The first terminal is a Cobham SeaTel 1500, a stabilised parabolic reflector antenna with a 1.5 m aperture, installed on the roof of a two-storey building within the ESTEC campus. The second terminal is a Kebni Satcom-On-The-Move (SOTM) low-profile Ka-band terminal designed for vehicular use, mounted on the roof of a van parked at ground level near the building. The two terminals differ not only in aperture and gain but also in their tracking architecture and their immediate electromagnetic environment, making them complementary instruments for the study of impairments that depend on the ground-level geometry, such as multipath from surrounding structures. A comparable challenge was identified in the FedSat experiment, where site obstructions were recognised as a significant source of measurement uncertainty [6].

The choice of ESTEC as a measurement site is motivated by the availability of infrastructure, the proximity to the De Bilt upper-air radiosonde station operated by the Royal Netherlands Meteorological Institute (KNMI), and the access to satellite scheduling through ESA operational channels. The radiosonde data collected at De Bilt, combined with surface meteorological records from the KNMI network, provides the atmospheric ground truth for the analysis. It is used to compute the theoretical gaseous attenuation correction [9] and to apply scintillation prediction models built on the Radiosonde Observation (RAOB) vertical profile of C_n^2 [14], enabling a quantitative comparison between measured and predicted channel behaviour.

1.4 Thesis Objectives and Structure

The primary objective of this thesis is the development and application of a complete signal processing pipeline for Ka-band CW data acquired from a LEO satellite, with the goal of characterising the propagation channel and identifying both atmospheric and instrumentation-related impairments. The specific objectives are the following.

- Organise and synchronise the multi-source dataset comprising Radio Frequency (RF) power samples, antenna telemetry from both terminals, satellite telemetry, surface meteorological records, and RAOB vertical profiles.
- Implement and validate a preprocessing chain covering FSPL compensation, satellite power amplifier gain correction, and gaseous attenuation removal based on ITU-R P.676-13.
- Identify and characterise instrumentation-related artefacts, specifically the 1 Hz pointing update cycle of the Cobham SeaTel 1500 and the multipath fading observed on the Kebni terminal.
- Apply the Rytov scintillation model to RAOB profiles and compare the predicted log-amplitude standard deviation with the measured value, assessing the validity of GEO-derived atmospheric models for the LEO propagation context.
- Provide a spectral and statistical characterisation of the scintillation observed on both terminals, with reference to the role of satellite velocity in determining the scintillation frequency content.

The remainder of this thesis is organised as follows. Chapter 2 provides the theoretical background necessary to interpret the measurement results, covering orbital mechanics and the Two-Line Element (TLE) format, the link budget and the FSPL, the Doppler effect in LEO passes, atmospheric propagation impairments, and multipath fading. Chapter 3 describes the experimental setup in detail, including the hardware specifications of both receiving terminals, the RF chain and signal digitisation, the acquisition software architecture, and the methodologies used for CW peak detection and Doppler computation. Chapter 4 presents the signal processing results and their analysis: the FSPL compensation, the satellite transmitter degradation, the pointing artefacts observed on the Cobham terminal, the multipath fading on the Kebni terminal, and the tropospheric scintillation measured against the radiosonde prediction. Chapter 5 draws the two stations together into a comparative analysis, summarises the key findings, and outlines directions for future work, including improvements to the ground terminal configuration and the development of propagation models adapted to the LEO geometry.

Chapter 2

Background

2.1 Satellite Communications and LEO Orbital Dynamics

LEO constellations are fundamentally reshaping the landscape of satellite communications. Unlike traditional GEO systems, LEO systems operate at altitudes typically ranging from 300 km to 2,000 km above the Earth's surface [1]. Because of this close proximity to the ground, LEO satellites inherently mitigate the severe propagation delay and path loss experienced in GEO, offering substantial benefits in terms of latency and bandwidth for global connectivity [2, 15].

Due to the laws of orbital mechanics, a satellite at these low altitudes must travel at extremely high velocities to counterbalance the Earth's gravitational pull and maintain a stable orbit. For instance, a typical LEO spacecraft (such as the LIDE satellite) travels at an orbital velocity of approximately 7.5 km/s [5]. Consequently, the satellite completes a full revolution around the Earth in a very short orbital period, usually ranging from 90 to 130 minutes depending on the specific altitude. For a ground station, this high velocity translates into a very brief visibility window during a satellite pass, typically lasting only a few minutes, which necessitates highly dynamic and precise antenna tracking systems.

To accurately track a fast-moving LEO satellite, ground stations rely on pre-calculated orbital parameters formatted as a TLE set. The TLE is a standard data format introduced to the user community alongside the Simplified General Perturbations 4 (SGP4) propagation models in the seminal Spacetrack Report No. 3 [16]. The TLE encodes essential orbital parameters such as inclination, eccentricity, the right ascension of the ascending node, argument of perigee, and mean anomaly. Furthermore, the format includes dynamic variables like the mean motion (the number of orbits completed per day) and its time derivatives, as well as the BSTAR drag term, which estimates the atmospheric decay rate of the satellite due to aerodynamic drag in the upper atmosphere (see **Table 2.1** and **Table 2.2**).

Table 2.1: TLE format, Line 1. [16]

Columns	Content
01	Line number
03–07	Satellite catalog number
08	Classification (U: unclassified, C: classified, S: secret)
10–11	International Designator (last two digits of launch year)
12–14	International Designator (launch number of the year)
15–17	International Designator (piece of the launch)
19–20	Epoch year (last two digits of year)
21–32	Epoch (day of the year and fractional portion of the day)
34–43	First derivative of mean motion, the ballistic coefficient (rev/day, per day)
45–52	Second derivative of mean motion (rev/day^3, decimal point assumed)
54–61	B^*, the drag term, or radiation pressure coefficient (units of 1/(Earth radii), decimal point assumed)
63	Ephemeris type (always zero, only used in undistributed TLE data)
65–68	Element set number. Incremented when a new TLE is generated for this object
69	Checksum (modulo 10)

Table 2.2: TLE format, Line 2. [16]

Columns	Content
01	Line number
03–07	Satellite Catalog number
09–16	Inclination (degrees)
18–25	Right ascension of the ascending node (degrees), in the ECI reference frame, measured from the vernal point
27–33	Eccentricity (unitless, decimal point assumed)
35–42	Argument of perigee (degrees)
44–51	Mean anomaly (degrees)
53–63	Mean motion (revolutions per day)
64–68	Revolution number at epoch (revolutions)
69	Checksum (modulo 10)

In a practical measurement campaign, such as the one conducted at ESTEC, the TLE data is ingested by an orbit propagator model (typically SGP4) to predict the satellite's exact position relative to the ground station over time. This approach, known as program tracking, generates a continuous schedule of azimuth and elevation coordinates that the ground station antenna controller uses to mechanically steer the parabolic dish and maintain the satellite within its main beam during the entire pass.

2.1.1 Keplerian Orbital Parameters and TLE Representation

The exact state of a satellite at a specific point in time, known as the epoch, is universally described using the TLE set format [16]. A TLE encodes the classical Keplerian orbital elements, which are mathematically necessary to uniquely define an unperturbed orbit in a three-dimensional Cartesian space relative to a central body (the Earth). The SGP4 propagation model uses these parameters to calculate the satellite's exact position and velocity vectors over time, allowing ground station antennas to maintain accurate program tracking.

To fully understand the orbital dynamics handled by the SGP4 algorithm, it is essential to define the six classical Keplerian elements extracted from the TLE:

- **Inclination (i):** The angle between the Earth's equatorial plane and the satellite's orbital plane. An inclination of $0°$ indicates an equatorial orbit, while $90°$ defines a polar orbit.
- **Right Ascension of the Ascending Node (Ω or RAAN):** This angle specifies the orientation of the orbital plane in space. It is measured in the equatorial plane from a reference direction (the First Point of Aries or Vernal Equinox) to the ascending node, which is the point where the satellite crosses the equator moving from the Southern to the Northern Hemisphere.
- **Eccentricity (e):** A dimensionless parameter defining the shape of the orbit. A value of $e = 0$ corresponds to a perfectly circular orbit, while $0 < e < 1$ denotes an elliptical orbit.
- **Argument of Perigee (ω):** This angle defines the orientation of the ellipse within the orbital plane. It is measured from the ascending node to the perigee (the point of closest approach to the Earth) in the direction of the satellite's motion.
- **Mean Anomaly (M):** An angular parameter that defines the specific position of the satellite along its orbit at the given epoch time. It is a mathematically constructed angle that assumes the satellite moves at a constant angular speed.
- **True Anomaly (ν or θ):** The actual geometric angle measured in the orbital plane from the perigee to the current physical position of the satellite. While the mean anomaly is used in TLE sets for computational simplicity, the true anomaly describes the exact real location of the spacecraft, accounting for the fact that its speed varies along the orbit according to Kepler's Second Law.

- **Mean Motion (n):** The number of complete orbits the satellite makes around the Earth in one solar day. In TLE formats, this value directly relates to the orbital period and the semi-major axis (a) via Kepler's Third Law.

Beyond the classical Keplerian elements, the TLE format also includes the B^* (BSTAR) drag term [16]. This is a pseudo-drag coefficient used by the SGP4 model to account for non-conservative perturbations, primarily atmospheric drag, which constantly causes the orbit of LEO satellites to decay.

Figure 2.1 provides a visual representation of these parameters relative to the Earth's equatorial plane and the vernal equinox vector.

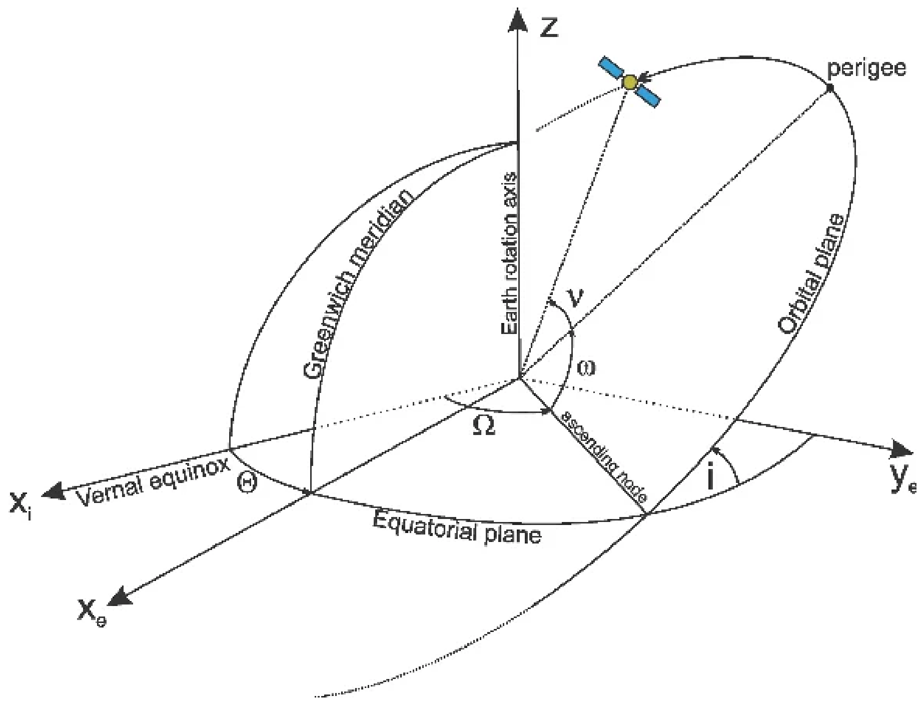

Figure 2.1: Geometric representation of the classical Keplerian orbital elements (Inclination, RAAN, Argument of Perigee, and Anomaly) defining a satellite's orbit around the Earth.

2.2 The Ka-Band Link Budget Analysis

The link budget is the fundamental system-level calculation used to determine the viability and performance of a satellite communication channel. It accounts for all the gains and losses in the telecommunications system, from the transmitter onboard the spacecraft to the receiver at the ground station. In recent years, the telecommunications industry has increasingly shifted towards the Ka-band to satisfy the growing demand for higher data rates.

The primary advantage of the Ka-band over legacy bands (like C or Ku, see **Table 2.3**) is the vast amount of available spectrum, which translates directly into significantly higher data capacity and broadband throughput capabilities [1, 2]. Furthermore, because the physical size of an antenna required to achieve a specific gain is inversely proportional to the frequency, the Ka-band allows for much smaller ground terminals and highly compact payload antennas on LEO satellites.

Table 2.3: Frequency bands for radar and space radiocommunications from ITU-R [17]. (1) For space radiocommunications K and Ka bands are often designated by the single symbol Ka.

Letter symbols	**Radar (GHz)**		**Space radiocommunications**	
	Spectrum regions	**Examples**	**Nominal designations**	**Examples (GHz)**
L	1-2	1.215-1.4	1.5 GHz band	1.525-1.710
S	2-4	2.3-2.5 2.7-3.4	2.5 GHz band	2.5-2.690
C	4-8	5.25-5.85	4/6 GHz band	3.4-4.2 4.5-4.8 5.85-7.075
X	8-12	8.5-10.5	-	-
Ku	12-18	13.4-14.0 15.3-17.3	11/14 GHz band 12/14 GHz band	10.7-13.25 14.0-14.5
K(1)	18-27	24.05-24.25	20 GHz band	17.7-20.2
Ka(1)	27-40	33.4-36.0	30 GHz band	27.5-30.0
V	-	-	40 GHz band	37.5-42.5 47.2-50.2

However, propagating signals at these high frequencies comes with a major mathematical drawback in the link budget: the FSPL. The FSPL represents the geometric spreading of electromagnetic energy over distance, and it is the most dominant loss component in any satellite link. It is defined in decibels (dB) by the following equation:

$$FSPL_{dB} = 20\log_{10}(d) + 20\log_{10}(f) + 20\log_{10}\left(\frac{4\pi}{c}\right) \tag{2.1}$$

where d is the slant range (the distance between the satellite and the ground station) and f is the operating frequency. As demonstrated by Equation 2.1, moving to the higher frequencies of the Ka-band inherently increases the baseline path loss.

In GEO links, the distance d is relatively constant (approximately 35 786 km), leading to a static FSPL. Conversely, in LEO satellite communications, the slant range d is highly dynamic and depends critically on the satellite's altitude H and its elevation angle ε relative to the ground station's horizon [18], see **Figure 2.2**. The slant range can be calculated using the following geometric relation:

$$d(\varepsilon) = R_E\left[\sqrt{\left(\frac{H+R_E}{R_E}\right)^2 - \cos^2(\varepsilon)} - \sin(\varepsilon)\right] \tag{2.2}$$

where R_E is the Earth's radius.

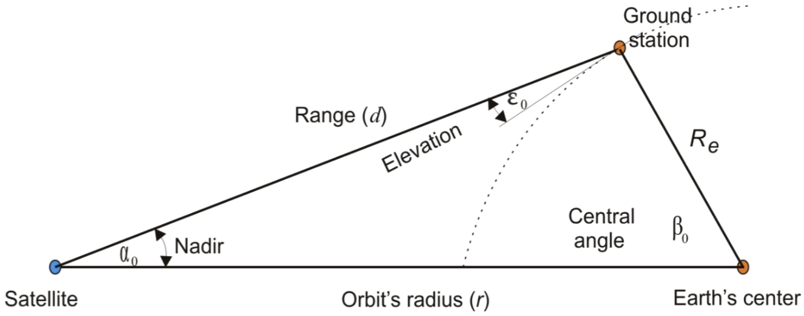

Figure 2.2: Geometry of the ground station and the satellite, adapted from [18].

During a typical LEO pass, the satellite first appears at a very low elevation angle near the horizon, where the slant range is at its maximum. As the satellite approaches the zenith (maximum elevation), the slant range decreases to its minimum value, which is simply the orbital altitude H. Because of the logarithmic dependence of FSPL on the distance d, this rapid variation in slant range causes drastic fluctuations in the free space attenuation during the short visibility window (typically 5 to 10 minutes). Consequently, the received signal power at the ground station continuously varies, demanding robust dynamic range capabilities from the receiving front-end to avoid saturation at the zenith and signal loss at the horizon.

2.3 Atmospheric and Solar Impairments

The propagation of radio waves from a satellite to a ground station is highly dependent on the physical properties of the medium through which they travel. For lower frequencies, such as the L or S bands, the Earth's atmosphere is mostly transparent. However, as telecommunications systems shift towards the Ka-band (17–40 GHz) to meet the demand for high data rates, the shorter wavelengths become extremely susceptible to various environmental impairments, primarily located in the troposphere (from the ground up to roughly 10–15 km) and, to a lesser extent, the ionosphere [19, 20]. It is convenient to separate these tropospheric impairments into three families according to how they behave in time and how they can be treated in the measurement chain. *Gaseous absorption* by oxygen and water vapour is a slowly varying, spatially smooth contribution that can be predicted from meteorological data and removed deterministically. *Cloud and rain attenuation* are caused by hydrometeors that are highly localised and intermittent in space and time, and therefore cannot be reconstructed from a single ground reference. *Tropospheric scintillation* is a fast, zero-mean fluctuation produced by clear-air turbulence, which is characterised statistically rather than removed. The following subsections describe each of these mechanisms, the relevant ITU-R recommendations, and the consequences of each for the LIDE measurement campaign.

2.3.1 Gaseous Absorption by Oxygen and Water Vapour

Even under perfectly clear-sky conditions the atmosphere is not transparent at Ka band. The signal is attenuated by molecular absorption, dominated by two species: diatomic oxygen (O_2), which possesses a dense complex of magnetic-dipole resonances clustered around 50–70 GHz plus an isolated line at 118.75 GHz, and water vapour (H_2O), whose first relevant electric-dipole rotational line sits at 22.235 GHz [9]. The LIDE downlink at 19.2 GHz lies on the low-frequency shoulder of the 22.235 GHz water-vapour line, so the gaseous budget at this frequency is dominated by water vapour and is therefore sensitive to the humidity of the air column, while the oxygen contribution is smaller and far more stable. **Figure 2.3** shows the specific (per-kilometre) attenuation of the two constituents as a function of frequency.

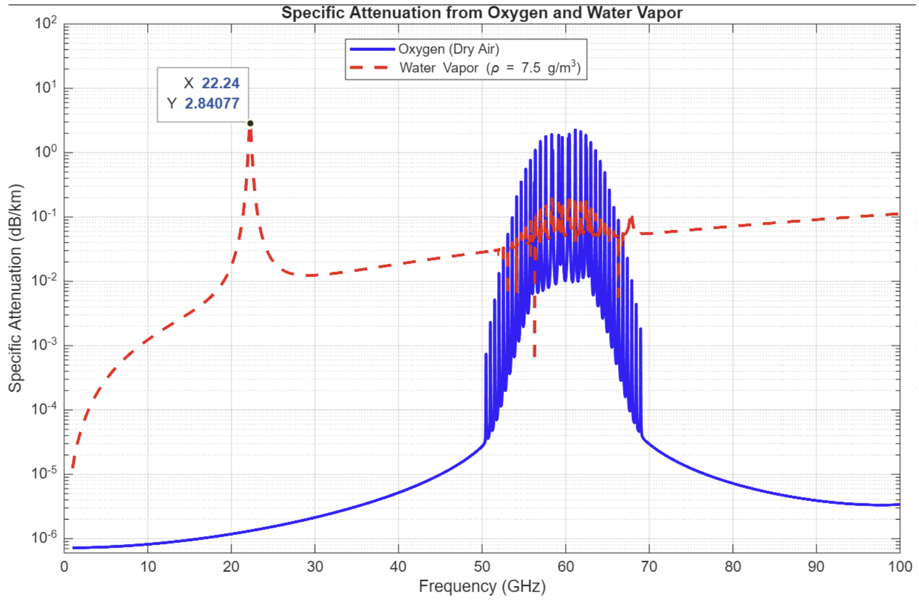

Figure 2.3: Specific attenuation due to atmospheric gases as a function of frequency. The data are calculated according to the approximate method described in Recommendation ITU-R P.676-13 [9], assuming standard atmospheric conditions at sea level (temperature 15 °C, pressure 1013.25 hPa, and water vapour density 7.5 g/m^3).

The reference framework for computing this contribution is Recommendation ITU-R P.676-13 [9], which provides two complementary methods. **Annex 1** is a line-by-line summation over the individual oxygen and water-vapour spectral lines; it is the most accurate formulation and, when fed with a vertical profile of pressure, temperature and humidity, yields the specific attenuation $\gamma(h) = \gamma_o(h) + \gamma_w(h)$ [dB/km] at each height. **Annex 2** is a simplified, computationally light approximation in which the oxygen and water-vapour specific attenuations are expressed in closed form as functions of the local pressure p, temperature T and water-vapour density ρ, valid up to 350 GHz. In both cases the total *zenith* (vertical) attenuation is obtained by integrating the specific attenuation over height,

$$A_{\text{zenith}} = \int_0^\infty \gamma(h)\,\mathrm{d}h \;\approx\; \gamma_o\,h_o + \gamma_w\,h_w, \tag{2.3}$$

where the right-hand side is the Annex 2 *equivalent-height* form: h_o and h_w are the equivalent heights of the oxygen and water-vapour columns, computed from the surface pressure and temperature. The attenuation along the slant path at elevation ε then follows the cosecant law for a plane-stratified atmosphere,

$$A_{\text{gas}}(\varepsilon) = \frac{A_{\text{zenith}}}{\sin\varepsilon}. \tag{2.4}$$

In this work the gaseous term is computed from the De Bilt radiosonde ascent using the ITU-R P.676-13 **Annex 1** line-by-line model: the specific attenuation $\gamma(h)$ is evaluated at every reported level from the measured pressure, temperature and humidity and integrated over height to give the zenith attenuation $A_{\text{zenith}} = \int_0^\infty \gamma(h)\,\mathrm{d}h$, which is then projected onto the slant path through Equation (2.4). The line-by-line model is preferred over the simplified Annex 2 closed form because the radiosonde resolves the actual vertical distribution of water vapour, which dominates the budget at 19.2 GHz. Because oxygen is uniformly mixed and the water-vapour field, although variable, changes smoothly and slowly over horizontal scales of tens of kilometres, the column represented by a single radiosonde is a faithful estimate of the gas the slant path actually traverses. The gaseous contribution is therefore a *deterministic*, predictable term. It is computed level by level and compensated in the preprocessing chain, where the predicted loss is added back to the measured received power, following the same philosophy adopted for multi-satellite propagation campaigns [21].

2.3.2 Cloud and Rain Attenuation

Above the clear-air gaseous floor, the largest impairments at Ka band are produced by liquid hydrometeors. **Cloud attenuation** arises from non-precipitating liquid water droplets, whose radii (a few to tens of micrometres) are far smaller than the 1.5 cm wavelength at 20 GHz, placing the interaction in the Rayleigh regime. In this regime the scatterer is much smaller than the wavelength, so the attenuation cross-section scales with the droplet *volume* rather than with its geometric cross-section; the cloud attenuation consequently depends only on the total liquid water content of the column and not on the detailed drop-size distribution. Recommendation ITU-R P.840 [10] formalises exactly this by modelling the cloud specific attenuation as proportional to the columnar liquid water content L [kg/m^2] through a temperature-dependent specific attenuation coefficient $K_l(f, T)$, giving a diffuse contribution of a few tenths of a decibel that is relevant mainly for high-availability link design.

Rain attenuation is the dominant impairment whenever it occurs. When the raindrop diameter becomes comparable to the wavelength (drops of 1–5 mm against 1.5 cm at 20 GHz), the wave is both absorbed, with its energy converted to heat inside the drop, and scattered out of the path. Heavy precipitation can therefore degrade the received power by tens of decibels [20]. The standard prediction framework is Recommendation ITU-R P.618 [11], which combines the specific attenuation power law of Recommendation ITU-R P.838,

$$\gamma_R = k\,R^{\alpha} \quad [\mathrm{dB/km}], \tag{2.5}$$

where R is the rain rate [mm/h] and the coefficients k and α depend on frequency and polarisation [22], with the rain-height model of Recommendation ITU-R P.839 [23] and the rainfall-rate statistics of Recommendation ITU-R P.837 [24]. The P.618 method introduces horizontal and vertical path-reduction factors to convert the specific attenuation into the slant-path attenuation $A_{0.01}$ exceeded for 0.01% of an average year. It must be stressed that these reduction factors are *long-term statistical* adjustments calibrated for the 0.01% point and are not meant to reconstruct the attenuation of an individual instantaneous event.

Why rain and cloud attenuation cannot be compensated in this campaign

This statistical character is precisely the difficulty for a single-pass, single-site measurement. Compensating an impairment in the received power requires knowing its instantaneous value *along the actual slant path*. For the gaseous term this is feasible (Section 2.3.1) because the absorbing gas is smooth and well mixed. Rain and clouds are the opposite: convective rain cells have characteristic horizontal scales of only a few kilometres and lifetimes of minutes, and cloud cover is broken and inhomogeneous, so the hydrometeor field is strongly intermittent in both space and time [25].

The geometry of a LEO pass aggravates this. The available ground truth is a single point: the surface weather record and the rain gauge of one station, plus a radiosonde ascent launched from De Bilt twice a day, which samples a vertical column at a fixed location that is not co-located with the link. During a pass, however, the satellite sweeps from horizon to horizon in a few minutes; the slant path swings across a wide range of azimuth and elevation and, projected onto the rain layer at a few kilometres of altitude, its piercing point migrates over many kilometres of sky. The beam therefore interrogates a large, horizontally inhomogeneous volume of troposphere that is mostly far from the ground reference. A point rain gauge cannot tell whether the segment of path actually traversed at a given instant passed through a rain cell, missed it, or clipped its edge, nor can it provide the rain rate *distributed along* that oblique path. The same argument applies to broken clouds.

For these reasons no attempt is made in this work to remove rain or cloud attenuation from the measured power: doing so from a single weather station would introduce errors larger than the effect being corrected. This limitation is intrinsic to the geometry rather than to the instrumentation: a faithful reconstruction of rain along a LEO slant path would require spatially resolved data, such as weather-radar fields or a synthetic rain-cell model of the kind used to simulate LEO fade dynamics [25], which lie outside the scope of the present campaign.

2.3.3 Tropospheric Scintillation

Unlike the bulk attenuation produced by gases and hydrometeors, tropospheric scintillation manifests as rapid, zero-mean fluctuations of the amplitude and phase of the received signal [26]. It is caused by small-scale turbulent mixing in the troposphere, which creates random fluctuations of the refractive index along the path; as the wavefront crosses these eddies it is alternately focused and defocused, producing fast fading at the antenna. Because the fluctuations have zero mean they cannot be removed by power control, only characterised statistically, and their intensity grows at low elevation, where the path through the turbulent layers is longest. The physical chain adopted here proceeds in three steps: (i) the refractive-index structure parameter C_n^2 is reconstructed as a function of height from a radiosonde ascent; (ii) the Rytov weak-fluctuation theory converts the C_n^2 profile into the variance of the log-amplitude fluctuations along the slant path; and (iii) the resulting prediction is applied to a LEO pass as a function of the instantaneous elevation.

Two features of these fluctuations recur throughout the later analysis and are worth defining at the outset. The first is the *inertial subrange* of the turbulence. Atmospheric turbulence organises the air into swirls, or eddies, of many different sizes: energy is injected at the largest (outer) scales by wind shear and convection, is handed down to progressively smaller eddies, and is finally dissipated into heat by viscosity at the smallest (inner) scale. Over the intermediate range of sizes, called the inertial subrange, energy is neither produced nor dissipated but simply cascades down the chain. There, Kolmogorov's theory of locally isotropic turbulence fixes the statistics of the refractive index field. As the wind, or the motion of the satellite, carries these eddies across the line of sight, each eddy size imprints a fluctuation on the received power at a corresponding rate: large eddies produce slow (low frequency) fluctuations and small eddies fast (high frequency) ones, so eddy size maps directly onto temporal frequency. The temporal power spectrum of the log-amplitude therefore inherits the cascade law: above a lower cut-off it rolls off as $f^{-8/3}$, where f is the fluctuation frequency of the received power in hertz [27, 28]. Observing that particular slope is the spectral signature of genuine turbulence.

The second feature is that lower cut-off itself, the scintillation *corner frequency* (or Fresnel frequency). Below it the fluctuation spectrum is essentially flat; above it the power rolls off as $f^{-8/3}$ as just described. Physically the corner frequency is set by how quickly the eddies are swept across the beam. It is roughly the transverse crossing speed divided by the width of the first Fresnel zone, the patch of atmosphere that the beam effectively samples, so a faster wind, or a faster satellite, pushes it towards higher frequencies. This dependence on the crossing speed is what makes the corner frequency, unlike the scintillation *variance*, sensitive to the orbital geometry, as discussed at the end of this section.

The refractive-index structure parameter C_n^2

At microwave frequencies the radio refractivity $N = (n - 1) \times 10^6$ of moist air depends on pressure, temperature and humidity [29]. Turbulent mixing stirs air parcels of different refractivity together, and the statistics of the resulting small-scale refractive-index field are quantified, within Kolmogorov's theory of locally isotropic turbulence, by the refractive-index structure parameter C_n^2 [m$^{-2/3}$], which varies with the height z above the surface. Following the classical derivation of Tatarski [28], C_n^2 is related to the vertical gradient of a conserved generalised refractivity M and to the outer scale of turbulence L_0 by

$$C_n^2(z) = a^2 \, \frac{K_H}{K_M} \, L_0^{4/3} \left(\frac{\partial M}{\partial z} \right)^2 , \tag{2.6}$$

where z is the altitude above the ground along the vertical profile, a^2 is an empirical constant of order unity, and K_H/K_M is the ratio of the turbulent exchange coefficients for heat and momentum. The symbols used throughout this section are collected in **Table 2.4**. The way the conserved variable in M is defined is the crucial modelling choice that distinguishes the two formulations used in this thesis.

Table 2.4: Symbols used in the tropospheric-scintillation model.

Symbol	**Quantity**	**Units**
z	altitude above the ground (vertical coordinate)	m
$C_n^2(z)$	refractive-index structure parameter	$\mathrm{m}^{-2/3}$
$n,\ N$	radio refractive index; refractivity $N = (n-1) \times 10^6$	–; N-units
M	generalised (moist) radio refractivity	N-units
$\partial M/\partial z$	vertical gradient of M	m^{-1}
a^2	empirical structure constant (≈ 2.8)	–
K_H/K_M	turbulent exchange-coefficient ratio (heat/momentum)	–
L_0	outer scale of turbulence (≈ 100 m)	m
T	absolute temperature	K
p	atmospheric pressure	hPa
q	specific humidity	kg/kg
θ	potential temperature, $\theta = T(1000/p)^{0.286}$	K
H	pseudopotential temperature, $H = T + (g/c_p)z$	K
g	gravitational acceleration	$\mathrm{m\,s^{-2}}$
c_p	specific heat of dry air at constant pressure	$\mathrm{J\,kg^{-1}\,K^{-1}}$
c	moist-refractivity coefficient ($\approx 7.8 \times 10^3$)	K
σ_χ^2	log-amplitude variance	Np^2
σ_χ	log-amplitude standard deviation	Np or dB
χ	log-amplitude fluctuation	Np
k	wavenumber, $k = 2\pi/\lambda$	$\mathrm{rad\,m^{-1}}$
$\lambda,\ f$	wavelength; carrier frequency	m; Hz
ε	elevation angle	deg
I	path integral $\int_0^\infty C_n^2(z)\, z^{5/6}\, \mathrm{d}z$	$\mathrm{m}^{7/6}$

Tatarski formulation. The original derivation builds M from the *pseudopotential* (or potential refractive) temperature obtained under the hydrostatic approximation, $H = T + (g/c_p)\, z$, where g is gravity and c_p the specific heat of dry air [28]. This dry-static-energy variable is a convenient surrogate for a conserved quantity within the boundary layer but is only approximately conserved aloft.

Cherubini & Businger formulation. Cherubini and Businger [30] showed that replacing H with the true *potential temperature*

$$\theta = T\left(\frac{1000}{p}\right)^{0.286} \tag{2.7}$$

removes a long-standing bias of the hydrostatic approximation. The two definitions agree closely below about 6–7 km. Above that height the atmosphere is increasingly stable, and near the tropopause the temperature itself levels off or even inverts; there the true potential temperature θ, which depends on the actual pressure through $(1000/p)^{0.286}$, keeps increasing sharply, while the pseudopotential $H = T + (g/c_p)z$ can only grow at the fixed dry-adiabatic rate built into its definition and so fails to capture this enhanced stability. Consequently $\partial\theta/\partial z$ exceeds $\partial H/\partial z$ aloft, and the pseudopotential formulation *underestimates* the refractivity gradient, and hence C_n^2, in the upper troposphere and lower stratosphere. Retaining the moisture term and expressing the gradient in terms of the potential temperature θ and the specific humidity q, the moist refractivity gradient takes the form (Cherubini & Businger [30], eqs. 17–20; applied to a Q-band slant link by Marziani *et al.* [14])

$$\frac{\partial M}{\partial z} = \frac{80\times 10^{-6}\,p}{T\,\theta}\left[\left(1+\frac{2c\,q}{T}\right)\frac{\partial\theta}{\partial z} - c\,\frac{\theta}{T}\,\frac{\partial q}{\partial z}\right], \qquad c \approx 7.8\times 10^{3}\ \mathrm{K}, \tag{2.8}$$

with p in hPa and T, θ in kelvin. The Tatarski profile is recovered by substituting the pseudopotential temperature H for θ in the gradient term of Equation (2.8), so that the comparison between the two isolates exactly the effect of the conserved variable.

In practice the radiosonde levels are first cleaned and interpolated onto a uniform vertical grid (50 m steps); at each mid-layer the gradients $\partial\theta/\partial z$ and $\partial q/\partial z$ are evaluated by finite differences and inserted into Equations (2.6)–(2.8), with the empirical constants $a^2 = 2.8$ and $L_0 \approx 100$ m adopted from [30]. The result is a vertical profile $C_n^2(z)$, peaking near the surface and in the moist mid-troposphere and decaying with height. **Figure 2.4** overlays the profile obtained with the potential temperature (Cherubini) and with the pseudopotential H (Tatarski) computed from the same ascent, reproducing the divergence above ∼6–7 km: on this sounding the Tatarski profile carries only about half of the Cherubini path integral above 6.5 km, and correspondingly predicts a lower σ_χ (0.32 against 0.34 dB).

From the C_n^2 profile to scintillation variance: the Rytov model

Once $C_n^2(z)$ is known, weak-fluctuation (Rytov) theory predicts the variance of the log-amplitude fluctuation χ of a spherical wave propagating along the slant path. Integrating the Kolmogorov spectrum over the path yields the closed-form result used in the radiosonde-based prediction methods of Vasseur [31] and Marzano & d'Auria [32]; the same frequency and elevation scaling laws underpin the scintillation model of Recommendation ITU-R P.618 [11]:

$$\sigma_\chi^2 = 0.5631\,k^{7/6}\,(\sin\varepsilon)^{-11/6}\int_0^\infty C_n^2(z)\,z^{5/6}\,\mathrm{d}z \quad [\mathrm{Np}^2], \tag{2.9}$$

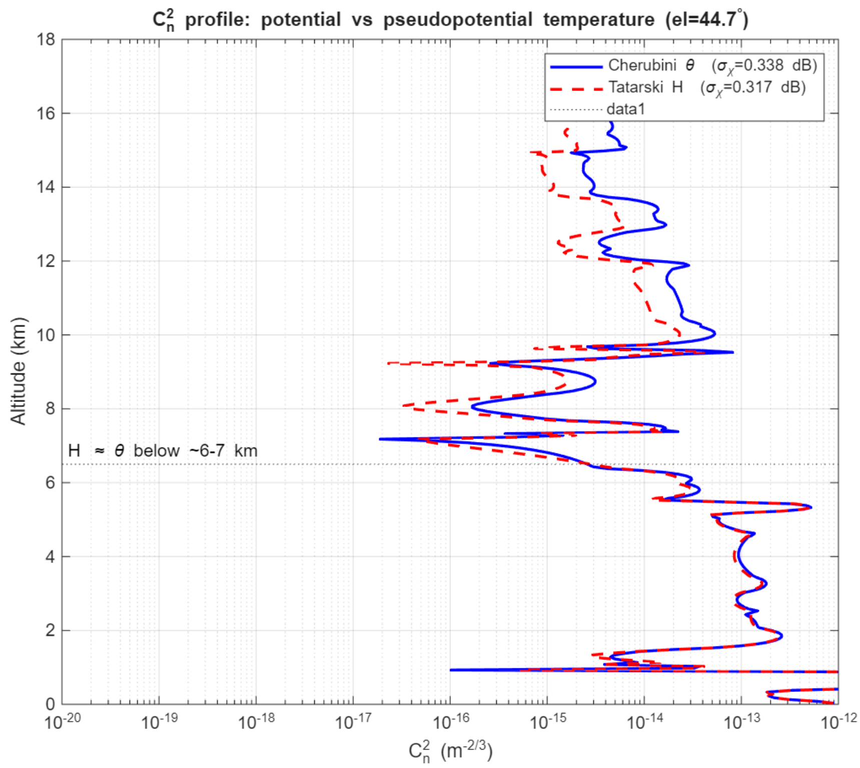

Figure 2.4: Vertical profile of the refractive-index structure parameter $C_n^2(z)$ computed from a De Bilt radiosonde ascent with the potential-temperature (Cherubini & Businger [30]) and the pseudopotential (Tatarski [28]) formulations. The two agree below ∼6–7 km; above it the fixed dry-adiabatic slope of the pseudopotential temperature H misses the enhanced static stability near the tropopause that the true potential temperature θ captures, so the pseudopotential formulation *underestimates* C_n^2 there and predicts a lower σ_χ than Cherubini's.

where $k = 2\pi/\lambda = 2\pi f/c$ is the wavenumber and ε is the elevation angle. The corresponding standard deviation expressed in decibels is

$$\sigma_\chi\,[\mathrm{dB}] = (20\log_{10} e)\,\sqrt{\sigma_\chi^2} = 8.686\,\sqrt{\sigma_\chi^2}. \tag{2.10}$$

Each factor in Equation (2.9) has a clear physical meaning. The $k^{7/6}$ dependence makes the scintillation grow with frequency, so that Ka band is markedly more affected than the lower bands. The geometric factor $(\sin\varepsilon)^{-11/6}$ accounts for the lengthening of the turbulent path as the elevation decreases and is the reason scintillation is most severe near the horizon. The weighting $z^{5/6}$ inside the integral gives progressively more influence to turbulence at higher altitude per unit C_n^2, a consequence of the spherical-wave Fresnel geometry. The leading constant 0.5631 follows from the $-11/3$ power law of the Kolmogorov spectrum. The model is valid only under weak scattering: the $(\sin\varepsilon)^{-11/6}$ term diverges as $\varepsilon \to 0$, and the weak-fluctuation assumption breaks down at low elevation, so the prediction is restricted to $\varepsilon \gtrsim 10°$ and is regarded as an upper bound below that floor [31, 32].

Application to a LEO pass from radiosonde data

The combination of Equations (2.6)–(2.9) defines a concrete procedure for predicting the scintillation along a LIDE pass directly from a De Bilt radiosonde ascent. The key observation is that the $C_n^2(z)$ profile is a purely meteorological quantity: it depends only on the state of the atmosphere, not on the geometry or motion of the satellite. The path integral

$$I = \int_0^\infty C_n^2(z)\, z^{5/6}\,\mathrm{d}z \tag{2.11}$$

is therefore computed *once* from the vertical column and, together with the fixed beacon wavenumber k, summarises the entire turbulence contribution of the atmosphere for that pass. This is legitimate because a single ascent is launched roughly every twelve hours whereas a LEO pass lasts only a few minutes, over which the clear-air profile is effectively frozen; if two soundings bracket the pass the profile can instead be interpolated in time.

The predicted log-amplitude standard deviation then becomes an explicit function of the instantaneous elevation alone,

$$\sigma_\chi(\varepsilon)\,[\mathrm{dB}] = 8.686\,\sqrt{0.5631\,k^{7/6}\,(\sin\varepsilon)^{-11/6}\,I}\,, \tag{2.12}$$

obtained by inserting Equation (2.11) into Equations (2.9)–(2.10). Evaluating Equation (2.12) over the elevation profile of the pass yields a predicted scintillation curve that rises towards the horizon and falls to its minimum at culmination, following the $(\sin\varepsilon)^{-11/6}$ law. This curve is the quantity compared, window by window, against the scintillation actually measured on the received power.

A subtlety specific to the LEO geometry must be emphasised. The Rytov *variance* of Equation (2.9) depends on the geometry only through the elevation and is otherwise independent of the satellite's motion; the rapid orbital slew does not change how strongly the signal scintillates. What the slew changes is the *spectrum* of the fluctuations: the effective transverse velocity at which the path cuts across the turbulent eddies is dominated by the satellite motion rather than by the wind, which shifts the scintillation corner frequency to substantially higher values than in the geostationary case [8, 33]. This is the central reason why GEO-derived scintillation *spectra* cannot be transferred directly to LEO links, even though the variance model remains applicable per window with the instantaneous elevation.

2.3.4 Solar Interference (Sun Outage)

Beyond atmospheric conditions, satellite communications can be severely disrupted by celestial bodies, most notably the Sun. A *sun outage* occurs when the apparent trajectory of the Sun places it directly behind the satellite, causing the Sun to enter the direct field-of-view (the main beam) of the ground station's receiving antenna [34] (see **Figure 2.5**).

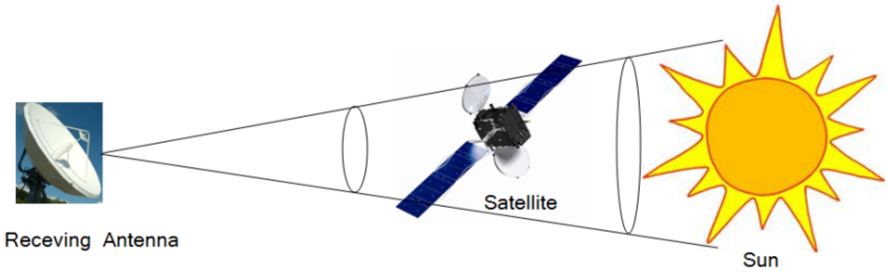

Figure 2.5: Geometric representation of Sun outage due to positioning of satellite.

The Sun is a massive broadband microwave noise source, radiating immense thermal energy across the entire RF spectrum. When the ground antenna points towards the Sun, the system's background noise temperature spikes dramatically, overwhelming the comparatively weak signal transmitted by the satellite [34]. While sun outages are highly predictable geometric events, they cause an unavoidable, temporary collapse of the Signal-to-Noise Ratio (SNR), resulting in complete communication link failures that typically last for several minutes.

This geometry was checked explicitly for the LIDE campaign. For every recorded pass, the angular separation between the Sun and the satellite, as seen from each antenna, was computed by pairing the SGP4 prediction with a solar ephemeris [35]. The Sun never entered the main beam of either terminal: the closest approach of the whole campaign was 1.8°, on the morning pass of 1 April 2026, seconds after the beacon had been switched off, and the closest approach during an active transmission was 2.0°, on 16 April 2026. These grazing transits are nonetheless visible in the data. On the van terminal, whose small aperture gives a beamwidth of a few degrees, the recorded noise floor rises by up to 0.9 dB exactly at the instant of closest approach, and the rise shrinks as the separation grows (about 0.3 dB at 3.9°). The rooftop Cobham, whose 1.5 m dish has a beamwidth of only about 0.7°, shows no measurable change at the same 1.8° separation. A full outage, with the Sun inside the main beam and the link overwhelmed, therefore never occurred during the campaign. The effect observed is limited to a brief, sub-decibel rise of the noise floor on the wide-beam terminal, too small to disturb the beacon measurements.

Figure 2.6 shows the 16 April event. In the top panel the received beacon power follows its normal profile, from acquisition to the commanded switch-off, and is not attenuated at any point: the Sun adds noise to the receiver, it does not block the signal. In the bottom panel the noise floor is drawn on a tight scale together with the angular distance between the Sun and the beam centre, that is, the satellite direction the antenna is tracking. As the separation falls towards its minimum of 2.0° the noise floor rises by about 0.8 dB and then falls back as the two directions part, mirroring the separation curve. This is the radiometric signature of the solar disc grazing the edge of the antenna beam.

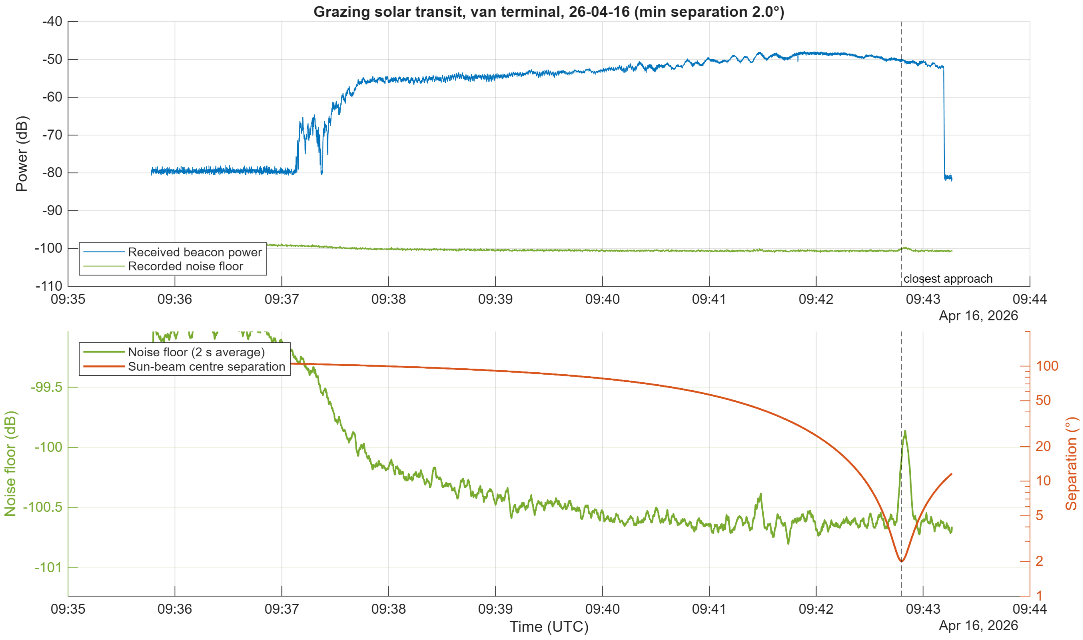

Figure 2.6: Grazing solar transit during the LIDE pass of 16 April 2026 (van terminal). Top: received beacon power and recorded noise floor over the pass; the dashed line marks the instant of closest Sun-satellite approach. Bottom: noise floor on a tight scale (2 s average, left axis) and angular distance between the Sun and the beam centre (right axis, logarithmic). The noise floor rises by about 0.8 dB exactly when the separation reaches its minimum of 2.0°, while the beacon power is unaffected: the Sun raises the system noise temperature but does not attenuate the signal.

2.4 The Doppler Effect in LEO Constellations

A fundamental characteristic that heavily differentiates LEO satellite communications from GEO links is the severe Doppler shift. The Doppler effect in telecommunications refers to the change in the observed frequency of an electromagnetic wave caused by the relative motion between the transmitter and the receiver. Because LEO satellites (such as LIDE) orbit at extremely high velocities relative to a fixed ground station, the resulting frequency offset represents a major impairment that must be rigorously compensated by the receiving hardware or software.

The absolute Doppler shift f_d observed at the ground station can be defined mathematically as:

$$f_d = -\frac{v_r}{c} f_c \tag{2.13}$$

where f_c is the nominal carrier frequency of the transmitted CW signal, c is the speed of light, and v_r is the relative radial velocity between the satellite and the ground station.

The severity of the Doppler shift is directly proportional to the carrier frequency f_c. Therefore, while L-band or S-band systems experience manageable frequency offsets, LEO systems operating in the Ka-band (e.g., above 18 GHz) face extreme Doppler shifts. In Ka-band, the maximum Doppler shift can easily exceed ± 400 kHz, with a rapid Doppler rate of change that can reach several kilohertz per second around the time of closest approach.

During a typical LEO overpass, the Doppler profile exhibits the iconic "S-curve" behaviour driven by the varying relative velocity (see **Figure 2.7**). This dynamic evolution can be broken down into three main phases:

- **Acquisition of Signal (AOS):** When the satellite rises above the horizon, its velocity vector is pointed strongly toward the ground station. The radial velocity v_r is highly negative (the distance is closing rapidly), resulting in a *positive* and maximum Doppler shift. The ground receiver must tune its local oscillator above the nominal f_c to capture the signal.

- **Time of Closest Approach (TCA):** As the satellite reaches its maximum elevation angle (zenith or the point of closest approach), its motion becomes purely tangential to the ground station. At this precise instant, the relative radial velocity v_r crosses zero. Consequently, the Doppler shift momentarily drops to exactly zero, and the received frequency matches the nominal f_c. This is also the point where the Doppler rate of change is at its maximum, creating the steepest slope in the frequency-time curve.

- **Loss of Signal (LOS):** As the satellite passes the zenith and begins to set, it moves away from the ground station. The radial velocity becomes positive, causing a *negative* Doppler shift. The received frequency drops below the nominal f_c and continues to decrease until the satellite drops below the horizon and the link is lost.

Failing to track and compensate this severe, non-linear frequency shift in Ka band leads to rapid synchronisation loss, severe inter-carrier interference in multi-carrier systems, and a complete breakdown of the communication link.

2.5 Multipath, Shadowing and Fast Fading Effects

While the FSPL and the atmospheric impairments define the macroscopic power budget of the satellite link, the local environment immediately surrounding the ground station introduces severe microscopic variations. These local propagation phenomena are primarily driven by *multipath interference*, which leads to a condition known as *fast fading* [36].

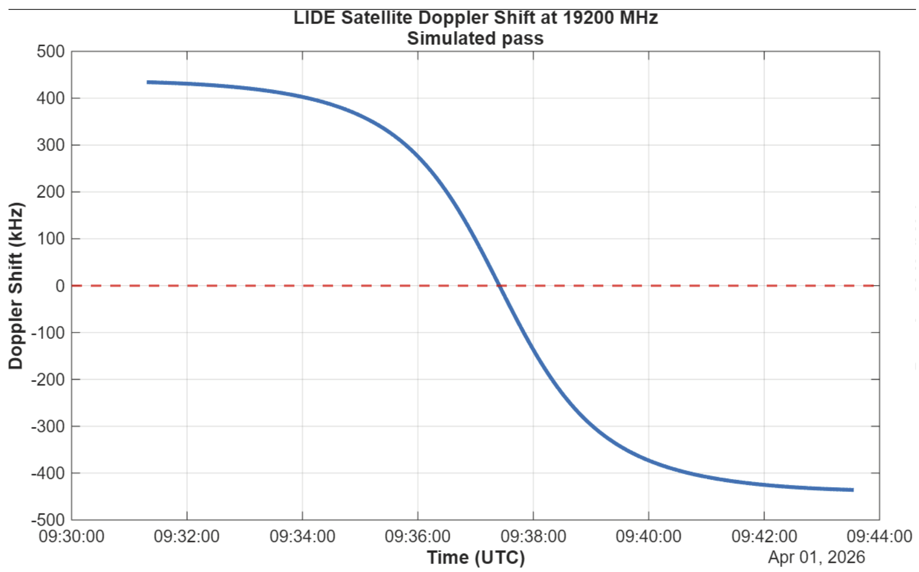

Figure 2.7: Doppler shift at 19.2 GHz of simulated LIDE pass.

In an ideal, unperturbed scenario, the communication channel between the LEO satellite and the ground station consists exclusively of a direct Line-of-Sight (LoS) signal. However, in practical deployments, such as the experimental setup involving the Kebni antenna mounted on a van or the SeaTel antenna on the ESTEC rooftop, the environment is rarely free of physical obstacles. Buildings, vehicles, metallic structures, and even the ground itself act as reflective surfaces for the highly directional Ka-band electromagnetic waves.

When the transmitted signal reaches the ground station, the receiving antenna captures not only the primary LoS ray but also multiple secondary rays that have bounced off nearby obstacles. Because these reflected rays travel slightly longer distances than the direct ray, they arrive at the receiver with different time delays and, consequently, different phases.

When the LoS signal and the multipath reflections superimpose at the receiving antenna, they interfere with one another. If the delayed rays arrive *in phase* with the direct signal, they combine constructively, causing a temporary enhancement of the received power. Conversely, if a strong reflected ray arrives 180° *out of phase* relative to the direct signal, the waves combine destructively. This destructive interference results in a sudden, sharp drop in the received signal power, an event defined as a *deep fade*.

Because the LEO satellite is moving at an extremely high velocity (changing its elevation and azimuth very rapidly), the geometry of the reflections changes continuously and dynamically. This rapid geometrical variation means that the phase relationships between the LoS ray and the multipath rays alternate from constructive to destructive within fractions of a second [37]. In the context of the ESTEC measurement campaign, these multipath effects are the cause of the anomalous and rapid power fluctuations observed in the acquired CW signal. They appear whenever the moving line of sight passes close to a reflecting or obstructing structure, so their timing is set by the local geometry of each pass rather than by the elevation alone, as the case study of Chapter 4 will show.

A related but physically distinct mechanism is *shadowing*. Whereas multipath fast fading arises from the interference of reflected rays with the direct ray, shadowing occurs when an opaque obstacle, such as a building, a mast, or vegetation, partially or totally blocks the LoS itself. The result is a comparatively slow but deep attenuation of the received power that persists for as long as the obstruction intercepts the path, frequently accompanied by diffraction at the obstacle edges. Because the Kebni terminal operates at ground level while the satellite sweeps a wide range of azimuth and elevation during a pass, the LoS can transit behind such obstacles, producing power drops that are clearly distinguishable both from the smooth FSPL trend and from the rapid multipath fluctuations. **Figure 2.8** shows the received power of a LIDE pass in which a shadowing event appears as a pronounced, sustained drop while the satellite transits behind a fixed obstacle.

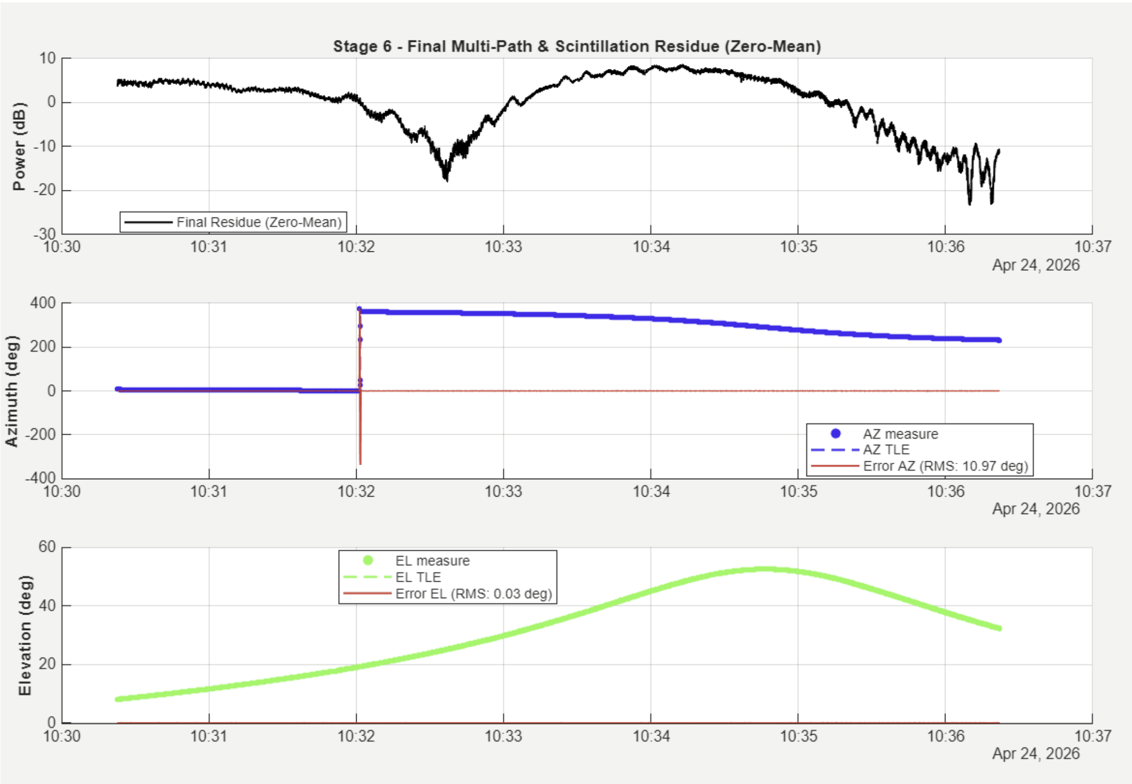

Figure 2.8: Measured received power during a LIDE pass affected by LoS shadowing.

Chapter 3

Experimental Setup and Instrumentation

3.1 ESTEC Measurement Campaign Architecture

The experimental foundation of this thesis is based on a fixed measurement campaign conducted at ESTEC in Noordwijk, the Netherlands. The primary objective of this experimental setup is to accurately capture and analyse the CW beacon signal transmitted in the Ka band by the LEO satellite LIDE.

To comprehensively evaluate the propagation channel and isolate the various impairments (such as the FSPL, atmospheric attenuation, and multipath interference), a dual-antenna architecture was deployed. This configuration allows for simultaneous data acquisition from two distinctly positioned receiving terminals, providing a comparative baseline to distinguish between satellite-induced anomalies (e.g., transmitter degradation) and ground-induced fast fading.

The architecture comprises two distinct subsystems, intended to operate synchronously during each satellite pass:

- **The Roof Installation:** The first terminal is a stabilised Cobham SeaTel 1500, a Very Small Aperture Terminal (VSAT) system. This antenna is securely mounted on the unobstructed roof of an ESTEC facility building. Thanks to its elevated position, the SeaTel antenna enjoys a clear, highly elevated LoS to the sky, ensuring that the received signal is largely immune to local ground reflections and physical blockages.

- **The Van Installation:** The second terminal is a low-profile Kebni SOTM antenna. Unlike the rigid roof system, this terminal is mounted on the roof of a mobile van. This deployment provides critical operational flexibility: the van can be positioned before each measurement to maximise the LoS for the specific azimuth and elevation trajectory of the upcoming LIDE pass. However, despite this optimisation, the van operates at ground level. The antenna therefore remains surrounded by local physical obstacles, such as adjacent buildings, parked vehicles, and metallic structures. If the positioning is insufficient for a specific pass, the low vantage point renders the signal highly susceptible to shadowing, multipath reflections and consequent fast fading.

The rationale of this dual-antenna design is that comparing the two data sets allows each impairment to be attributed to its origin according to whether it is common to both terminals or exclusive to one. An unexpected power drop recorded simultaneously by both the SeaTel and the Kebni antennas points to a systemic cause, such as LIDE transmitter degradation or a macro-atmospheric event (e.g., rain fade); conversely, rapid and severe power fluctuations observed exclusively by the ground-level Kebni antenna are attributable to fast fading caused by multipath interference from the surrounding obstacles. It must be stated from the outset, however, that in this campaign the comparison could be exploited only qualitatively: as detailed in Chapter 5, the Cobham terminal's difficulty in tracking the fast LEO satellite (Section 3.2) meant that relatively few passes were recorded cleanly by both antennas at the same time.

3.2 Antenna Systems and Hardware Specifications

To properly characterise the Ka band propagation channel during the LEO satellite passes, the measurement campaign relies on two highly specialised receiver systems. Each antenna presents distinct mechanical and radio frequency characteristics, which directly influence the signal acquisition performance and the subsequent data analysis.

3.2.1 Cobham SeaTel 1500 Reference Terminal

The primary reference receiver is the Cobham SeaTel 1500, a highly stabilised VSAT deployed on the ESTEC facility roof. This system is centred around a 1.5 m parabolic reflector, enclosed within a rigid 81 inch radome that protects the internal mechanics from wind and severe environmental stressors [38]. Because of its substantial aperture size, the SeaTel antenna delivers an exceptionally high gain of 46 dBi. This high gain translates into a very narrow beamwidth, which is fundamental to isolate the CW beacon of the LIDE satellite from surrounding thermal noise and adjacent frequency interference.

However, despite its excellent radio frequency performance, this architecture imposes significant mechanical constraints during LEO operations. The SeaTel 1500 was originally designed for maritime environments, prioritising the compensation of slow vessel wave motion rather than the extreme angular velocities of low orbiting targets. The software program tracking updates the pointing coordinates at a rate of 1 Hz, which is relatively slow for LEO dynamics. Consequently, the heavy mechanical rotators struggle to keep pace with the satellite when it passes directly overhead at the zenith, and the terminal frequently loses lock during a pass. As a result, only a minority of the recorded passes yielded clean Cobham data and the campaign is, in practice, dominated by the more agile van terminal; the quantitative consequences of this imbalance for the two-station comparison are discussed in Chapter 5.

3.2.2 Kebni SOTM Low Profile Terminal

Operating in parallel with the roof installation is the Kebni SOTM terminal, which is mounted on the mobile van. In sharp contrast to the bulky SeaTel system, the Kebni antenna is engineered specifically for vehicular mobility and rapid deployment. It features an innovative low-profile architecture built around a tuneless wideband waveguide array, measuring only 28.5 cm in height (60 cm radome diameter) and weighing less than 35 kg [39]. Its two-axis direct-drive brushless gimbal slews at up to 200°/s in azimuth and 100°/s in elevation, with accelerations as high as 700°/s^2, and maintains a pointing accuracy better than 0.2° [39]. This high slew rate is precisely what allows the Kebni to hold a solid lock on the fast-moving LIDE satellite even near the zenith and at ground level, where the line of sight is dynamically varying and frequently obscured by local obstacles. The compact design inevitably results in a lower absolute gain of 32 dBi compared to the 1.5 meter parabolic dish, making the van mounted system inherently more sensitive to the fast fading and multipath effects [39].

3.2.3 Radio Frequency Specifications Summary

Considering that the CW beacon downlink from the LIDE satellite operates at exactly 19.2 GHz, the RF chains of both systems are strictly configured to capture this frequency. The Low Noise Block (LNB) downconverters mounted on the antennas play a crucial role in amplifying the weak received signal while minimising the introduction of thermal noise, which directly impacts the antenna gain-to-noise-temperature ratio (G/T). Table 3.1 provides a comprehensive summary of the key RF specifications for both receiver systems used during the campaign.

Table 3.1: Summary of the hardware and radio frequency specifications for the Cobham and Kebni receiver systems.

Parameter	**Cobham SeaTel 1500**	**Kebni SOTM**
Antenna Aperture	1.5 m (Parabolic Reflector)	Low profile SOTM
Antenna Gain (RX)	46 dBi (@ 19.7 GHz)	32 dBi (mid band)
LNB RX RF Frequency	18.45 – 19.45 GHz	19.20 – 21.20 GHz
LNB LO Frequency	17.50 GHz	18.25 GHz
LNB Gain	55 dB (min)	60 dB (min)
Noise Figure	1.3 dB to 1.7 dB (typ)	1.6 dB max (1.3 dB typ)
G/T	23.2 dB/K @ 19.70 GHz	8.5 dB/K (typ @ mid-band)

3.3 Radio Frequency Chain and Digitisation

The journey of the signal from the satellite to the digital domain requires a robust sequence of analogue downconversion and precise digital sampling. As discussed, the CW beacon transmitted by the LIDE satellite propagates through the atmosphere at exactly 19.2 GHz. Processing a signal at such extreme frequencies directly in the digital domain is technically unfeasible with standard instrumentation, so an intermediate analogue downconversion stage is required before any software analysis can occur.

The first stage of this conversion takes place directly at the focal point of the receiving antennas, where the LNBs are mounted. The LNB is responsible for amplifying the weak incoming satellite signal and simultaneously mixing it with a Local Oscillator (LO) frequency. The Cobham SeaTel 1500 system is equipped with an LNB operating with an LO frequency of 17.50 GHz, while the Kebni SOTM antenna uses a slightly different LNB with an LO frequency of 18.25 GHz. Through analogue mixing, the original 19.2 GHz Ka band signal is heterodyned down to a much lower Intermediate Frequency (IF) falling within the L band range. Specifically, the Cobham system outputs an IF of 1.7 GHz, whereas the Kebni system outputs an IF of 950 MHz.

Once downconverted, the analogue L band signals travel to the Software Defined Radios (SDRs), where the critical digitisation phase takes place. For this experiment, the analogue-to-digital conversion is performed by two distinct SDR platforms manufactured by Ettus Research.

The SeaTel antenna chain is connected to a high performance Universal Software Radio Peripheral (USRP) X300, a wideband SDR designed for demanding signal intelligence and high throughput applications [40]. Conversely, the Kebni antenna chain is routed into a USRP N210, which provides excellent dynamic range and stability suitable for the narrower bandwidth requirements of the mobile terminal [41]. The primary task of these USRP devices is to sample the analogue IF signals, converting the continuous voltage variations into discrete In-phase and Quadrature (I/Q) complex data streams. Digitisation is performed by the on-board analogue-to-digital converters at their native rate, 200 MS/s on the USRP X300 and 100 MS/s on the USRP N210, both at 14 bit resolution [40, 41]. The on-board FPGA then digitally downconverts the band around the beacon and decimates it to the 4 MS/s complex baseband stream that is forwarded to the host.

Finally, both the USRP X300 and the USRP N210 stream the digitised I/Q samples via high speed Gigabit Ethernet connections to a Linux workstation. This dedicated machine manages the continuous data flow, temporarily stores the large volume of samples in raw binary format, and executes the bespoke software algorithms required for the peak detection and the Doppler shift calculation.

3.4 Software Architecture and Doppler Shift Computation

Following the analogue downconversion and the digitisation performed by the Ettus Research USRP platforms, the raw I/Q samples are streamed via Gigabit Ethernet to a dedicated Linux host. Here, the core data processing is executed through a customised SDR workflow developed in Python using the GNU Radio framework.

3.4.1 Peak Detection and Scalloping Loss Mitigation

The fundamental objective of the software architecture is to continuously monitor the received signal power at a frame rate of 20 Hz, while accurately tracking the extreme frequency dynamics of the LEO satellite. To achieve this, the complex baseband signal, sampled at 4 MHz, is passed through a large Fast Fourier Transform (FFT) block with a size of 32,768 bins. This transform shifts the data from the time domain into the frequency domain, providing a continuous spectral representation where the CW beacon appears as a distinct peak above the thermal noise floor.

A dedicated Python embedded block identifies the precise bin index containing the maximum log power value, allowing for real time peak detection. However, a pure CW tone subject to continuous Doppler shift rarely aligns perfectly with the centre of a discrete FFT bin [42]. As the satellite accelerates and the received frequency slides between adjacent bins, the spectral energy spreads, causing a well known digital signal processing artefact called *scalloping loss* [42]. If the software strictly recorded only the absolute maximum bin, the resulting power plot would exhibit severe artificial amplitude modulations, masking the true propagation effects.

To mitigate this mathematical artefact, the Python script implements an integrated peak power approach. Instead of observing a single bin, the algorithm defines a small integration window centred around the detected peak, extracting three adjacent bins. The logarithmic power values are temporarily converted back to linear power, summed together to capture the true total energy of the tone, and then converted back to decibels. As visual analysis during the experimental campaign confirmed, integrating the central bin with its immediate neighbours significantly cleans the data, producing a stable and less noisy power readout that is immune to scalloping drops. Extending the integration window beyond three bins yields diminishing returns, as it begins to integrate surrounding thermal noise without capturing additional signal energy, see **Figure 3.1**. Simultaneously, the software samples a separate, fixed region of the spectrum far from the beacon to estimate and log the baseline noise power.

3.4.2 Doppler Shift Computation

Alongside the absolute power measurements, the software must compute the Doppler shift induced by the LIDE satellite's extreme orbital velocity. Because the FFT block is configured to centre the spectrum around the nominal IF, the theoretical transmission frequency of the satellite corresponds exactly to the centre bin of the FFT (e.g., bin 16,384 for a 32,768 size transform).

The absolute Doppler shift is calculated directly from the FFT peak index. By subtracting the theoretical centre bin index from the actual observed peak bin index, the algorithm determines the bin offset. This offset is then multiplied by the frequency resolution of a single bin, defined as the sample rate divided by the FFT size (approximately 122 Hz per bin in this configuration), see **Figure 3.2**.

This approach provides a direct, highly accurate, and computationally lightweight measurement of the Doppler shift. The resulting data comprises a precise Coordinated Universal Time (UTC) timestamp, the integrated peak power in dB, the noise

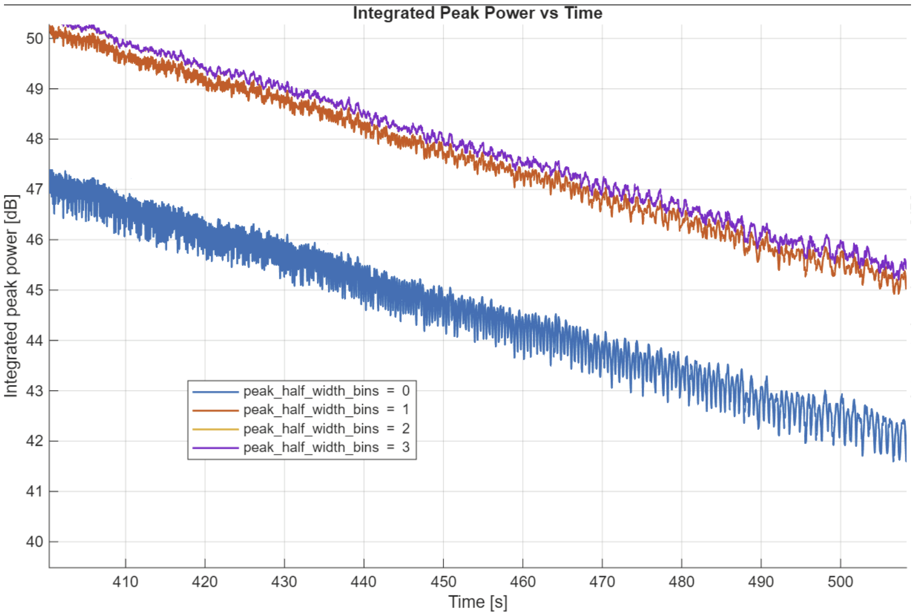

Figure 3.1: Close-up of the received power for different sizes of the bin integration window.

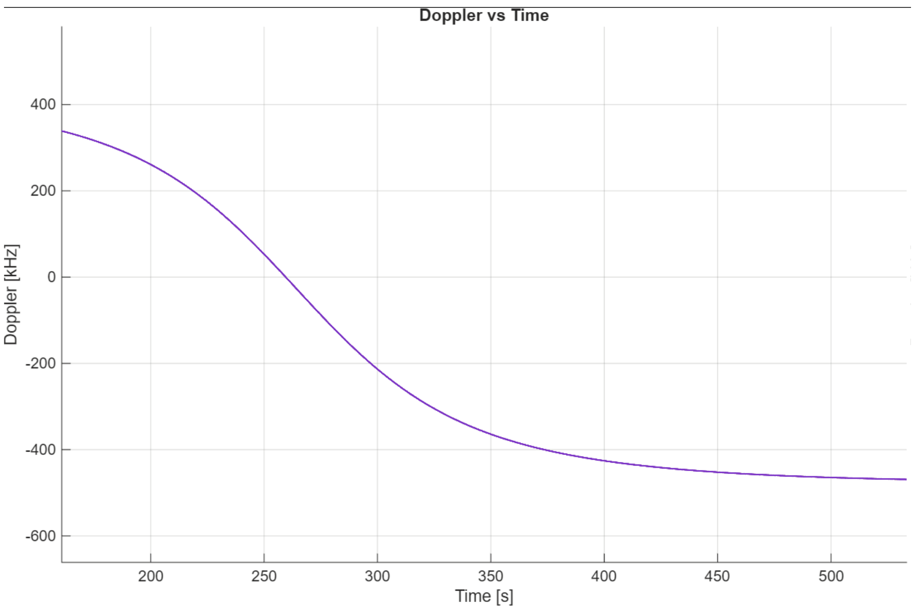

Figure 3.2: Doppler shift of the 18 March pass. The beacon was switched on at 40° of elevation and switched off at 10°.

floor power, and the Doppler shift in Hz. It is formatted into an 80 byte ASCII string and written to a Comma-Separated Values (CSV) file, ready for the post processing propagation analysis detailed in the following chapters.

The software architecture remains identical for both the Cobham SeaTel 1500 and the Kebni SOTM terminals. The only operational difference resides in the configuration of the centre frequencies, which are adjusted in the GNU Radio top block variables to match the respective 1.7 GHz and 950 MHz IFs output by the two different LNBs.

Chapter 4

Signal Processing and Results Analysis

This chapter presents the processing of the CW data acquired during the LIDE campaign and the analysis of the propagation and instrumentation effects that emerge once the deterministic terms of the link budget are removed. The raw product of the acquisition software described in Chapter 3 is, for each terminal and each pass, a time series sampled at 20 Hz containing a UTC timestamp, the integrated beacon power, the noise-floor power and the measured Doppler shift. The analysis proceeds in two conceptual phases. In the first, *deterministic* phase, the three predictable contributions to the received power, (the satellite transmit-power variation, the FSPL, and the clear air gaseous attenuation) are reconstructed from independent ancillary data and compensated. This leaves a residual referenced to a constant transmit level and free of the geometric trend. In the second, *statistical* phase, this residual is segmented to the interval in which the beacon is actually received, and what remains is characterised against a clear-sky reference (excess attenuation) and through its fast fluctuations (scintillation), while the artefacts of each terminal, the Cobham pointing cycle and the Kebni ground multipath, are isolated and attributed to their physical origin.

4.1 The Processing Pipeline

The processing is implemented as a single parametric routine applied identically to both terminals, differing only in the per-terminal configuration established in Chapter 3 (station coordinates, and the terminal-specific notch settings). For each pass the routine ingests and time-aligns the multi-source dataset. The sources are: the 20 Hz power series of the two stations, the antenna azimuth/elevation telemetry, the LIDE transmit-power telemetry, the TLE set, the KNMI surface weather record, and the De Bilt radiosonde ascent. The samples are mapped onto a common uniform time grid, and then the cascade summarised in **Figure 4.1** is executed:

1. **Deterministic compensation** (Section 4.2): the satellite power-amplifier variation, the FSPL, and the gaseous attenuation are added back to the measured power, in this order.

2. **Pass segmentation** (Section 4.3): the record is trimmed to the interval in which the beacon is actually being received, discarding the receiver-noise tails that bracket the pass.

3. **Statistical conditioning and analysis**: the compensated power feeds the final analyses, the clear-sky reference and excess attenuation (Section 4.4, aggregated over the campaign in Section 4.5), and the high-pass scintillation extraction (Section 4.8). A pointing-error mask and a signal-to-noise gate run alongside the chain to flag unusable samples.

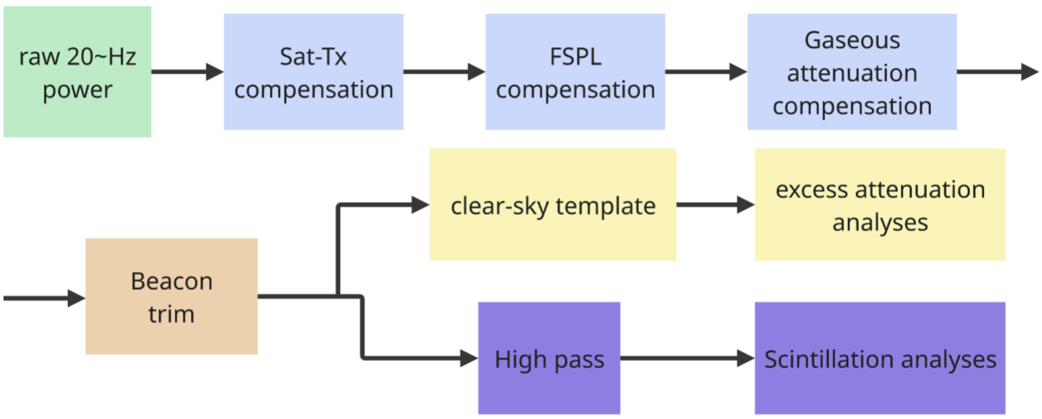

Figure 4.1: Block diagram of the processing chain applied to each terminal, from the raw 20 Hz power to the attenuation and scintillation analyses.

4.2 Deterministic Compensation of the Received Power

The measured beacon power carries, superimposed on the atmospheric and local effects of interest, three contributions whose values are known *a priori* from sources external to the propagation channel. Each of these is power the signal genuinely lost (or a transmit level it genuinely lacked), so the compensation consists of *adding the predicted term back* to the measured power, restoring the level the link would have had with a constant transmitter, no distance, and a transparent dry atmosphere. The sign convention is therefore additive throughout: a larger predicted loss lifts the compensated curve by the same amount. Denoting the measured received power by $P_{\mathrm{rx}}(t)$ and writing $\varepsilon(t)$ for the instantaneous elevation, the three stages produce P_0, P_1 and finally the compensated power P_{comp}, as detailed below. This deterministic philosophy, in which only the predictable and spatially smooth terms are removed while the intermittent ones are left to the statistical analysis, follows the preprocessing approach adopted for multi-satellite Medium Earth Orbit (MEO) and LEO campaigns [21].

4.2.1 Step 0: Satellite Power Amplifier Compensation

The first correction is not a propagation term at all but a property of the transmitter. The LIDE beacon is radiated by an on-board power amplifier (PA) whose output is not constant: over the few minutes of a pass it drifts by one to two decibels.

The PA output power is recorded in the LIDE transmit telemetry and is used in two ways. First, it delimits the pass: the beacon is considered active only while the PA output exceeds a threshold of 25 dBm, which provides the coarse temporal window of the transmission. Second, it drives the compensation. The telemetry is interpolated onto the 20 Hz grid and the per-pass maximum $P_{\mathrm{PA}}^{\max}$ is taken as the reference transmit level; the deficit relative to that maximum is then added back to the measured power,

$$P_0(t) = P_{\mathrm{rx}}(t) + \Delta_{\mathrm{PA}}(t), \qquad \Delta_{\mathrm{PA}}(t) = P_{\mathrm{PA}}^{\max} - P_{\mathrm{PA}}(t) \ \geq \ 0. \tag{4.1}$$

Because $\Delta_{\mathrm{PA}}(t)$ is non-negative by construction, the operation normalises every sample to the strongest transmit level of the pass and removes the satellite-side droop, so that the surviving slow variation can be attributed unambiguously to the ground segment and the atmosphere. **Figure 4.2** overlays the raw and PA-compensated power together with the telemetry profile.

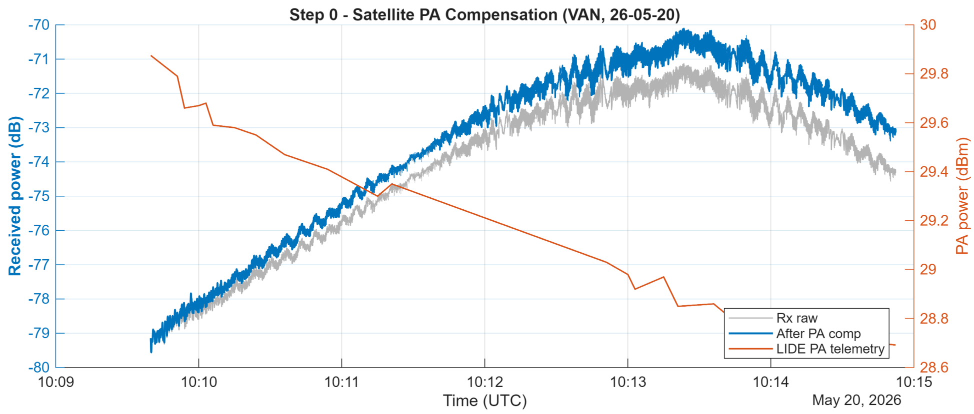

Figure 4.2: Step 0 satellite power-amplifier compensation: the received power, the LIDE transmit-power telemetry, and the compensated profile.

4.2.2 Step 1: Free Space Path Loss Compensation

The dominant feature of every raw power profile is the deep, symmetric trough imposed by the FSPL: the received power is lowest near the horizon, where the slant range is maximal, and peaks at culmination, where the range collapses to the orbital altitude. For each 20 Hz sample the satellite position is obtained by propagating the corresponding TLE set with the SGP4 model [16], the slant range $d(t)$ is computed from the known station coordinates through the geometry of Equation (2.2), and the path loss of Equation (2.1) is evaluated at the beacon frequency of 19.2 GHz and added back:

$$P_1(t) = P_0(t) + FSPL_{\mathrm{dB}}\big(d(t), f\big). \tag{4.2}$$

To keep the computation light despite the per-sample orbit propagation, the SGP4 state is evaluated on a coarse half-second grid and interpolated to the full 20 Hz resolution with a shape-preserving piecewise cubic, which reproduces the smooth range evolution at a fraction of the cost.

Figure 4.3 shows a representative pass before and after this operation. The raw curve follows the theoretical FSPL almost exactly over most of the pass, which both confirms the geometric model and validates the SGP4 pointing prediction; the compensated curve is, to first order, flat, so that the residual departures from a constant level, of the order of a few decibels, become visible and can be attributed to the non-geometric mechanisms analysed in the remainder of this chapter.

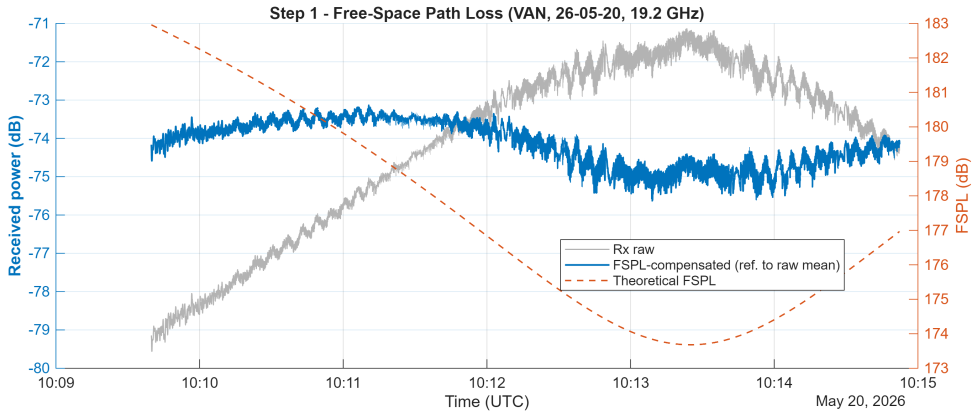

Figure 4.3: Step 1 FSPL compensation: raw received power, theoretical FSPL from the SGP4 slant range, and the resulting compensated profile.

4.2.3 Step 2: Gaseous Attenuation Compensation

The third deterministic term is the clear air absorption by atmospheric gases introduced in Section 2.3.1. It is computed from the De Bilt radiosonde ascent closest in time to the pass, using the ITU-R P.676-13 Annex 1 line-by-line model [9]: the specific attenuation $\gamma(h) = \gamma_o(h) + \gamma_w(h)$ is evaluated at every reported level from the measured pressure, temperature and humidity and integrated over height to give the zenith attenuation $A_{\text{zenith}} = \int_0^\infty \gamma(h)\,\mathrm{d}h$. The slant value is obtained from the cosecant law of Equation (2.4) and added back to the power:

$$P_{\text{comp}}(t) = P_1(t) + A_{\text{gas}}\big(\varepsilon(t)\big), \qquad A_{\text{gas}}(\varepsilon) = \frac{A_{\text{zenith}}}{\sin\varepsilon}. \tag{4.3}$$

The gaseous term scales as $1/\sin\varepsilon$ and is therefore largest near the horizon: at the low-elevation edges of the pass it amounts to of the order of one decibel, against roughly a tenth of a decibel at culmination. Adding it back consequently *lifts the low-elevation extremities* of the compensated curve while leaving the high-elevation centre almost unchanged. As argued in Section 2.3.1, the single radiosonde column is a faithful estimate of the gas the slant path actually traverses, so the term can legitimately be treated as deterministic. The output $P_{\text{comp}}(t)$ is the fully compensated power that feeds the statistical analysis. **Figure 4.4** illustrates the effect at the pass edges.

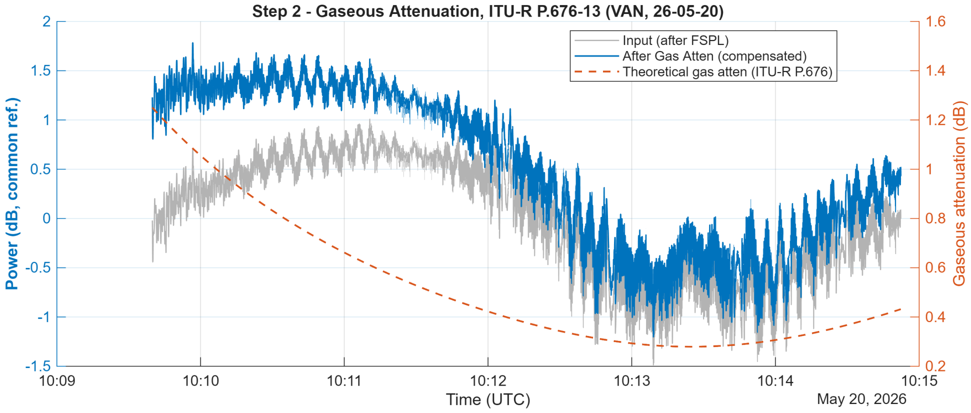

Figure 4.4: Step 2 gaseous attenuation compensation: the $1/\sin\varepsilon$ slant loss lifts the low-elevation extremities of the compensated power.

4.2.4 The Pointing-Error Quality Mask

Running in parallel with the signal chain, but not modifying the power, is a quality flag that identifies the samples acquired while the antenna was not correctly pointed at the satellite. For each terminal the actual antenna azimuth and elevation, read from the pedestal telemetry, are compared with the commanded SGP4 direction, and the three-dimensional angular separation is thresholded: samples whose pointing error exceeds a fixed bound are flagged as invalid. These flagged samples are excluded from the atmospheric statistics, where a pointing fade would otherwise masquerade as attenuation, but are retained for the dedicated study of the pointing artefacts in Section 4.6.

4.3 Pass Segmentation: Beacon-On Detection

The acquisition is started before the beacon is switched on and stopped after it is switched off, so each record is bracketed by intervals in which the receiver registers only its own thermal noise floor, some thirty decibels below the beacon. If these tails were left in, the running clear-sky reference of Section 4.4 would straddle the abrupt on/off transition and translate it into a spurious attenuation event, corrupting the statistics. The record must therefore be trimmed to the interval in which the beacon is genuinely received.

A naive amplitude threshold is inadequate, because on the Kebni terminal a deep but *fast* multipath fade momentarily drops the received level below any fixed threshold for a few seconds; a level-based detector then mistakes the fade for the loss of the pass and cuts away genuine signal around the peak. The segmentation instead exploits the one property that cleanly separates receiver noise from a fade: the noise is flat, at the floor, for many seconds, whereas a fade is a brief excursion from which the signal recovers. The detector is built around this observation and operates on the SNR:

$$\mathrm{SNR}(t) = P_{\mathrm{rx}}(t) - P_{\mathrm{noise}}(t). \tag{4.4}$$

Three running statistics. The detector never thresholds the instantaneous signal, which would be at the mercy of a single noisy sample or a fast fade. It instead uses three running statistics over three different time windows, each with one specific job:

- a **3 s running median** fixes the two reference levels, the noise floor ℓ_{off} and the beacon level ℓ_{on};
- a **5 s running mean** decides which stretches of the record are noise (this is the multipath-safe step);
- a **1 s running median** locates the precise on/off instants.

The two reference levels are taken as the extremes of the 3 s median, $\ell_{\mathrm{off}} = \min$ and $\ell_{\mathrm{on}} = \max$, which makes them independent of *how long* each state lasts. They define the contrast $\Delta = \ell_{\mathrm{on}} - \ell_{\mathrm{off}}$ and a *near-floor band* $f_{\mathrm{n}}\Delta$,

$$\tau = \ell_{\mathrm{off}} + f_{\mathrm{n}}\,\Delta, \qquad f_{\mathrm{n}} = 0.25, \tag{4.5}$$

against which the noise test is made. The default window lengths and thresholds are listed in Table 4.1.

The procedure. The beacon-on interval is then found in the following steps:

1. **Is there a step at all?** If the contrast Δ is below `step_db`, the record shows no clear on/off transition (the beacon is on throughout, or never present) and it is kept untrimmed. Otherwise the trimming proceeds.

2. **Classify the noise: the "flat for many seconds" test.** A sample is labelled *noise* when the 5 s running mean of the SNR falls inside the near-floor band,

$$\overline{\mathrm{SNR}}_{5\mathrm{s}}(t) < \tau.$$

   A fast fade cannot satisfy this: averaging a brief dip with the surrounding beacon keeps the 5 s mean well above τ. Only genuine noise stays below it, because it is flat at the floor for tens of seconds. This is what makes the detector immune to multipath.

3. **Group into sustained-noise blocks.** Consecutive noise samples are joined into runs, and only runs longer than `min_noise_s` are retained, discarding isolated noise-like samples.

4. **Snap the edges to the true transition.** Because the 5 s mean lags the real on/off step by about half its window, the provisional edges would leave $\sim$2 s of noise at the end and clip $\sim$2 s of beacon at the start. Each edge is therefore moved to the first (last) sample at which the *sharp* 1 s median exceeds τ, searched only inside the provisional beacon span. Restricting the search to that span means a mid-pass fade, which dips below τ but is bracketed by beacon on both sides, cannot move the edges; only the genuine on/off transitions can.

5. **Apply a guard.** A short margin `guard_s` is trimmed inward at each edge to discard the on/off ramp itself. The surviving index range is the beacon-on window, and all later processing (geometry, gaseous compensation, clear-sky reference, scintillation) is rebuilt on it.

As a safeguard, if no beacon segment survives between the two noise blocks, for example an all-noise record, the routine returns the untrimmed record and flags the pass as not detected, so that a failed detection never silently produces an empty pass.

Table 4.1: Default parameters of the beacon-on segmentation (20 Hz sampling).

Parameter	Role	Default
`step_db`	minimum ON–OFF contrast Δ to act on a pass	15 dB
persistence window	3 s running median for the levels $\ell_{\text{off}}, \ell_{\text{on}}$	3 s
averaging window	5 s running mean for the flat-noise test	5 s
f_n	near-floor band as a fraction of the contrast Δ	0.25
`min_noise_s`	minimum duration of a sustained-noise block	8 s
`guard_s`	inward trim at each edge (on/off ramp)	2 s

Each statistic therefore plays a distinct part: the 3 s median pins the reference levels, the 5 s mean rejects fast fades, and the 1 s median places the cut. **Figure 4.5** shows the diagnostic for a strongly multipath-affected Kebni pass, where the kept window correctly spans the beacon across two deep fades while removing the flat-noise tails at both ends.

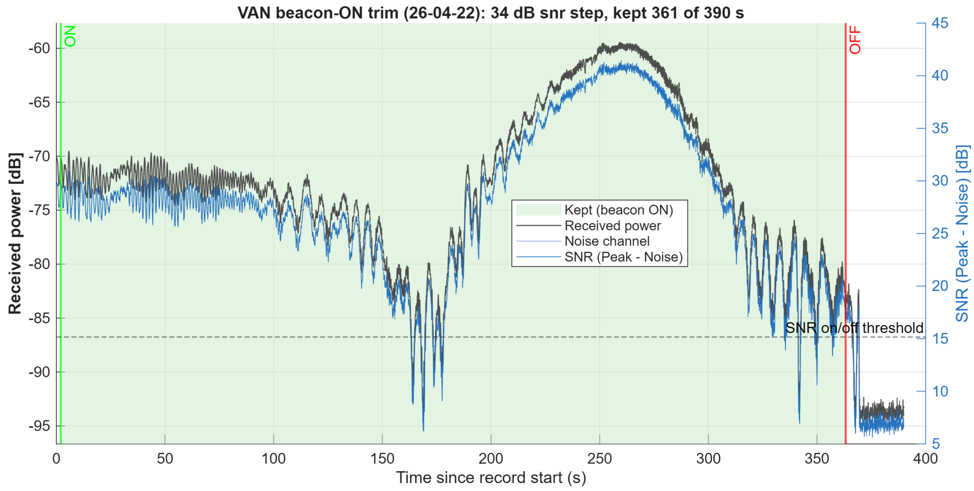

Figure 4.5: Beacon-on segmentation for a multipath-affected Kebni pass: the kept window (shaded) spans the deep fades while the flat-noise tails are removed.

4.4 Clear-Sky Reference and Excess Attenuation

After the three deterministic terms have been removed and the pass has been trimmed, the compensated power $P_{\text{comp}}(t)$ retains only a slow residual variation, on top of which sit the impairments of interest. To quantify attenuation one needs the *clear-sky* level, that is, the power the link would have had in the absence of clouds, rain and scintillation. This template is estimated directly from the data as a sliding high-percentile (95th) envelope of the compensated power, exploiting the fact that the atmosphere only ever *reduces* the signal, so that the upper envelope traces the clear sky [21]. The excess attenuation is then

$$A_{\text{exc}}(t) = P_{\text{ref}}(t) - P_{\text{comp}}(t) \;\geq\; 0, \tag{4.6}$$

clipped to non-negative values and reported only on samples that survive the SNR gate of Equation (4.4), so that deep pointing fades or dropouts are not mis-counted as atmospheric loss. It must be stressed that P_{ref} is a *per-pass, relative* reference: its zero-decibel baseline is the clear-sky level of that individual pass, not an absolute calibration. Since the elevation-dependent gas term has already been compensated in Step 2, A_{exc} in clear air is close to zero at all elevations, and departs from zero only during cloud or rain events. **Figure 4.6** shows a representative pass with its template, the excess-attenuation series, and the simultaneous elevation and azimuth, so that every feature of $A_{\text{exc}}(t)$ can be read against the geometry.

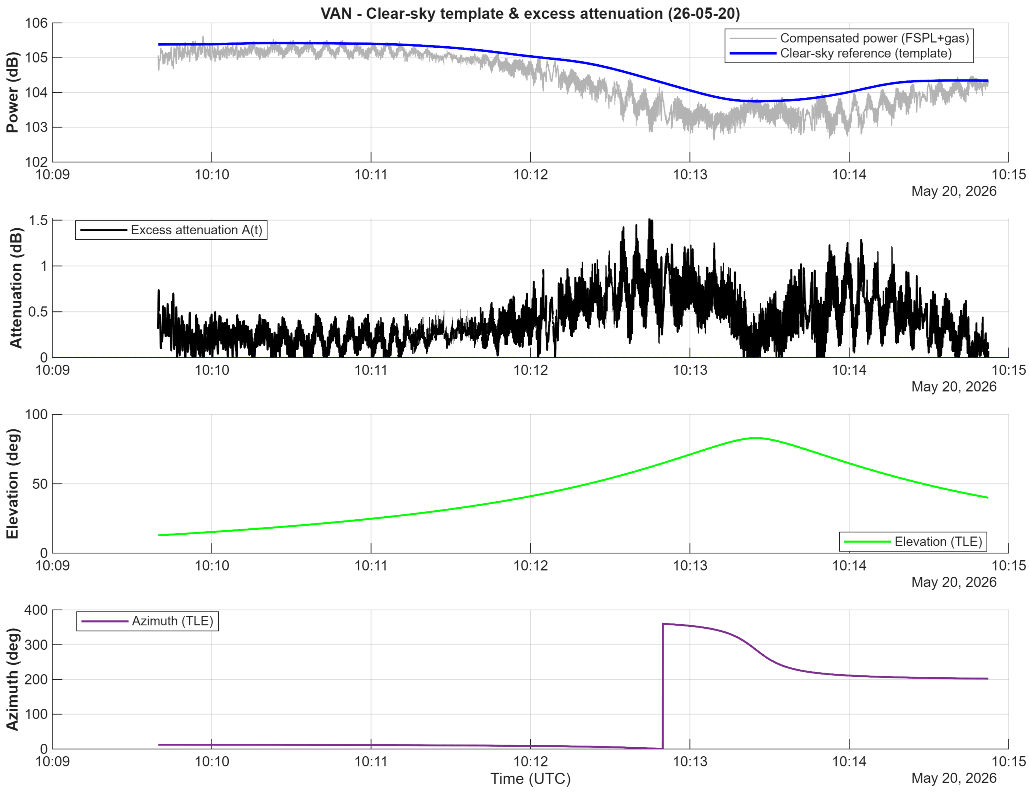

Figure 4.6: Per-pass clear-sky template and excess attenuation $A_{\text{exc}}(t)$, shown together with the pass geometry.

4.5 Campaign Attenuation Statistics

In principle, aggregating the excess attenuation of Section 4.4 over all the passes yields the standard propagation deliverable: the complementary cumulative distribution function (CCDF) $P(A_{\text{exc}} > x)$, the fraction of time a given attenuation is exceeded, which translates directly into the link margin required for a target availability. (For the campaign-wide aggregations, here and in Section 4.8, the gaseous term of Step 2 is computed with the lighter Annex 2 surface model of Equation (2.3), driven by the KNMI record, in place of the per-pass Annex 1 integration; the difference is small, slow and smooth, so it is absorbed by the per-pass clear-sky template and removed entirely by the scintillation high-pass.) **Figure 4.7** shows this distribution, conditioned on elevation class. It is included here, however, not as a propagation result but to make the limitations of the present dataset explicit: with this campaign the CCDF is *not* statistically meaningful, and the rest of this section explains why.

The clearest symptom is that the elevation ordering is unphysical. A genuine atmospheric loss must be *largest* at low elevation, where the slant path through the atmosphere is longest; yet in the figure the lowest-elevation class (10–15°) shows *less* attenuation than the higher-elevation classes. Such an inversion cannot be real, so it is a symptom of insufficient and biased sampling rather than of propagation. Two mechanisms produce it:

- **Sparse, biased low-elevation bins.** A pass sweeps through low elevation quickly, in fact the elevation rate is greatest near the horizon and falls to zero at culmination. For this reason, the antenna dwells only briefly around 10–15° (and the beacon is often not yet switched on there), leaving that bin thinly populated. It is also biased *low*. At low elevation the signal is weakest, since the path is longest, so the SNR gate, which discards any sample whose beacon-to-noise ratio falls below the 10 dB threshold to prevent noise-floor samples being read as attenuation, preferentially removes the *deepest* fades there. What survives is a subset skewed toward small attenuation. On the van, any ground multipath or shadowing event that happens to fall in this thin bin (Section 4.7) scatters the few surviving samples further. The 10–15° tail therefore rests on a handful of non-representative samples. (These effects act at the low-elevation end; at the high-elevation end the opposite bias applies to the Cobham, whose zenith slew pointing fades of Section 4.6 inject spurious attenuation into its high-elevation bins, so the two biases together exaggerate the apparent inversion.)

- **No rain classification.** As argued in Section 2.3.2, a single surface rain gauge cannot resolve the rain along an oblique, fast-moving LEO path. No rain/clear split is therefore attempted at all: every sample is pooled into a single clear air distribution with no rain tail, and the KNMI record is used only for the gaseous compensation of Step 2, not to condition the statistics.

The figure should therefore be read as evidence of what the campaign *cannot* yet support, not as a channel statistic. Too few passes were collected, and in particular too few low-elevation, well-pointed, and rainy samples. It is not possible to build a meaningful elevation-conditioned attenuation CCDF. A trustworthy distribution would require spatially resolved rain information, such as a weather-radar field instead of a point gauge, and a substantially larger and more weather-diverse set of passes. With the present data the honest conclusion is that the excess-attenuation statistic is not yet meaningful; it is reported only to delimit the scope of the campaign and to motivate the requirements of a future one.

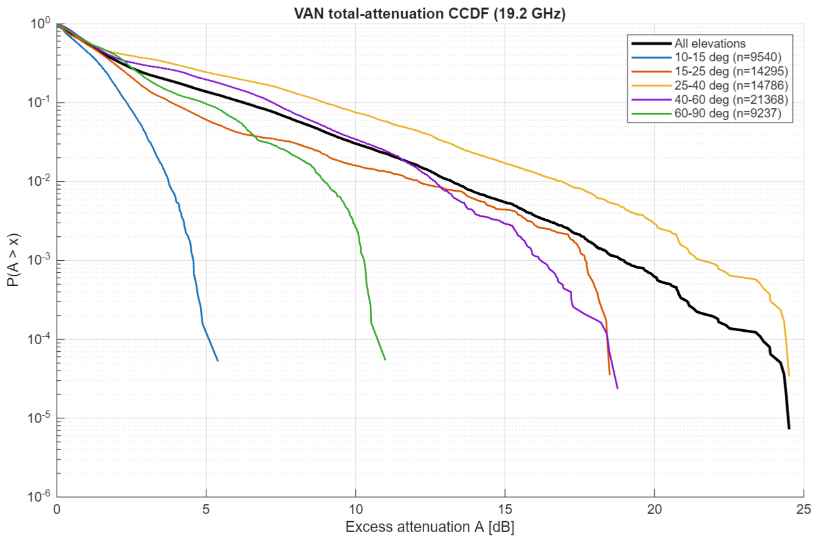

Figure 4.7: Campaign CCDF of the excess attenuation by elevation class. The unphysical ordering (lower attenuation at 10–15° than at higher elevation) and the absence of a populated rain class indicate that the dataset is too small and too clear air to yield a meaningful attenuation statistic; the figure is shown to expose this limitation, not as a result.

4.6 Cobham Tracking Anomalies and the 1 Hz Pointing Artefact

With the deterministic terms removed, the residual of the Cobham SeaTel 1500 reveals two instrumentation signatures that are a direct consequence of the mechanical and control limitations discussed in Section 3.2.

The first is a set of sharp power drops concentrated around culmination. As the satellite approaches the zenith its apparent angular velocity is maximal, and the pedestal, designed for the slow motion of a ship, cannot follow it; the antenna lags, the satellite leaves the centre of the narrow 1.5 m beam, and the received power falls by up to 8 dB before the tracking loop recovers. These dropouts are absent at low elevation, where the angular rate is modest, which confirms their kinematic origin and justifies flagging them through the pointing mask of Section 4.2.4.

The second signature is a periodic ripple at around 1 Hz, the rate at which the program-tracking software updates the commanded pointing. Because the pedestal is repositioned in discrete one-second steps rather than continuously, the beam centre oscillates about the true direction at the update frequency, modulating the received power. The effect is made explicit by the power spectral density (PSD) of the residual, which exhibits a strong line near 1 Hz together with a few harmonics (**Figure 4.8**). These lines are narrow but their exact frequencies drift from pass to pass, forming a cluster in the 0.8–1.05 Hz band with harmonics near 1.8 and 2.4 Hz, so a fixed notch bank at a single frequency cannot track them. They are instead removed by an adaptive notch that, for each pass, detects the narrow peaks standing well above the local continuum and places a zero-phase notch on each. As **Figure 4.8** shows, this suppresses the tones by roughly two orders of magnitude. Because the lines fall squarely inside the band used for the scintillation analysis, this removal is applied to the Cobham signal before the windowed σ_χ is computed, so that the instrumental ripple is not mistaken for a clear air fluctuation.

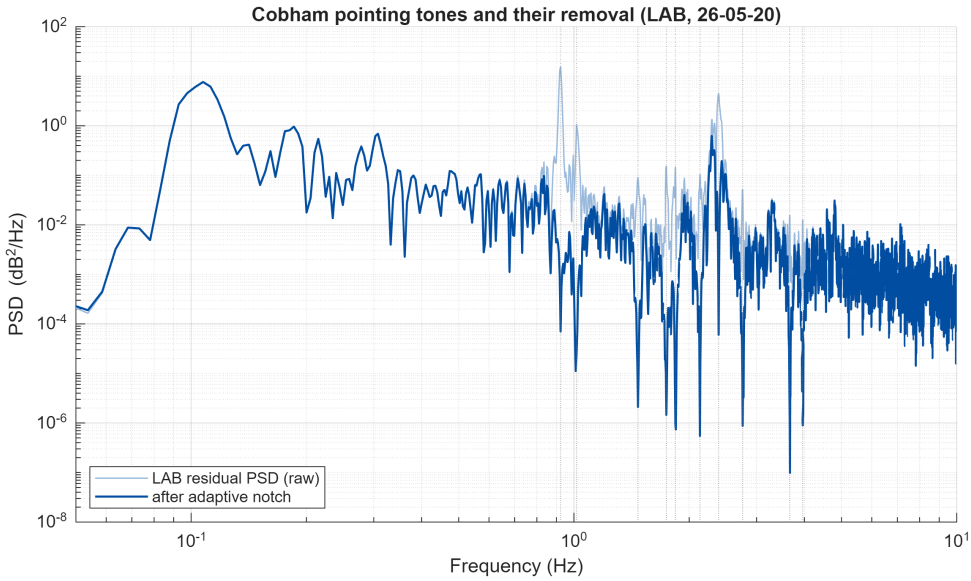

Figure 4.8: Cobham SeaTel 1500 pointing artefact: the power spectral density of the LAB residual before (light) and after (dark) the adaptive pointing-tone notch. The program-track lines cluster near 1 Hz with harmonics up to about 2.4 Hz; the notch, which finds and removes them per pass, suppresses each by roughly two orders of magnitude. The dotted vertical lines mark the tones detected on this pass.

4.7 Kebni Ground Multipath: A Clean and an Obstructed Pass

The Kebni residual is qualitatively different from the Cobham's. Thanks to its agile direct-drive gimbal and high slew rate (Section 3.2) it shows neither the zenith dropouts nor the 1 Hz line of the Cobham; what affects it instead is its ground-level environment, anticipated in Section 2.5. Operating on the van, the terminal receives, besides the direct ray, reflections from the surrounding buildings, vehicles and metallic structures, and its line of sight can be intercepted by fixed obstacles. This is not a smooth, elevation-ordered effect like the atmosphere: it is a *localised* event, tied to the geometry of a particular obstacle, that appears whenever the moving line of sight crosses that obstacle, at whatever elevation the crossing happens to occur, not preferentially at low elevation. For this reason, with only a handful of van passes there is no ensemble from which to fit a statistical fading model. What the data *does* support is a direct comparison of an unobstructed and an obstructed pass, presented here as a case study.

Two passes illustrate the two situations. On 2026-05-20 the van had a clear view of the satellite throughout: the compensated power is smooth, and the fast residual fluctuation stays at the few tenths of a decibel level characteristic of clear air scintillation and receiver noise. On 2026-04-24, by contrast, a pole stood in the line of sight, and as the satellite passed behind it, at an elevation of about 30°, the compensated power collapsed into a sharp, deep fade flanked by rapid oscillations: the pole blocked the direct ray while its edges diffracted the field, so that for a short interval the received signal was dominated by the weaker, randomly-phased scattered rays. These oscillations are consistent with a two-ray-with-Doppler channel model [37], and the whole event occurs at *mid* elevation, set by the obstacle geometry rather than by any low-elevation regime.

The shape of the 2026-04-24 received power closely resembles the multipath events documented in the FedSat Ka-band propagation campaign, where obstructions at the ground station produced the same characteristic signature [6]; the same event was already shown, from the raw received power, in **Figure 2.8**.

Figure 4.9 places the two passes side by side. The upper panels show the compensated power of the clean (2026-05-20) and the obstructed (2026-04-24) pass, the latter with the pole event clearly visible around 30° elevation; the lower panel overlays the histograms of their fast fluctuation residuals. The contrast is unambiguous: the clean pass yields a narrow distribution confined to a fraction of a decibel, whereas the obstructed pass yields a much broader distribution with a pronounced tail toward deep fades of several decibels. This clean versus obstructed comparison is the level of statement the present dataset can honestly support; a full statistical characterisation of the fading, for instance a Rician K-factor as a function of geometry, would require many more passes and is left to future work (Chapter 5).

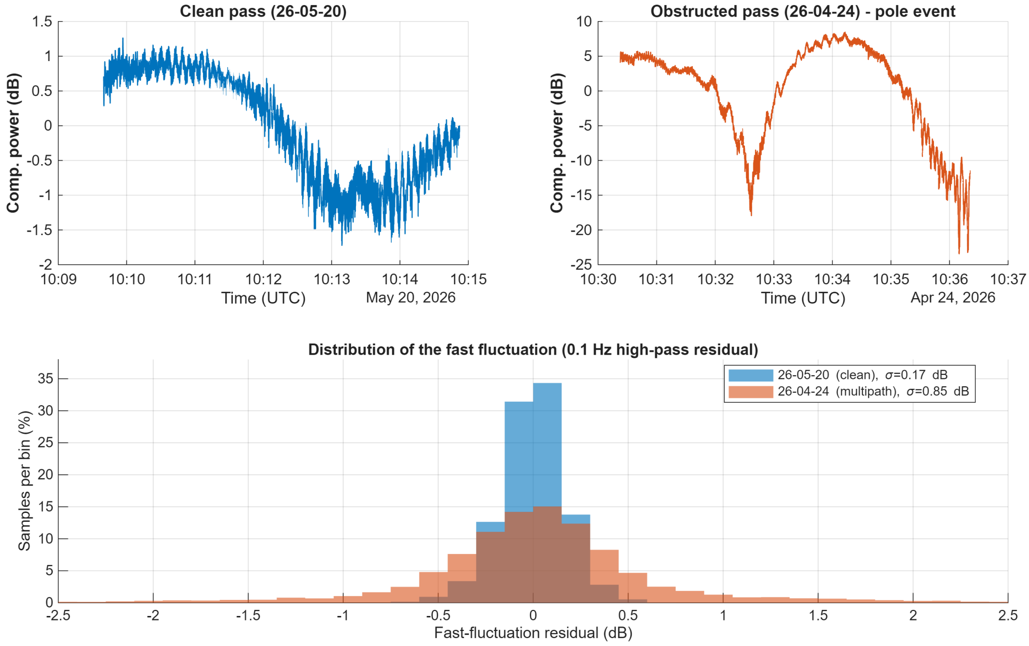

Figure 4.9: Comparison of a clean van pass (2026-05-20) and one affected by an LoS obstruction (2026-04-24). Top: compensated power relative to a sliding clear-sky baseline. The pole-diffraction multipath event, a fade of about 16 dB flanked by rapid oscillations, is visible at about 30° elevation on the obstructed pass, while the clean pass stays flat. Bottom: the distribution of the fast-fluctuation residual (0.1 Hz high-pass) of the two passes, showing the marked broadening produced by the multipath (σ rising from 0.17 to 0.85 dB).

4.8 Tropospheric Scintillation: Measurement versus RAOB Prediction

The clear air scintillation is extracted from the conditioned residual following the procedure of Section 2.3.3. The compensated power is high-pass filtered with a fourth-order zero-phase Butterworth filter at a cut-off of 0.04 Hz, chosen to lie above the slow deterministic trend (periods of tens of seconds) yet below the scintillation corner frequency (defined in Section 2.3.3), so that the slow drift is removed without eroding the turbulent fluctuation. The filtered log-amplitude is divided into short sliding windows; within each window the standard deviation σ_χ is computed and associated with the window-mean elevation. In parallel, the $C_n^2(z)$ profile is reconstructed from the De Bilt radiosonde with the Cherubini & Businger formulation of Equation (2.8), its path integral I of Equation (2.11) is evaluated once, and the predicted $\sigma_\chi(\varepsilon)$ of Equation (2.12) is obtained as a function of elevation.

Figure 4.13 compares the measured and predicted σ_χ as a function of elevation, pooled over the passes of the campaign. For each elevation class the boxes show the median and interquartile range of the per-window values, one box for the measurement and one for the radiosonde based Rytov prediction (Section 2.3.3, Equation (2.12)); both are drawn from the same population of windows. The median measurement and the prediction agree closely once the elevation exceeds about 25°: the class medians differ by less than 0.05 dB across the 25–40°, 40–60° and

60–90° bins, and the fraction of windows lying within 0.5 dB of the prediction rises from about three-quarters at 25–40° to essentially all of them above 60°. Where the weak turbulence Rytov model is expected to hold, that is for a well pointed beam at moderate to high elevation, the radiosonde based prediction therefore reproduces the measured scintillation.

The two part company only at low elevation. Below about 25° the measured median falls *below* the prediction (for instance 0.49 dB against 1.06 dB in the 10–15° bin), for two reasons introduced earlier. First, the weak-fluctuation prediction rises steeply towards the horizon, as $(\sin\varepsilon)^{-11/6}$, and overshoots once the true scintillation enters the saturation regime, past the model's $\sim$10° validity floor. Second, near the horizon the van beacon is weak and frequently obstructed, so that the SNR gate and the beacon-on trimming of Section 4.3 discard most of these samples; the low-elevation classes are consequently sparse and biased towards the cleaner windows that survived.

It matters that this comparison uses the *median* rather than the mean. The outliers visible above each box are the ground-multipath and pointing events of Section 4.7, which are worst at the intermediate elevations where the moving line of sight most often crosses an obstacle. Those rare events pull the *mean* well above the typical value and would spuriously suggest that the measurement exceeds the prediction; the median is insensitive to them and recovers the genuine, turbulence-dominated level, which is what tracks the Rytov curve. The temporal *spectrum*, however, is not conclusive from these measurements.

Figure 4.11 makes the reason explicit. It compares the temporal PSD of the log-amplitude residual of the two LEO terminals with that of the fixed-geometry Alphasat beacon, all obtained by Welch's method [43] after the same high-pass conditioning. A genuine turbulence spectrum should reveal the inertial subrange defined in Section 2.3.3, the band over which the log-amplitude PSD follows the Kolmogorov $f^{-8/3}$ law, but, as the figure shows, none of the three records exposes one.

On the low-gain van, however, this band is simply not accessible. The scintillation is weak, a fluctuation of at most a few tenths of a decibel, so above a few tenths of a hertz the spectrum flattens onto a constant, horizontal plateau. That plateau is the *receiver-noise floor*. Every receiver adds its own random electronic noise which, unlike turbulence, is white, and therefore appears on the spectrum as a flat shelf. Two facts identify this shelf as noise and not signal: it is featureless and frequency-independent, which is exactly what electronic noise looks like, and its level agrees with the receiver-noise power recorded independently on each pass. Where the turbulence signal would keep falling with frequency, it instead sinks beneath this floor, so that no inertial subrange can be seen above it. On the high-gain Cobham the carrier-to-noise ratio is far better, but its fast residual is corrupted by the antenna's own pointing tones. Once these are removed by the adaptive notch of Section 4.6 (the solid Cobham curve in the figure; the faint curve is the raw residual, with the tones), what remains is still not turbulence: a broadband pedestal jitter that sits above the receiver noise but shows no inertial-range roll-off.

Only the Alphasat beacon, recorded at a fixed geometry and with a high SNR (about 38 dB), shows a well-behaved spectrum: a low-frequency plateau up to a few tenths of a hertz, followed by a genuine roll-off that extends, clear of the noise, over more than a decade. Even this roll-off, however, is *shallower* than $f^{-8/3}$: the measured curve lies *above* the reference line throughout the resolved band. Because this is a high-SNR record, the shallowness is a real property of the spectrum and not a noise artefact, so the $-8/3$ law can be read only as an approximate tendency, and is not confirmed by these data. What the data *does* confirm, on the high-SNR windows of the geostationary Alphasat where the fluctuation stands well clear of the noise, is the *amplitude* statistic: the log-amplitude there is close to Gaussian, as weak scintillation requires. **Figure 4.10** compares the amplitude distribution of the three residuals. Only the GEO one is a clean Gaussian, with an excess kurtosis near zero; the van residual is also roughly Gaussian, but this says nothing about turbulence, since the receiver noise that dominates it is itself Gaussian; and the Cobham residual is not Gaussian at all, being shaped by the periodic pointing tones of Section 4.6 and by frequent losses of lock (its excess kurtosis is strongly positive). The amplitude test therefore supports a turbulence origin only where the SNR is high enough to expose it: on the geostationary reference, not on the LEO passes.

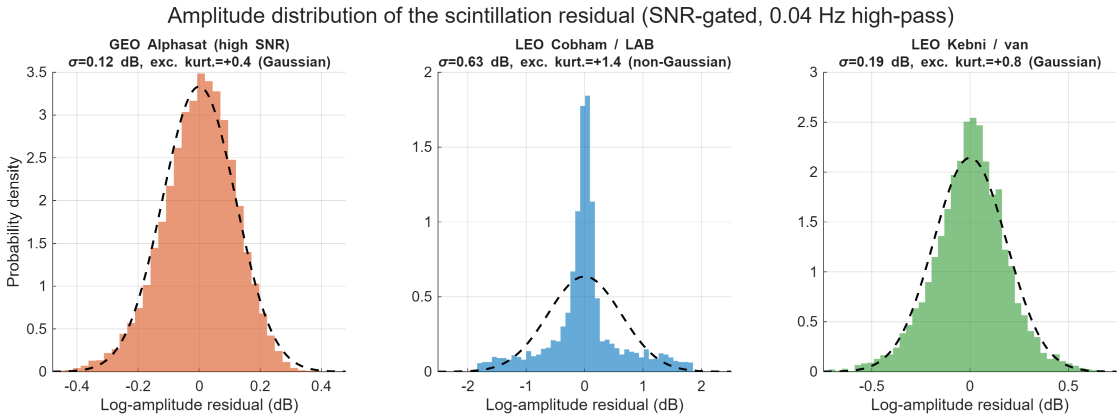

Figure 4.10: Amplitude distribution of the SNR-gated, high-pass log-amplitude residual for the three records, each with a Gaussian of the same standard deviation (dashed). Only the high-SNR geostationary Alphasat residual is a clean Gaussian (excess kurtosis ≈ 0), the signature of weak scintillation. The Kebni/van residual is also roughly Gaussian, but this merely reflects the Gaussian receiver noise that dominates it; the Cobham/LAB residual is markedly non-Gaussian (excess kurtosis clearly positive), shaped by the pointing tones and by lock losses.

It would be expected, for a LEO, that a corner-frequency shift arises from the orbital slew, which adds a large transverse velocity across the eddies and would move the whole spectrum to higher frequency while leaving its *variance* unchanged (Section 2.3.3). This remains a physically well-founded expectation that our short, noise- and artefact-limited LEO records cannot verify. Testing it would require long-duration beacon campaign. The one-year Svalbard campaign, for example, tracked the Thor 7 20.198 GHz Ka-band beacon at a nominal 2.6° elevation continuously for a year [44]. Such records accumulate the long, stationary time series, unavailable from a brief, fast LEO pass, from which a scintillation spectrum and its corner frequency could be extracted with confidence. (At Svalbard's extreme 2.6° the scintillation is itself *strong*, with summer fluctuations exceeding 10 dB, that is, the saturated regime rather than the weak turbulence sampled here, and that campaign reports scintillation *statistics* rather than a spectral slope.) Confirming the LEO corner-frequency prediction is therefore left to a dedicated acquisition of this kind (Chapter 5) [8, 33].

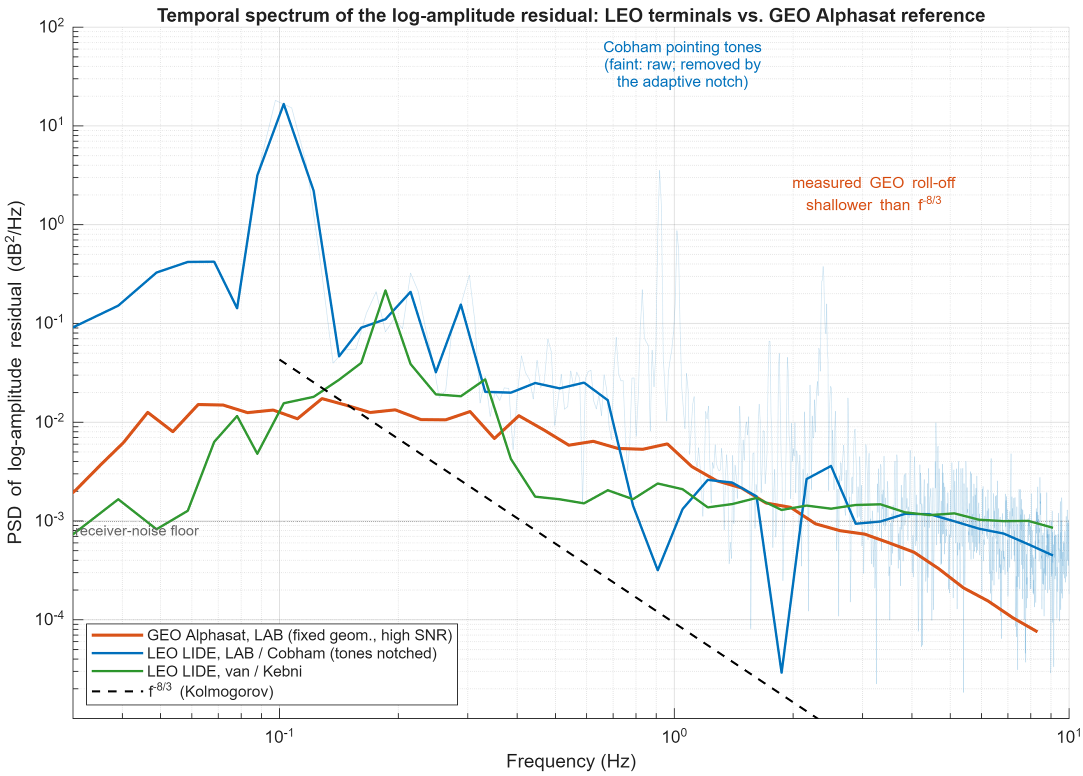

Figure 4.11: Temporal PSD of the log-amplitude residual (Welch estimate, after the 0.04 Hz high-pass used throughout the scintillation analysis) for the two LEO terminals and, as a reference, the fixed-geometry Alphasat GEO beacon; the dashed line is the Kolmogorov $f^{-8/3}$ law anchored on the GEO roll-off. Both LEO spectra fail to expose an inertial subrange: the Kebni because its scintillation is too weak and sinks into the white receiver-noise floor above about 0.3 Hz, the Cobham because its fast residual, even after the pointing tones are notched out (the faint curve is the raw residual with the tones), is broadband pedestal jitter rather than turbulence. Only the high-SNR Alphasat beacon shows a well-behaved spectrum, a low-frequency plateau followed by a roll-off, and even that roll-off is shallower than the $f^{-8/3}$ reference, remaining above it throughout the resolved band.

Because a LEO pass offers only a brief, continuously varying elevation, an independent single-elevation check was attempted on the geostationary Alphasat Aldo Paraboni beacon recorded at the same site. As noted in Section 2.3.3, only about twenty minutes of Alphasat data is available; the measured σ_χ is of the same order as the Rytov prediction at the fixed elevation, which is encouraging, but the record is far too short for a statistically meaningful correlation and is reported here only as a qualitative consistency check. **Figure 4.12** shows it: the compensated Alphasat power fluctuates by only a few tenths of a decibel about a stable clear-sky mean, and the windowed σ_χ holds steady at a median of about 0.12 dB across the twenty-minute record, a weak, stationary fluctuation of the kind expected from clear air turbulence at this fixed elevation. A conclusive validation would instead require the kind of sustained fixed-elevation acquisition already discussed above, in the spirit of the radiosonde-based Alphasat analysis of Marziani *et al.* [14].

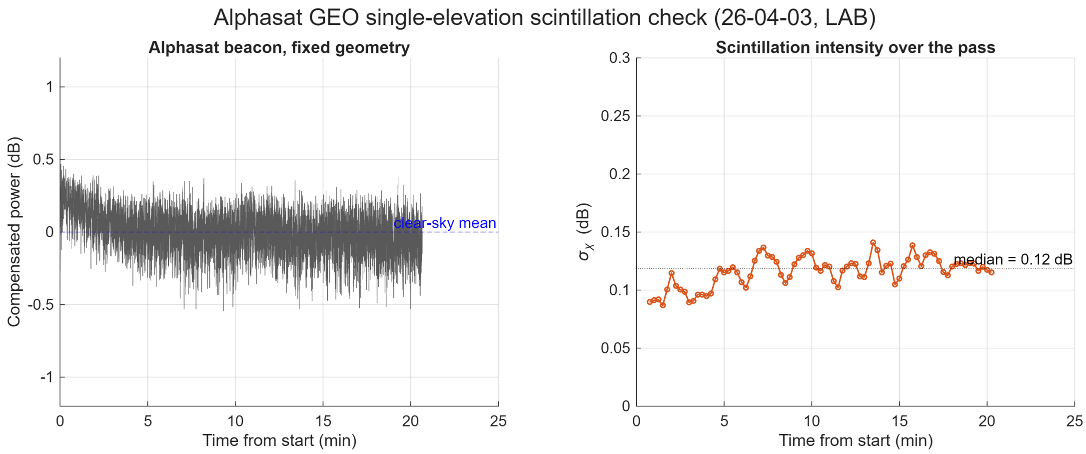

Figure 4.12: Single-elevation scintillation check on the geostationary Alphasat beacon (fixed elevation $\approx 27°$, LAB terminal). Left: the compensated received power fluctuates by a few tenths of a decibel about a constant clear-sky mean. Right: the log-amplitude standard deviation σ_χ, computed in 30 s windows, holds steady at a median of about 0.12 dB over the record, a weak, stationary clear air scintillation. The record (about twenty minutes) is far too short for a statistical validation against the Rytov prediction and is shown only as a qualitative consistency check.

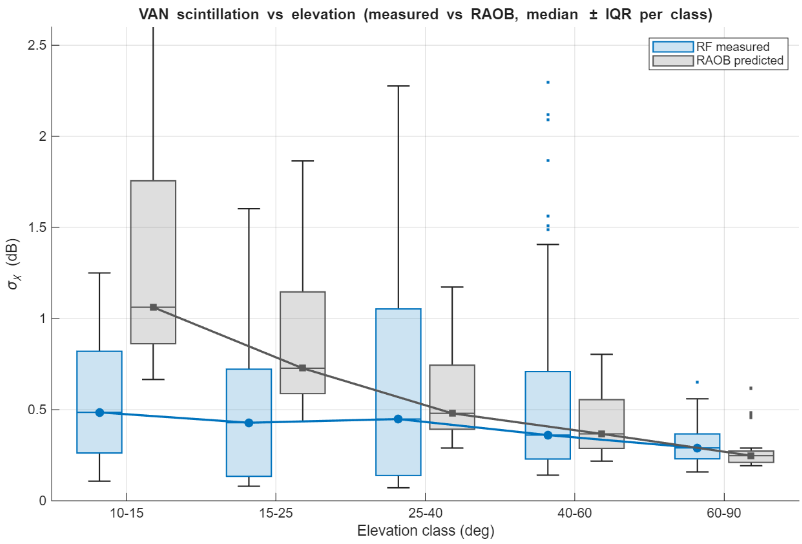

Figure 4.13: Measured versus radiosonde-predicted log-amplitude scintillation σ_χ as a function of elevation (Kebni/van). In each elevation class the boxes give the median and interquartile range of the per-window σ_χ, blue for the measurement and grey for the RAOB/Rytov prediction, computed over the same population of windows (rare multipath outliers appear as dots and extend above the axis cap). The median measurement tracks the prediction closely above about 25°: the medians agree to within 0.05 dB, and every 60–90° window falls within 0.5 dB of the prediction. It diverges only below about 25°, where the weak-fluctuation prediction overshoots in the saturation regime and the sparse low-elevation windows are thinned by the SNR gate. The solid lines join the class medians, blue for the measurement and grey for the prediction, to make the elevation trend clear.

Chapter 5

Conclusions and Future Work

This thesis set out to develop and apply a complete signal processing pipeline for Ka-band CW data acquired from a LEO satellite. Its twofold aim was to remove the deterministic terms of the link budget and to study the residual atmospheric and instrumentation effects, in particular the clear air scintillation, under the fast, continuously varying geometry that distinguishes a LEO pass from the geostationary case. The work was carried out on the LIDE beacon at 19.2 GHz, recorded at ESTEC by two co-located terminals, a high-gain Cobham SeaTel 1500 on the roof and a low-profile Kebni SOTM on a van. This closing chapter summarises what was built, what was found, where the approach is limited, and which directions the work naturally opens.

5.1 Summary of the Work

The core contribution is a physically grounded, reproducible processing chain that turns the raw 20 Hz beacon power into a clean atmospheric measurement (Chapter 4). It proceeds in two phases. In the first, *deterministic* phase, three predictable contributions are reconstructed from independent ancillary data and added back to the received power (Section 4.2): the satellite power amplifier variation, from the on-board transmit telemetry; the FSPL, from the SGP4-propagated slant range; and the clear air gaseous attenuation, integrated over the De Bilt radiosonde profile with the ITU-R P.676-13 line-by-line model. The result is a signal referenced to a constant transmit level and free of the dominant geometric trend.

A dedicated, multipath-safe segmentation step (Section 4.3) then trims each record to the interval in which the beacon is genuinely received, identifying the receiver-noise tails by the property that they are flat, at the floor, for many seconds, so that even the deep fast fades of the van terminal are preserved while the noise is removed. On the conditioned signal, two complementary analyses are run: a clear-sky reference and excess-attenuation series, and a high-pass extraction of the scintillation. In parallel, the refractive-index structure parameter $C_n^2(z)$ is reconstructed from the radiosonde and converted, through the Rytov weak-fluctuation model, into a predicted log-amplitude standard deviation that can be compared, window by window, with the value measured on the received power. The instrumentation artefacts of each terminal, the Cobham pointing cycle and the Kebni ground multipath, were isolated and attributed to their physical origin by exploiting the co-location of the two antennas. In this way the objectives set out in Chapter 1 were met: the multi-source dataset was organised and synchronised, the preprocessing chain was implemented and validated, the instrumentation artefacts were characterised, and the scintillation was compared against an independent radiosonde-based prediction.

5.2 Key Findings

The principal results of the analysis can be grouped as follows. Because the two terminals observe the same satellite, the same slant range and the same atmosphere but differ in their pedestal and their ground-level environment, an effect that appears on both stations at once must be satellite- or atmosphere-borne, whereas one confined to a single terminal is instrumental or local. In short, the part common to both records is the sky and the part that differs is the antenna. **Table 5.1** classifies each observed effect on this basis, and the paragraphs below discuss them in turn.

Table 5.1: Origin of each observed effect, classified by its appearance on the two terminals.

Effect	**Cobham**	**Kebni**	**Origin**
Slow 1–2 dB power droop	yes	yes	Satellite PA (thermal)
Zenith slew dropouts	yes	no	Cobham pedestal kinematics
1 Hz periodic ripple	yes	no	Cobham 1 Hz program track
Fast fading / shadowing	no	yes	Ground multipath (van)
Clear air scintillation	yes	yes	Tropospheric turbulence

- **The deterministic terms are removable and behave as physics predicts.** The FSPL follows the SGP4 slant range; the gaseous term scales as $1/\sin\varepsilon$, of the order of one decibel near the horizon and a tenth of a decibel at the zenith; and a slow one-to-two-decibel droop, common to both terminals, is correctly identified as the thermal drift of the satellite amplifier rather than a propagation effect, precisely because no ground-local mechanism could reproduce it simultaneously at two separated sites.

- **The dominant residual impairments are instrumental, not atmospheric.** The Cobham terminal exhibits a periodic ripple at its 1 Hz program-track update rate and sharp pointing dropouts near culmination, where the heavy maritime pedestal cannot follow the fast azimuth slew of an overhead LEO pass. The van terminal, at ground level, is instead dominated by rapid multipath fast fading and occasional shadowing produced by obstacles in its line of sight; unlike the atmosphere this is a *localised* effect, tied to the geometry of a particular obstruction and appearing at whatever elevation the line of sight happens to cross it rather than preferentially near the horizon, as the clean-versus-obstructed pass comparison of Section 4.7 made explicit. The recurrent apparent power loss around the high-elevation part of a pass was shown, from the weak correlation between the two co-located records, to be a station-specific pointing effect rather than a satellite-side one: a common cause would appear identically on both antennas, which it does not.

- **The clear air scintillation matches the prediction where the geometry is clean.** In the calm, well pointed body of a pass the measured log-amplitude standard deviation is of the order of a few tenths of a decibel, consistent with published Ka-band results. Above about 25° of elevation its median agrees with the radiosonde based Rytov prediction to within a few hundredths of a decibel. Below that, the weak fluctuation model overshoots in its saturation regime, while the SNR gate thins the surviving samples, so the comparison degrades. The amplitude statistics support a turbulence origin, since the log-amplitude of the clean high-SNR windows is close to Gaussian, as weak scintillation requires. The temporal *spectrum*, however, could not be pinned down: on the van it is limited by receiver noise, on the Cobham by pedestal jitter that persists even after the pointing tones are notched out, and even the clean, high-SNR Alphasat reference rolls off more shallowly than the Kolmogorov $f^{-8/3}$ asymptote. The corner frequency shift expected from the LEO geometry, in which the orbital slew raises the transverse velocity across the eddies, and hence the scintillation frequency content, while leaving the *variance* unchanged (Section 2.3.3), is thus a physically sound expectation that the present records cannot confirm. It is one reason a GEO-derived temporal spectrum should not be transferred directly to a LEO link, even where the variance model remains applicable per elevation [8, 33].

- **Honest limitations were identified rather than hidden.** The campaign-wide attenuation statistic proved not to be meaningful. With a small, clear air dataset that is poor in low elevation samples, the elevation-conditioned CCDF shows an unphysical ordering and is reported only to delimit the scope of the campaign. Likewise, the radiosonde based C_n^2 prediction reproduces the right order of magnitude but carries the well known uncertainty of its outer scale and empirical constants, as the geostationary Alphasat cross-check illustrated.

5.3 Limitations and Possible Improvements

Several of the limitations above point directly to concrete improvements for a future campaign.

The most consequential is the ground segment. The high-gain Cobham terminal, the natural instrument for the atmospheric measurement, is limited by a maritime pedestal that tracks the fast LEO poorly and frequently loses lock: only about five passes were recorded cleanly by both antennas at once, against roughly a dozen usable passes from the van alone. This imbalance is itself a finding, but it also blunts the co-located design. Even on those few simultaneous passes, each record is so dominated by its own terminal's pointing behaviour and local multipath, and the genuinely shared atmospheric signal so small (a few tenths of a decibel, with essentially no rain during the campaign), that a quantitative station-to-station correlation of the atmosphere could not be extracted. Equipping the high-gain terminal with a mount matched to LEO slew rates would remove the pointing artefacts, greatly increase the number of simultaneous passes, and finally enable the spatial-diversity correlation that the co-located design was meant to exploit.

On the processing side, a quantitative study of rain attenuation would require spatially resolved precipitation data, such as a weather-radar field, since a single surface gauge cannot reconstruct the rain distributed along an oblique, fast-moving slant path. For the scintillation prediction, calibrating the outer scale and empirical constants of the C_n^2 model against a controlled measurement, for instance a sustained acquisition of a geostationary beacon at a constant, low elevation in the spirit of long-term high-latitude campaigns [44], would tighten the agreement between the radiosonde-derived and RF-measured variances. Finally, a larger and more weather-diverse set of passes, together with a tighter telemetry-to-ephemeris time synchronisation (which would make the pointing mask usable for the scintillation windows), would give the statistics the significance the present dataset lacks.

5.4 Directions for Future Work

Beyond consolidating the measurement, the natural next step is to move from *characterising* the channel to *predicting* its short-term behaviour, and in particular the power outages that a LEO Ka-band link suffers. The analysis in this thesis showed that these outages have heterogeneous origins: some are deterministic and geometry-driven (the pointing loss at the zenith keyhole is, in principle, predictable from the ephemeris alone), others are quasi-stochastic but strongly conditioned on the local geometry and environment (ground multipath and shadowing, tied to specific obstacles in the line of sight and appearing at whatever elevation the obstruction is crossed), and others are atmospheric (scintillation, and rain when present). This mixture, and the fast LEO dynamics for which the classical fixed-elevation ITU-R models are poorly suited [2, 3], make the problem a good candidate for a data-driven approach.

A concrete objective is therefore a **machine-learning model for the prediction of power outages**. The rich multi-source dataset already assembled here provides the natural feature space: the predicted elevation, azimuth and their rates from the propagated TLE; the recent history of received power, SNR and Doppler; the surface and radiosonde meteorology; and a running scintillation index derived as in this work. Trained to estimate the probability that the received power will drop below a usable threshold within a short look-ahead horizon, such a model could anticipate a fade a few seconds before it occurs and drive proactive mitigation: adaptive coding and modulation, a make-before-break handover to another satellite, or, for the mobile terminal, a repositioning of the van to restore the LoS. A particularly valuable outcome would be a model that learns to *separate* the predictable, geometry-driven outages (which can be scheduled around) from the genuinely stochastic ones (which must be absorbed by margin), complementing the physics-based propagation models rather than replacing them. Feeding such a model with data from a longer campaign, and with the improved terminal and calibration described above, would turn the present offline characterisation into an operational, predictive tool.

A more targeted extension, which connects directly to that predictive goal, concerns the ground multipath itself. The clean versus obstructed comparison of Section 4.7 isolated the signature of a diffraction and multipath event unambiguously, a deep, mid-elevation fade flanked by rapid oscillations. With only a handful of van passes, however, it could be reported as a case study rather than as a statistical model. With a larger set of vehicular passes the fast fading could be characterised quantitatively through a Rician amplitude model [36].

5.5 Concluding Remarks

This work delivered a complete and physically motivated pipeline that isolates the atmosphere from the satellite and the ground instrumentation in Ka-band LEO beacon data, and applied it to characterise the propagation channel and its impairments under a dynamic geometry that classical, GEO-derived models do not fully describe. Its clearest single message is methodological: with two co-located terminals and a chain of physically justified compensations, the effects of the transmitter, the geometry, the atmosphere, and each antenna's own tracking and environment can be told apart and attributed to their origin. Its clearest limitation is equally plain: the scientific value of the atmospheric statistics is bounded not by the method but by the size of the dataset and by the tracking performance of the high-gain terminal. Addressing those two constraints, and building on the assembled multi-source dataset towards a predictive, machine-learning model of the link's outages, is the promising path that this thesis opens for the experimental study of LEO Ka-band propagation.

Bibliography

[1] G. Maral, M. Bousquet, and Z. Sun. *Satellite Communications Systems: Systems, Techniques and Technology*. Wiley, 2020. ISBN 978-1-119-38208-9. URL https://books.google.nl/books?id=kdjIDwAAQBAJ.

[2] Oltjon Kodheli, Eva Lagunas, Nicola Maturo, Shree Krishna Sharma, Bhavani Shankar, Jesus Fabian Mendoza Montoya, Juan Carlos Merlano Duncan, Danilo Spano, Symeon Chatzinotas, Steven Kisseleff, Jorge Querol, Lei Lei, Thang X. Vu, and George Goussetis. Satellite communications in the new space era: A survey and future challenges. *IEEE Communications Surveys & Tutorials*, 23 (1):70–109, 2021. doi: 10.1109/COMST.2020.3028247.

[3] Víctor Monzón Baeza, Eva Lagunas, Hayder Al-Hraishawi, and Symeon Chatzinotas. An Overview of Channel Models for NGSO Satellites. *arXiv preprint arXiv:2209.10173*, 2022.

[4] ESA. Esa artes 5g, lide. https://connectivity.esa.int/archives/projects/esa-artes-5g, 2025. Accessed: 2026-06-16.

[5] Shkelzen Cakaj. *Ground Station Design and Analysis for LEO Satellites: Analytical, Experimental and Simulation Approach*. Wiley, 1 edition, 2022. ISBN 978-1-119-89925-9. doi: 10.1002/9781119899280.

[6] T. Kostulski. *Ka Band Propagation Experiments on the Australian Low-Earth Orbit Microsatellite "FedSat"*. PhD thesis, University of Technology, Sydney (UTS), 2008. URL https://opus.lib.uts.edu.au/handle/2100/1166. PhD thesis; A/Prof. S. Reisenfeld, supervisor.

[7] C K Chen and Kevin Tseng. ROCSAT-1 Ka-Band experimental communication P/L operations. In *SpaceOps 2002 Conference*, Reston, Virginia, October 2002. American Institute of Aeronautics and Astronautics.

[8] Weiwen Liu and David G. Michelson. Effect of turbulence layer height and satellite altitude on tropospheric scintillation on ka-band earth–LEO satellite links. *IEEE Transactions on Vehicular Technology*, 59(7):3181–3192, 2010. URL https://ieeexplore.ieee.org/document/5415556.

[9] International Telecommunication Union. Attenuation by atmospheric gases and related effects. Recommendation ITU-R P.676-13, 2022. URL https://www.itu.int/rec/R-REC-P.676/en.

[10] International Telecommunication Union. Attenuation due to clouds and fog. Recommendation ITU-R P.840-9, 2023. URL https://www.itu.int/rec/R-REC-P.840/en.

[11] International Telecommunication Union. Propagation data and prediction methods required for the design of earth-space telecommunication systems. Recommendation ITU-R P.618-13, 2017. URL https://www.itu.int/rec/R-REC-P.618/en.

[12] B. Arbesser-Rastburg and G. Brussaard. Propagation research in Europe using the OLYMPUS satellite. *Proceedings of the IEEE*, 81(6):865–875, 1993. URL https://ieeexplore.ieee.org/document/257700.

[13] R. K. Crane and P. C. Robinson. ACTS propagation experiment: Rain-rate distribution observations and prediction model comparisons. *Proceedings of the IEEE*, 85(6):946–958, 1997.

[14] A. M. Marziani, C. Riva, F. Consalvi, E. R. Restuccia, and F. S. Marzano. Clear-air scintillation analysis of q-band Alphasat link at spino d'adda using radiosounding data. In *2017 11th European Conference on Antennas and Propagation (EuCAP)*, pages 15–19, 2017. doi: 10.23919/EuCAP.2017.7928865.

[15] Ruibo Wang, Mustafa A. Kishk, and Mohamed-Slim Alouini. Ultra reliable low latency routing in leo satellite constellations: A stochastic geometry approach. *IEEE Journal on Selected Areas in Communications*, 42(5):1231–1245, 2024. doi: 10.1109/JSAC.2024.3365884.

[16] F. R. Hoots and R. L. Roehrich. Spacetrack report no. 3: Models for propagation of NORAD element sets. Technical report, Aerospace Defense Center, Peterson Air Force Base, 1980.

[17] International Telecommunication Union. Nomenclature of the frequency and wavelength bands used in telecommunications. Recommendation ITU-R V.431-8, 2015. URL https://www.itu.int/rec/R-REC-V.431/en.

[18] Shkelzen Cakaj, Bexhet Kamo, Vladi Koliçi, and Olimpjon Shurdi. The range and horizon plane simulation for ground stations of low earth orbiting (leo) satellites. *Int. J. Commun. Netw. Syst. Sci.*, 4(9):585–589, 2011.

[19] Nadirah Abdul Rahim, Hanis Nabilah, and Khairayu Badron. Study of tropospheric scintillation effects in ku-band frequency for satellite communication system. *International Journal of Electrical and Computer Engineering*, 10:3136–3144, 2020. doi: 10.11591/ijece.v10i3.pp3136-3144.

[20] Khairunnisa Aziding. Rain attenuation investigation for ka band satellite transmission in tropical region. *Journal of Information Systems Engineering and Management*, 10:426–430, 2025. doi: 10.52783/jisem.v10i7s.874.

[21] Marlene Brás, Susana Mota, and Armando Rocha. Preprocessing Procedure for Propagation Measurements on MEO Satellite Links. *International Journal of Satellite Communications and Networking*, 2026. doi: 10.1002/sat.70044.

[22] International Telecommunication Union. Specific attenuation model for rain for use in prediction methods. Recommendation ITU-R P.838-3, 2005. URL https://www.itu.int/rec/R-REC-P.838/en.

[23] International Telecommunication Union. Rain height model for prediction methods. Recommendation ITU-R P.839-4, 2013. URL https://www.itu.int/rec/R-REC-P.839/en.

[24] International Telecommunication Union. Characteristics of precipitation for propagation modelling. Recommendation ITU-R P.837-7, 2017. URL https://www.itu.int/rec/R-REC-P.837/en.

[25] Weiwen Liu and D.G. Michelson. Fade slope analysis of ka-band earth-LEO satellite links using a synthetic rain field model. *IEEE Transactions on Vehicular Technology*, 58(8):4013–4022, 2009. ISSN 0018-9545, 1939-9359. doi: 10.1109/TVT.2008.2004766. URL http://ieeexplore.ieee.org/document/4608731/.

[26] A. M. Marziani. *Tropospheric scintillation and attenuation on satellite-to-Earth links at Ka and Q band: modeling, validation and experimental applications.* PhD thesis, Sapienza University of Rome, 2019.

[27] A. N. Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Doklady Akademii Nauk SSSR*, 30(4): 301–305, 1941. Reprinted in Proc. R. Soc. Lond. A 434 (1991) 9–13.

[28] V. I. Tatarski. *Wave Propagation in a Turbulent Medium.* McGraw-Hill, New York, 1961. Translated by R. A. Silverman.

[29] International Telecommunication Union. The radio refractive index: its formula and refractivity data. Recommendation ITU-R P.453-14, 2019. URL https://www.itu.int/rec/R-REC-P.453/en.

[30] T. Cherubini and S. Businger. Another look at the refractive index structure function. *Journal of Applied Meteorology and Climatology*, 52(2):498–506, 2013. doi: 10.1175/JAMC-D-11-0263.1.

[31] H. Vasseur. Prediction of tropospheric scintillation on satellite links from radiosonde data. *IEEE Transactions on Antennas and Propagation*, 47(2):293–301, 1999. doi: 10.1109/8.761069.

[32] F. S. Marzano and G. d'Auria. Model-based prediction of amplitude scintillation variance due to clear-air tropospheric turbulence on earth-satellite microwave links. *IEEE Transactions on Antennas and Propagation*, 46(10):1506–1518, 1998.

[33] Danielle Vanhoenacker-Janvier and Carlos Pereira. Impact of Tropospheric Turbulence on Non-Geostationary Communication Systems. In *XXXIII Finnish URSI Convention on Radio Science and SMARAD Seminar*, pages 13–16, 2013.

[34] Lihua Ma, Chao Hu, Jun Pei, Xinghua Ma, and Yanben Han. A novel model for predicting sun outage in satellite communication. *Earth Science Research*, 7:106, 2018. doi: 10.5539/esr.v7n1p106.

[35] M. Blanco-Muriel, D. C. Alarcón-Padilla, T. López-Moratalla, and M. Lara-Coira. Computing the solar vector. *Solar Energy*, 70(5):431–441, 2001. doi: 10.1016/S0038-092X(00)00156-0.

[36] Bernard Sklar. Rayleigh fading channels in mobile digital communication systems. *IEEE Communications Magazine*, 35:148–155, 1997. doi: 10.1109/MCOM.1997.621038.

[37] Weimin Zhang. Two-ray channel models with doppler effects for leo satellite communications. *ITU Journal on Future and Evolving Technologies*, 5:243–259, 2024. doi: 10.52953/UAMU6411.

[38] Cobham SATCOM. *Sea Tel 1500 Ku/Ka Band VSAT Antenna Installation Manual*. Cobham SATCOM, 2020. Official technical documentation detailing the 1.5-meter parabolic reflector specifications, radome dimensions, and mechanical tracking capabilities.

[39] Kebni AB. Kebni sotm low profile ka-band otm terminal product sheet. Technical report, Kebni, 2023. Datasheet providing the physical dimensions, weight, and base radio frequency performance of the low-profile vehicular antenna.

[40] Ettus Research. *USRP X300 and X310 Series*. Ettus Research, 2019.

[41] Ettus Research. *USRP N200/N210 Networked Series*. Ettus Research, 2019.

[42] Richard Lyons. Reducing fft scalloping loss errors without multiplication. *IEEE Signal Processing Magazine*, 28:112–116, 2011. doi: 10.1109/MSP.2010.939845.

[43] P. D. Welch. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, 15(2):70–73, 1967. doi: 10.1109/TAU.1967.1161901.

[44] James Nessel, Xavier Boulanger, Laurent Castanet, Torgeir Prytz, and Antonio Martellucci. Results from one year of Ka-band beacon measurements at Svalbard. In *24th Ka and Broadband Communications Conference / 36th Int. Communications Satellite Systems Conference (ICSSC)*, Niagara Falls, Canada, 2018. URL https://ntrs.nasa.gov/citations/20190001549.

Acknowledgements

I want to thank my colleagues at ESTEC and specifically the Telecommunication Laboratory keepers, Augusto Marziani, Enrico Bassetti and Giovanni Serafini that helped me with everything. Thanks to my Superviso Professor Vanelli, my co-supervisor Stefano Cioni and Head Of Section Alberto Ginesi, because they made this experience possible.

Grazie a tutti.

Declaration on the Use of Artificial Intelligence

During the preparation of this thesis, AI-based tools (Claude, Anthropic) were used to assist with several supporting tasks: drafting and refactoring MATLAB and Python analysis code, generating diagnostic plots from the recorded data, searching and summarising literature for citation purposes, and copy-editing the English prose of the manuscript for clarity, grammar and consistency of style. All the underlying measurements, the design of the processing pipeline, the physical interpretation of the results, and the scientific conclusions presented in this work are the author's own. Every AI-assisted output, whether code, figure, or text, was reviewed, verified against the recorded data or the cited literature, and edited by the author before inclusion in this thesis. The author takes full responsibility for the content of this document.