

One Researcher's Noise is Another's
Data: Using Radio Astronomy
Observations to Study Small-Scale
Structures in the Earth's Ionosphere

by

BEN BOYDE

A thesis submitted to the University of Birmingham

for the degree of

DOCTOR OF PHILOSOPHY

School of Engineering

College of Engineering and Physical Sciences

University of Birmingham

March 2025

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

The LOw Frequency ARray (LOFAR) provides a unique opportunity to observe the mid-latitude ionosphere. Although it is intended for astronomical observation, the signals from astronomical radio sources are distorted as they pass through the ionosphere. The nature of these variations in intensity and phase can provide detailed information on structures in the ionosphere at a range of scales, and the distributed network of LOFAR stations provides opportunities to characterise the propagation, spatial extent and temporal evolution of these structures. Because of the frequencies observed, wide bandwidth, spatially distributed network and high sensitivity provided by LOFAR, the observations of ionospheric phenomena made using LOFAR are highly complementary to other established techniques for observing the ionosphere.

Using just intensity measurements, it is shown that LOFAR can detect structures on spatial scales which are inaccessible with many other common techniques. A case study is presented in which LOFAR broadband intensity measurements from a bright cosmic radio source are used to infer the presence of a travelling ionospheric disturbance with a wavelength of ~ 20 km. This is achieved by comparing the observed intensity as a function of frequency and time to the modelled intensity, considering the ionosphere as a 1-dimensional thin screen containing a sinusoidal perturbation. Extensions to the model are also used to identify deviations from a simple sine wave shape. Other data sources such as Global Navigation Satellite System (GNSS) Total Electron Content (TEC) and ionosonde measurements are used to confirm that the model parameters used are physically reasonable.

The calibration process for interferometric observations provides another means of studying the ionosphere with LOFAR, as it produces a direct measure of TEC. This is

sensitive to disturbances up to an order of magnitude smaller than those detected using GNSS signals. A method is presented for identifying waves in this data and calculating their wavelengths, amplitudes and propagation directions, including robust characterisation of the uncertainties in these parameters. The range of wave parameters which can be identified in this way is shown to be complementary to previous work using networks of GNSS receivers, extending to lower amplitudes and shorter timescales. Using this analysis method with $\sim 2,700$ hours of observations, the statistics of ionospheric waves observed using LOFAR are investigated, showing a dominant population with directions tracking the neutral wind, which are associated with atmospheric gravity waves (AGWs). The precise relationship between wave direction and wind direction is found to be period dependent, indicating a previously unreported feature of AGW propagation in the thermosphere. A secondary population of field-aligned plasmaspheric disturbances is also found to be present in the observations.

Within the regular ionospheric scintillation observations made with LOFAR, a class of scintillation features referred to here as ‘U-shaped scintillation’ are identified. These are shown to occur on northward lines of sight at low elevation. They show an increased occurrence rate at night during winter relative to other seasons and times, and an elevated occurrence rate during geomagnetically active conditions. Detailed investigation of one of these events indicates that the structures responsible are strongly field-aligned and likely occur within the auroral oval or the poleward wall of the mid-latitude trough. Possible explanations are proposed for the frequency dependent behaviour of this scintillation.

Acknowledgements

This work would not have been possible without the support of many people. First among them is my supervisor Dr. Alan Wood, whose expertise, advice and above all ceaseless encouragement throughout the last three and a half years were utterly invaluable. If I have achieved anything, it is entirely thanks to him. A debt of gratitude is also due to all my co-supervisors (both official and unofficial), Dr. Gareth Dorrian, Dr. David Themens and Dr. Richard A. Fallows, for all their advice and support.

I must also thank everyone in the SERENE research group for creating such a welcoming and friendly work environment. The last three and a half years of work have been a pleasure thanks not only to the many scientific discussions and suggestions, but also the positive and supportive atmosphere, camaraderie, and regular trips to the Bratby.

I would further like to thank my family, for their support not only during the work represented in this thesis, but at every stage of my life leading up to it. Finally, I would like to thank my partner Debby, for supporting and tolerating me for the past nine years, including agreeing to move across the country for me to do this PhD.

This thesis uses data obtained with the International LOFAR Telescope (ILT). LOFAR (van Haarlem et al., 2013) is the Low Frequency Array designed and constructed by ASTRON. It has observing, data processing, and data storage facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the ILT foundation under a joint scientific policy. The ILT resources have benefited from the following recent major funding sources: CNRS-INSU, Observatoire de Paris and Université d'Orléans, France; BMBF, MIWF-NRW, MPG, Germany; Science Foundation Ireland (SFI), Department of Business, Enterprise

and Innovation (DBEI), Ireland; NWO, The Netherlands; The Science and Technology Facilities Council (STFC), UK; Ministry of Science and Higher Education, Poland. LOFAR data are available at <https://lta.lofar.eu/>.

This thesis uses data from the Juliusruh Ionosonde which is owned by the Leibniz Institute of Atmospheric Physics, Kühlungsborn. The responsible Operations Manager is Jens Mielich. The ionosonde data was accessed through the Global Ionospheric Radio Observatory (GIRO: Reinisch and Galkin, 2011) accessible at <http://giro.uml.edu/>. The manual scaling of ionosonde data used in this thesis was carried out by Dr. David Themens. The services of the Natural Environment Research Council (NERC) British Isles continuous GNSS Facility (BIGF), www.bigf.ac.uk, in providing archived GNSS data (and/or products) to this study, are gratefully acknowledged. GPS TEC data products and access through the Madrigal distributed data system are provided to the community by the Massachusetts Institute of Technology under support from US National Science Foundation grant AGS-1952737. Data for the TEC processing is provided from the following organizations: UNAVCO, Scripps Orbit and Permanent Array Center, Institut Geographique National, France, International GNSS Service, The Crustal Dynamics Data Information System (CDDIS), National Geodetic Survey, Instituto Brasileiro de Geografia e Estatística, RAMSAC CORS of Instituto Geográfico Nacional de la República Argentina, Arecibo Observatory, Low-Latitude Ionospheric Sensor Network (LISN), Topcon Positioning Systems, Inc., Canadian High Arctic Ionospheric Network, Institute of Geology and Geophysics, Chinese Academy of Sciences, China Meteorology Administration, Centro di Ricerche Sismologiche, Système d’Observation du Niveau des Eaux Littorales (SONEL), RENAG : REseau National GPS permanent, GeoNet - the official source of geological hazard information for New Zealand, GNSS Reference Networks, Finnish Meteorological Institute, SWEPOS - Sweden, Har-tebeesthoek Radio Astronomy Observatory, TrigNet Web Application, South Africa,

Australian Space Weather Services, RETE INTEGRATA NAZIONALE GPS, Estonian Land Board, Virginia Tech Center for Space Science and Engineering Research, and Korea Astronomy and Space Science Institute. The a_p , K_p and F10.7 data used in this thesis was obtained from GFZ Potsdam (<https://doi.org/10.5880/Kp.0001>). The SuperMAG collaborators (<https://supermag.jhuapl.edu/info/?page=acknowledgement>) are gratefully acknowledged for the SME data used here. The data used from the Swarm satellites was provided by the European Space Agency (ESA).

This thesis was supported by the Leverhulme Trust under Research Project Grant RPG-2020-140.

Contents

1	Introduction	1
2	Thermospheric and Ionospheric Physics	5
2.1	The Thermosphere	5
2.1.1	Atmospheric Gravity Waves	7
2.2	Ionospheric Processes	12
2.2.1	Production	13
2.2.2	Recombination	15
2.2.3	Transport	17
2.3	Ionospheric Vertical Structure	22
2.3.1	D-Region	24
2.3.2	E-Region	24
2.3.3	F-Region	25
2.4	Ionospheric Latitudinal Structure	27
2.5	Ionospheric Variability	29
2.5.1	Travelling Ionospheric Disturbances (TIDs)	30
2.5.2	Ionospheric Instabilities and Irregularities	40
3	Ionospheric Radio Propagation	54
3.1	Physics of Ionospheric Radio Propagation	54

3.2	Modelling Ionospheric Radio Propagation	58
3.2.1	Raytracing	58
3.2.2	Phase Screen Modelling	58
3.3	Ionospheric Effects on Radio Signals	60
3.3.1	Scintillation	61
3.3.2	Refractive Delays and Determination of Total Electron Content	66
3.3.3	Refractive Position Shifts	68
3.3.4	Focusing and Defocusing	69
4	LOFAR: The LOw Frequency ARray	73
4.1	Ionospheric Results from Observations of Astronomical Radio Sources .	75
4.1.1	Early Observations	76
4.1.2	Observations Using LOFAR	78
4.1.3	Observations Using Other Radio Telescopes	83
4.2	LOFAR Ionospheric Scintillation Observations	88
5	Case Study: Inferring Small-Scale TIDs with LOFAR	91
5.1	LOFAR Observations	92
5.2	Modelling	95
5.2.1	Analytic Modelling	95
5.2.2	Applying the Model	100
5.2.3	Numerical Modelling	106
5.3	Discussion	111
5.4	Conclusions	115
6	A Method for Identifying Wave Signatures in Interferometric Calibration Solutions	117
6.1	Differential TEC and Absolute TEC	118

6.2	Sample Data	121
6.3	A Method for Characterising Waves using Wavelet Analysis	126
6.3.1	Noise Estimation	129
6.3.2	Noise Fit Validation	133
6.3.3	Noise Correlation	135
6.3.4	Fit Method	138
6.4	Validation Using Synthetic Data	140
6.5	Conclusions	145
7	A Climatology of Waves Derived from Interferometric Calibration Solutions	148
7.1	Data and Processing	148
7.1.1	Data Flagging and Interpolation	150
7.1.2	Event Identification	151
7.2	Results	153
7.2.1	Wave Occurrence Rates	153
7.2.2	Wave Parameters	157
7.2.3	Propagation Directions	159
7.2.4	Wave Amplitude	170
7.3	Discussion	171
7.3.1	Occurrence Rates and Possible Wave Sources	171
7.3.2	Period Dependence of Propagation Direction and Inferred Altitudes	173
7.4	Conclusions	176
8	U-Shaped Scintillation Features	179
8.1	Occurrence Statistics	180
8.2	Possible Physical Explanations	189

8.2.1	Irregularity Scale	190
8.2.2	Large Scale Refraction	194
8.3	Case Studies: 1st April 2022 and 21st January 2023	196
8.3.1	Timescales of the Scintillation	199
8.3.2	Power Spectrum Analysis	201
8.3.3	Scintillation Power	208
8.3.4	Cross-Correlation Analysis	213
8.3.5	Simultaneous Measurements from the Swarm Satellites	221
8.4	Discussion	226
9	Conclusions: LOFAR as an Ionospheric Observatory	232
A	LOFAR Interferometric Calibration and Derivation of dTEC	238
B	dTEC Noise Correlation Estimation	241
C	PCA Bias Correction	243
D	Synthetic dTEC Data	245
E	Wave Event Characterisation	247
F	Short Period Waves: Possible Plasmaspheric Structures	252
G	Additional U-Shape Phase Screen Modelling	260
H	Cross-Correlation Fitting	269
	Bibliography	276

Chapter 1

Introduction

The ionosphere is the portion of the Earth's upper atmosphere that exists in an ionised state. It becomes significant above an altitude of around 60 km and peaks in density around 250 – 350 km in altitude. Ionisation arises primarily from solar radiation in the extreme ultraviolet (EUV) and soft X-ray wavelengths, although other mechanisms such as energetic particle precipitation are significant in specific regions and under certain conditions. The behaviour of the ionosphere is highly latitude dependent: the high-latitudes are strongly coupled to the Earth's magnetosphere and the solar wind (Dungey, 1961; Keskinen and Ossakow, 1983), and the equatorial latitudes receive the most intense solar radiation and their dynamics are strongly influenced by the quasi-horizontal magnetic field. By contrast, the mid-latitude region is typically considered to be relatively quiescent, except during large scale events such as geomagnetic storms (e.g. Basu et al., 2002).

The ionosphere is highly dynamic, as it is coupled to many different systems simultaneously. For example, ionospheric irregularities are often driven by solar variability (e.g. Afraimovich et al., 2009; Coster et al., 2017), geomagnetic activity (e.g. Ding et al., 2007; Shimeis et al., 2015), or convection in the troposphere (e.g. Lay et al., 2013), but

can also be caused by individual impulsive events such as seismic or volcanic activity (e.g. Calais and Minster, 1998; Themens et al., 2022). On larger timescales the ionosphere is influenced by diurnal and seasonal variations in solar radiation distribution and by the solar cycle, which affects both solar radiation intensity and the frequency of events such as solar flares and coronal mass ejections (Hathaway, 2015). Variations in ionospheric conditions can have significant impacts on a range of systems (Cannon et al., 2013), such as Global Navigation Satellite System (GNSS, e.g. GPS) positioning (e.g. Banville and Langley, 2013; Priyadarshi, 2015) and long-range HF communications for aviation (e.g. Jodalen et al., 2001).

The LOw Frequency ARray (LOFAR: van Haarlem et al., 2013) is a radio telescope centred in the Netherlands with stations distributed across Europe. It is intended for work such as surveying extragalactic radio sources (e.g. de Gasperin et al., 2021; Heald et al., 2015; Shimwell et al., 2017), searching for astrophysical transients (e.g. Carbone et al., 2016; Kuiack et al., 2021a) and observing highly redshifted signals from the early universe (e.g. Hothi et al., 2021; Yatawatta et al., 2013). However, the ionosphere at the frequencies observed has a significant impact on radio propagation and so distorts the signals. To a radio astronomer, this distortion is noise that should be removed (de Gasperin et al., 2019), but for ionospheric research, it makes LOFAR a continent sized ionospheric observatory (e.g. Beser et al., 2022; Dorrian et al., 2023; Fallows et al., 2020).

Compared to other common methods of observing the ionosphere, LOFAR has several key differences. Firstly, it observes in a frequency range (10 – 240 MHz) not widely used for trans-ionospheric communications such as GNSS. The lower frequencies used by LOFAR mean that ionospheric effects can be much more severe, and it also observes signals simultaneously in a broad frequency band. This broadband coverage allows LOFAR to be used to characterise the observed ionospheric effects based on their frequency

dependence and identify characteristic frequencies (e.g. focal frequencies, see Chapter 5) which provide information on the type of structures present. Using the broadband frequency coverage of LOFAR along with its distributed network of stations, it can also be possible to identify where in the ionosphere the structures causing the signal distortion are located (e.g. Fallows et al., 2020).

The high sensitivity of LOFAR as an instrument enables structures to be identified at extremely small scales and amplitudes. For example, in Chapter 5 it is shown that LOFAR can identify signatures of small-scale travelling ionospheric disturbances (TIDs), ionospheric waves with wavelengths of ~ 20 km, which are far less studied than medium- or large-scale TIDs with wavelengths of 100s to 1000s of km (e.g. Habarulema et al., 2013; Hunsucker, 1982; Oinats et al., 2016; Terra et al., 2020; Themens et al., 2022). This is achieved by comparison of the observed data to the results of a simple model of the ionosphere as a thin screen distorting the incident radio waves.

In order to mitigate ionospheric and instrumental distortions to images, an extensive calibration process is applied to astronomical data taken using LOFAR (de Gasperin et al., 2019). The calibration parameters themselves therefore include measurements of the ionospheric conditions above LOFAR, made with extremely high precision (Mevius et al., 2016). This calibration data gathered by astronomers is a large but largely unexplored source of ionospheric information. In Chapter 6, this data source is used to develop a method for identifying and quantifying ionospheric waves. Then, in Chapter 7, this method is used to develop a climatology of daytime ionospheric waves observed using LOFAR and investigate their connection to phenomena in the neutral atmosphere.

In addition to transient phenomena, LOFAR observations can also reveal large scale recurring ionospheric structures. For example, observations looking North at low elevation reveal a common feature characterised by rapid variations in signal intensity (scintillation) above some cutoff frequency with almost no variation below this frequency.

These features are discussed in detail in Chapter 8. The cutoff frequency varies with time creating a characteristic ‘U’ shape in the dynamic spectrum (intensity vs. frequency and time). These features occur preferentially during disturbed geomagnetic conditions and are tentatively associated with kilometre-scale ionospheric irregularities aligned with the geomagnetic field. Data from the Swarm satellites (Friis-Christensen et al., 2008) indicates that these irregularities are probably in the auroral ionosphere.

This work starts in Chapter 2 by summarising the basic physics of the ionosphere, including relevant details of the neutral thermosphere which exists in the same altitude range. This is followed in Chapter 3 by the theory of radio propagation in the ionosphere and approaches to modelling it. Then in Chapter 4 LOFAR’s capabilities as an ionospheric observatory are described, along with a summary of existing ionospheric results from LOFAR and other radio telescopes. Next, in Chapter 5, a case study of a single LOFAR observation is shown, indicating the presence of small-scale waves in the ionosphere and demonstrating the applicability of a very simple model for ionospheric radio propagation. Subsequently in Chapter 6 a method for extracting signatures of ionospheric waves from the calibration solutions routinely produced by astronomers is developed, including analysis of the uncertainties in these estimated wave parameters. This is followed in Chapter 7 by an application of this method to produce a climatology of daytime ionospheric waves observed by LOFAR. Chapter 8 then provides a description of a recurring type of scintillation feature in LOFAR observations, including a more in depth case study of two such events. Finally, Chapter 9 provides a summary of this work and promising directions for future study.

Chapter 2

Thermospheric and Ionospheric Physics

2.1 The Thermosphere

The thermosphere is the uppermost of the four regions of the Earth's neutral atmosphere defined based on the vertical temperature profile shown in Figure 2.1. It extends from the mesopause at roughly 80 – 85 km to around 500 km (Forbes, 2007). The mesopause is defined as the level above which atmospheric temperatures begin to increase with altitude again due to absorption of solar ultra-violet radiation. Just above this is another important level in the atmosphere at roughly 100 km, known as the turbopause (Yiğit et al., 2016). The turbopause is the level above which the atmosphere is no longer well mixed, as molecular diffusion separates out the constituent gases faster than turbulence is able to mix them together. Below this level, where turbulence dominates, the relative densities of different major species are approximately constant (except for species with significant sources and sinks within this region such as H_2O). However, above this they diverge, with densities of heavier molecular species decaying




Figure 2.1. The typical vertical temperature profile of the neutral atmosphere, with the different atmospheric layers labelled. Adapted from Kelley (2009).

much faster with altitude than those of lighter atomic species.

The motion of the neutral gas is a significant driver of ionospheric variations, and so it is useful to understand some of its characteristics. One of the most important parameters is the Brunt-Väisälä frequency, which is the characteristic frequency of oscillation of neutral gas in the vertical direction. It describes the response of the atmosphere to adiabatic vertical displacements with buoyancy acting as the restoring force, and its typical variation with altitude is shown in Figure 2.2 along with the acoustic cutoff frequency and the speed of sound. These quantities restrict the possible spectrum of waves that can propagate in the thermosphere.

On large scales, the wind patterns of the thermosphere are dominated by atmospheric tides (Vincent, 2015). Those with the largest amplitudes are the solar tides due to heating, but lunar gravitational tides are also present (Zhang and Forbes, 2014).




Figure 2.2. The typical variation of speed of sound (s , dotted line), acoustic frequency (ω_a , dashed line) and Brunt-Väisälä frequency (ω_b , solid line) with altitude. Figure reproduced from Hargreaves (1992) after I. Tolstoy and P. Pan, *J. Atmos. Sci.*, **27**, 31, 1970.

These tides directly influence various aspects of the ionosphere, such as longitudinal variation in equatorial plasma density (Immel et al., 2006) and the altitude of sporadic-E layers (Haldoupis, 2012; Pignalberi et al., 2014, see Section 2.3.2 of this chapter for an explanation of sporadic-E layers).

2.1.1 Atmospheric Gravity Waves

The most significant category of thermospheric waves for causing localised ionospheric variability are atmospheric gravity waves (AGWs). In general, these waves involve a combination of gravitational and compressional restoring forces, meaning they combine properties of both pure gravity waves and pure acoustic waves. The dispersion relation can be derived assuming that: variations of pressure and density are small (i.e. a perturbation treatment can be applied), there is no loss of energy (i.e. zero viscosity), the

effect of the Earth's rotation can be neglected, and a plane wave solution is considered (Hargreaves, 1992). The resulting relation is

$$k_z^2 = \left(1 - \frac{\omega_a^2}{\omega^2}\right) \frac{\omega^2}{s^2} - k_x^2 \left(1 - \frac{\omega_b^2}{\omega^2}\right), \quad (2.1)$$

where k_z and k_x are the vertical and horizontal components of the wavevector respectively, ω is the angular frequency of the AGW, ω_a is the acoustic frequency, ω_b is the Brunt-Väisälä frequency and s is the speed of sound. The acoustic and Brunt-Väisälä frequencies can be expressed as

$$\begin{aligned} \omega_a &= \frac{\gamma g}{2s}, \\ \omega_b &= \frac{\sqrt{\gamma - 1}g}{s}, \end{aligned} \quad (2.2)$$

where γ is the adiabatic constant and g the gravitational field strength. The dispersion relation in equation (2.1) can be reduced to the dispersion relation for a sound wave by neglecting the second term on the right hand side (equivalent to assuming $g = 0$) or to that of a Boussinesq gravity wave by ignoring the first term (which ignores density variations except when the gravitational field is involved; Spiegel and Veronis, 1960). Under the Boussinesq approximation the condition for propagation (i.e. for k_x and k_z to be real) is $\omega < \omega_b$. While this is only an approximation, it holds in general for many AGWs provided that $\omega^2 \ll s^2 k_x^2$ and is a commonly used approximation as a criterion for AGW propagation (e.g. Hoogeveen and Jacobson, 1997; Jacobson and Erickson, 1992; Lan et al., 2018). One distinctive characteristic of AGWs is that the vertical phase propagation has the opposite sign to the vertical energy and momentum flux, meaning that an AGW propagating upwards from the lower atmosphere will have a downward phase progression (Hines, 1960).

In so far as the approximation that viscosity is negligible holds, which it does for

much of the lower atmosphere, an important consequence is that the amplitude of AGWs grows rapidly with increasing altitude. This is because the neutral density ρ decays with a scale height ~ 10 km in the lower atmosphere, and in order to conserve energy the amplitude must increase $\propto \rho^{-\frac{1}{2}}$ (Hargreaves, 1992). This indicates that even waves with small amplitudes when excited in the troposphere can have significant impacts when they reach the thermosphere.

This exponential growth in amplitude with altitude breaks down either when viscosity becomes significant or when the amplitude of the wave causes the perturbation treatment used to derive equation (2.1) to break down. Viscous effects tend to introduce significant dissipation on upward propagating AGWs at altitudes around 60 – 110 km depending on wave parameters and background atmospheric conditions (Hines, 1960). This means that AGWs excited in the lower atmosphere generally cannot propagate into the thermosphere themselves, but the energy and momentum they deposit in the mesosphere and lower thermosphere can excite secondary gravity waves which can propagate upwards to higher altitudes (Vadas et al., 2003).

AGWs in the thermosphere have a range of sources, both within the thermosphere itself and lower in the atmosphere. Lower atmospheric sources include wind flow over orography (e.g. Liu et al., 2019), convective systems (e.g. Lane et al., 2001), large scale structures such as jet streams and the polar vortex (e.g. Buss et al., 2004; Sato and Yoshiki, 2008), as well as impulsive events such as tsunamis, earthquakes and volcanoes (e.g. Inchin et al., 2020; Laughman et al., 2017; Wright et al., 2022). Generation in the thermosphere can occur due to Joule heating during geomagnetic storms (e.g. Hocke and Schlegel, 1996; Hunsucker, 1982), excitation of higher-order waves due to the dissipation of waves from lower in the atmosphere (e.g. Vadas et al., 2003), and non-linear wave-wave interactions (e.g. Wüst and Bittner, 2006). Upward propagating AGWs from the lower atmosphere transfer energy and momentum upwards, and are

critical to understanding vertical coupling in the neutral atmosphere (Yigit et al., 2016), and the circulation and thermal structure of the middle atmosphere in particular (Fritts and Alexander, 2003).

AGWs are strongly influenced by the background atmospheric state. For example, the neutral winds apply a strong filtering effect which typically becomes significant around the mesosphere and lower thermosphere, such that AGWs propagating against the wind direction can propagate to higher altitudes and AGWs propagating with the wind cannot (Cowling et al., 1971). As the winds in the mesosphere and lower thermosphere are largely governed by atmospheric tides, this leads to temporal and spatial variations in the permitted spectrum of AGWs in the thermosphere (Taylor et al., 1993). As well as the neutral wind, variations in density, composition and temperature of the neutral atmosphere with altitude can lead to AGW dissipation due to kinematic viscosity or thermal diffusivity (e.g. Vadas, 2007), as well as altering the propagation characteristics of non-dissipated AGWs through changes in s , ω_a and ω_b in equation (2.1). These effects are dependent on wave parameters such as period and vertical and horizontal wavelengths, meaning that they also limit the range of wave parameters expected to be observed in the thermosphere.

Direct measurements of AGWs in the upper atmosphere are difficult to obtain, but there are some techniques which enable them to be studied. One obvious source of data is in-situ measurements from satellites in Low Earth Orbit (LEO). Most satellites do not carry dedicated instruments for measuring the neutral atmosphere, but some carry high-precision accelerometers designed for studying the Earth’s gravitational field, which can provide information on neutral densities and cross-track winds (e.g. Park et al., 2014; Xu et al., 2024). These data sets are useful, but are limited by coarse time resolution (~ 10 s) which for the orbital speeds of the satellites correspond to spatial scales of ~ 80 km (e.g. Park et al., 2014; Xu et al., 2024). This limits the range of AGWs that

can be detected to only the longer horizontal wavelengths. Because the satellite only samples along a single track, it is also limited in the range of propagation directions that can be detected, as even if density perturbations and cross-track winds are combined it is not possible to reliably identify AGWs propagating parallel or perpendicular to the satellite’s direction of motion (Xu et al., 2024). A final limitation of satellites the altitude range they cover (above ~ 250 km, and more typically above ~ 400 km, e.g. Fedorenko and Kryuchkov, 2011; Xu et al., 2024), which is generally above the altitude of most interest for ionospheric research around the altitude of peak electron density in the F-region (see Section 2.3). Satellites which reach lower altitudes are also often on highly elliptical orbits (e.g. Fedorenko and Kryuchkov, 2011), which significantly limits the coverage of measurements at these lower altitudes in terms of latitude, longitude, local time and season.

Another technique used for measuring AGWs in the upper atmosphere is airglow imaging (e.g. Hecht et al., 2004; Yue et al., 2019), which uses observations of the intensity of light emitted at certain wavelengths associated with the de-excitation of specific energised molecules in the neutral atmosphere (Ingham, 1971; Meinel, 1950). The emissions at a given wavelength typically arise from a well defined altitude range at which the relevant species is sufficiently abundant and the excitation rate is significant, meaning that the intensity is a good proxy for variations in the neutral density at that altitude (Ingham, 1971). However, these emissions are too faint to be detected from the ground during the day, which limits these observations to only studying night-time AGWs (e.g. Hecht et al., 2004; Yue et al., 2019). Another limitation is that the emission altitudes tend to cluster in the mesosphere and lower thermosphere around 80 – 100 km (Ding et al., 2004; Mangogna et al., 2016). In combination with the satellite measurements described above, this leaves a gap between roughly 100 – 250 km in which no direct measurements of AGWs are available. These altitudes are of particular interest for

AGW propagation as many theoretical models predict that this is the altitude range at which many AGWs should dissipate (e.g. Vadas, 2007).

The neutral wind and density perturbations induced by AGWs cause perturbations in the electron density of the ionosphere, which are observed as Travelling Ionospheric Disturbances (TIDs, see Section 2.5.1). Due to the difficulty of observing the neutral thermosphere directly, TIDs are often used as proxies to understand AGW propagation (e.g. Chum et al., 2021; Crowley et al., 1987; Medvedev et al., 2015; Nicolls and Heinselman, 2007; Oinats et al., 2016). To a certain extent this fills the observational gap for AGWs between $\sim 100 - 250$ km mentioned above (e.g. Lan et al., 2018; Medvedev et al., 2015; Nicolls and Heinselman, 2007). However, any conclusions that can be drawn about AGW characteristics from TID observations must be qualified by the fact that the ionospheric response to AGWs is not straightforward, and hence the observations will be biased towards those AGWs which are able to generate the most significant ionospheric perturbations (e.g. Sivakandan et al., 2021, and see Section 2.5.1).

2.2 Ionospheric Processes

The evolution of ionospheric density is encapsulated by the continuity equation for ion density N , given by

$$\frac{dN}{dt} = q - L - \vec{\nabla}(N\vec{v}), \quad (2.3)$$

where t is time, q is the ionisation rate, L is the loss rate and \vec{v} is the plasma velocity (Hargreaves, 1992). The three terms on the right represent ion production, ion loss and ion transport respectively. The details of each of these terms will be described in the following subsections. In practice, this equation must be evaluated separately for each ion species, meaning that q and L include charge transfer from one species of ion to form another as well as net changes in ionisation density.

In order to apply this equation to understanding ionospheric density profiles (addressed in Section 2.3) a steady-state is assumed. For much of the lower ionosphere in the mid-latitudes, it is also a reasonable approximation to neglect transport (Hargreaves, 1992). This is because loss rates are higher at lower altitudes due to the higher neutral density, meaning that ions are typically lost before they can be transported any significant distance. However, at higher altitudes, including the peak of plasma density, the effect of diffusion is not negligible.

2.2.1 Production

The main driver of ionisation in the Earth’s atmosphere is EUV and soft X-ray radiation from the Sun. Other sources of ionisation are significant in certain altitude and latitude ranges, such as high energy particles precipitating from the solar wind and the Earth’s magnetosphere at high latitudes. For the mid-latitudes in which LOFAR is situated, photo-ionisation is the dominant production process under almost all conditions.

The photo-ionisation process can be represented for a molecule A as



where hf represents a photon with frequency f (Hargreaves, 1992). Naturally, this requires that the photon has sufficient energy to overcome the binding energy between the outermost electron and the molecule, which is why the EUV and soft X-ray radiation is the major driver. Each neutral species has a different minimum energy, and this means that certain minor constituents of the neutral atmosphere can be major contributors to ionisation if their ionisation energy is lower than that of the major constituents. An example of this is nitric oxide which has an ionisation energy just below the Lyman-alpha emission line of the solar spectrum at 121.5nm (Laštovička, 2001), providing a

significant amount of ionisation despite making up only a tiny fraction of the total neutral atmospheric number density (Norton and Barth, 1970).

The photo-ionisation rate q in equation (2.3), for a given species and assuming monochromatic radiation can therefore be expressed as

$$q = \eta \sigma n_{neut} I, \quad (2.5)$$

where η is the ionisation efficiency, σ is the absorption cross-section, n is the neutral species density, n_{neut} is the neutral number density and I is the intensity of ionising radiation (Hargreaves, 1992). The spatial and temporal variation of the production rate is therefore determined by the variations in n_{neut} and I . The neutral number density n_{neut} decays with increasing altitude, meaning its direct contribution is to increase the production rate towards lower altitudes. However, I will decrease as the radiation propagates downwards through the atmosphere due to absorption. The variation in the production rate is then determined by the balance of these two effects.

The variation of I with altitude can be derived from the expression for change in density dI due to absorption

$$dI = -\sigma I n_{neut} ds, \quad (2.6)$$

where ds the increment along the propagation direction. This can be recast in terms of altitude z , and, by assuming an exponential decrease in n_{neut} with z characterised by a scale height H , this becomes

$$dI = -\sigma I n_0 \exp\left(\frac{-z}{H}\right) \sec(\chi) dz, \quad (2.7)$$

where n_0 is the neutral number density at $z = 0$, and χ is the solar zenith angle (SZA, i.e. the angle between the position of the Sun and the local vertical) (Chapman,

1931). Note that this change from s to z assumes a plane stratified atmosphere, and so becomes inaccurate at high SZA where the curvature of the Earth becomes significant (Chapman, 1931). The altitude profile of I is derived by integrating with respect to z , and is given by

$$I = I_{\infty} \exp \left(-\sigma H \sec(\chi) n_0 \exp \left(\frac{-z}{H} \right) \right) = I_{\infty} \exp(-\tau), \quad (2.8)$$

where I_{∞} is the intensity incident at $z = \infty$ and $\tau = \sigma H \sec(\chi) n_0 \exp \left(\frac{-z}{H} \right)$ is the optical depth.

From these simple assumed profiles for n and I , equation (2.5) can be used to calculate the resulting profile of q . Neglecting constant coefficients, this gives

$$q \propto \exp \left(\frac{-z}{H} \right) \exp(-\tau). \quad (2.9)$$

Differentiating with respect to z , it can be shown that the production rate q maximises where the optical depth τ is equal to 1 (Hargreaves, 1992). The simple picture here suggests a single maximum of q with altitude. However, accounting for the fact that the solar radiation is not monochromatic and that the various neutral species have their own scale heights and absorption coefficients, in practice this altitude at which $\tau = 1$ will be wavelength dependent.

2.2.2 Recombination

Balancing this ionisation rate in most of the ionosphere is recombination, the processes by which ions recombine with electrons. The dominant route for this is dissociative recombination, in which an ionised molecule interacts with an electron, splitting into its constituent parts in the process (Fox, 1993). For an ionised molecule AB^+ this can

be represented as



This process naturally requires an ionised molecule rather than an atom, meaning that for atomic ions recombination is a two step process in which they first must exchange charge with a neutral molecule to create a molecular ion, which then recombines with an electron. Considering an atomic ion C^+ , this charge exchange can be represented as



For the case of the ion AB^+ undergoing dissociative recombination, the loss term L in equation (2.3) is

$$L = \alpha[\text{AB}]N, \quad (2.12)$$

where α is the recombination coefficient, $[\text{AB}]$ is the number density of AB^+ ions and N is the electron density (plasma density) (Hargreaves, 1992). For an atomic ion C^+ , the rate of charge exchange L' with neutral molecules AB to produce molecular ions AB^+ is

$$L' = \beta([\text{AB}_{\text{neut}}])[C], \quad (2.13)$$

where $\beta([\text{AB}_{\text{neut}}])$ is the attachment coefficient as a function of neutral species density $[\text{AB}_{\text{neut}}]$ and $[C]$ is the number density of C^+ ions (Hargreaves, 1992). Assuming that species C^+ is the dominant ion being produced, then there are two possible situations to consider for the overall rate of ionisation loss. If $L < L'$ then the overall ionisation loss rate is restricted by equation (2.12) and is described as ‘alpha-type’, whereas if $L > L'$ then the overall rate is restricted by equation (2.13) and is described as ‘beta-type’ (Hargreaves, 1992).

2.2.3 Transport

Diffusion

Transport of the plasma from one region to another is accounted for by the remaining term in equation (2.3). Assuming that velocity \vec{v} is purely due to diffusion (i.e. no external forcing from electric fields or neutral wind) and the ionosphere is vertically stratified, the diffusion coefficient for a single species can be expressed as

$$D = \frac{k_B T}{m\nu}, \quad (2.14)$$

where k_B is Boltzmann's constant, T is the temperature, m is the mass of the particle and ν is the neutral collision frequency (Hargreaves, 1992). The velocity due to diffusion is

$$\vec{v} = -\frac{D}{N} \left(\frac{\partial N}{\partial z} \vec{\hat{z}} + \frac{1}{H_N} \right), \quad (2.15)$$

where $H_N = \frac{k_B T}{mg}$ is the scale height of the single species and $\vec{\hat{z}}$ is the vertical unit vector.

This describes the behaviour of a single species gas, but the plasma consists of both electrons and ions. The equations above suggest that they would separate due to their differing masses, but this separation produces an electric field drawing the ions and electrons together. Assuming the electron mass can be neglected relative to the ion mass, this coupling can be represented by

$$D_p = k_B \frac{T_e + T_i}{m_i \nu_i}, \quad (2.16)$$

and

$$H_p = k_B \frac{T_e + T_i}{m_i g}, \quad (2.17)$$

where D_p is the ambipolar diffusion coefficient, H_p is the plasma scale height and

subscripts i and e refer to ions and electrons respectively (Hargreaves, 1992). D_p and H_p can be substituted for D and H_N respectively in equation (2.15) to find the drift velocity due to diffusion for a plasma with a single ion species. This velocity can then be used in the third term of equation (2.3) to determine the rate of change of density due to diffusion.

Charged Particle Motion in a Magnetic Field

The motion of the ionosphere in response to external forcings is controlled by the fact that it is a magnetised plasma. Charged particles moving in the presence of a magnetic field are affected by the Lorentz force \vec{F}_B , given by

$$\vec{F}_B = q\vec{v} \times \vec{B}, \quad (2.18)$$

where q is the charge of the particle, \vec{v} is the velocity of the particle and \vec{B} is the magnetic field. For the purposes of this discussion it is assumed that \vec{B} points along the positive z -direction.

From equation (2.18) it is clear that any velocity component along z will not give rise to any Lorentz force, and so motion along the magnetic field is unconstrained (for the remainder of this discussion it is assumed that $v_z = 0$). For any motion normal to the magnetic field, the resulting force will be normal to both that motion and the magnetic field. Assuming the field can be considered spatially uniform, the resulting motion will be an orbit around the magnetic field in the x - y plane, referred to as gyromotion. By equating the Lorentz force to the force required to maintain circular motion $F = \frac{mv^2}{r}$, the radius of the orbit (the gyroradius r_B) is given by

$$r_B = \frac{m|\vec{v}|}{|q|B}, \quad (2.19)$$

where m is the particle mass. The (angular) gyrofrequency $\Omega_B = \frac{|\vec{v}|}{r_B}$ is therefore

$$\Omega_B = \frac{|q|B}{m}. \quad (2.20)$$

The gyrofrequency is independent of particle energy, and purely dependent on the species and the field strength. This means that it provides a characteristic timescale for charged particle motions in the presence of a magnetic field.

Given that the Lorentz force does not arise from motion along the z -direction, any component of an external force \vec{F}_{ext} along z will cause a response identical to if the magnetic field was not present. Hence, for the purposes of discussing the effect of the magnetic field it can be assumed without loss of generality that \vec{F}_{ext} lies along the x -direction. This means that the total force \vec{F} experienced by the particle is

$$\vec{F} = \vec{F}_B + \vec{F}_{ext} = (qBv_y + F_{ext})\hat{x} - qBv_x\hat{y}. \quad (2.21)$$

This can be expressed as a pair of coupled differential equations in v_x and v_y :

$$\begin{aligned} m \frac{dv_x}{dt} &= qBv_y + F_{ext}, \\ m \frac{dv_y}{dt} &= -qBv_x. \end{aligned} \quad (2.22)$$

Differentiating the second of these with respect to t provides a second expression for $\frac{dv_x}{dt}$, and it can therefore be shown that

$$\frac{d^2v_y}{dt^2} = -\Omega_B^2 v_y - \frac{qB}{m} F_{ext}. \quad (2.23)$$

Assuming a solution consisting of the sinusoidal motion corresponding to the gyromo-

tion plus a constant term, this constant term v_d along the y -direction is

$$\begin{aligned} v_d &= \frac{-F_{ext}}{qB}, \\ \vec{v}_d &= \frac{\vec{F}_{ext} \times \vec{B}}{q|\vec{B}|^2}, \end{aligned} \tag{2.24}$$

where the second line applies in the general case for arbitrary orientations of \vec{F}_{ext} and \vec{B} (the solution to equation (2.22) for v_x is identical to the case in the absence of \vec{F}_{ext} and hence there is no drift).

In the ionosphere, the two main external forcings to consider are neutral winds and electric fields. The above discussion has ignored any effect of collisions, which in principle would prevent the neutral wind having any effect on the plasma. However, provided the frequency of collisions is much smaller than the gyrofrequency then the collisions do not interrupt the gyromotion, and so this collisionless approach can be used to consider the effect of the neutral wind. The following subsection explores this in more detail and considers the impact of cases where this is not satisfied.

For the neutral wind, the force applied to the electrons and ions has the same direction, but in general a different magnitude due to the difference in collision frequencies. Due to the factor of q in the denominator of equation (2.24), the resulting motion of electrons and ions will be in opposite directions, meaning that the wind induces a current to flow along the $\vec{F}_{ext} \times \vec{B}$ direction. Conversely, for an electric field no current is generated as the force applied to electrons and ions is in opposite directions ($\vec{F}_{ext} = q\vec{E}$). The drift velocity resulting from an applied electric field \vec{E} , obtained from equation (2.24), is

$$\vec{v}_d = \frac{\vec{E} \times \vec{B}}{|\vec{B}|^2}. \tag{2.25}$$

As this is only dependent on the electric and magnetic fields, and not on the particles' charges or masses, this drift velocity is the same for the electrons and all ion species

and represents a bulk drift of the plasma.

Effect of Collisions

In the ionosphere, neglecting collisions in this way is typically not appropriate. This is especially true at lower altitudes as the higher neutral density leads to a higher collision frequency. The effect of collisions is to add an additional force on the charged particles, referred to in this context as neutral drag F_{drag}^{\rightarrow} , given by

$$F_{drag}^{\rightarrow} = -m\nu\vec{v}, \quad (2.26)$$

where ν is the collision frequency (in general this is different for electrons and ions), and \vec{v} here is the velocity of the charged particle relative to the neutrals (Hargreaves, 1992). Clearly, this acts at a right angle to the Lorentz force (ignoring any motion along \vec{B}). Hence, if F_{drag}^{\rightarrow} dominates over \vec{F}_B , the drift motion induced by F_{ext}^{\rightarrow} will occur in the direction of F_{ext}^{\rightarrow} , as this is normal to the drift expected in the absence of F_{drag}^{\rightarrow} (see Hargreaves (1992) for a more thorough explanation).

The relative significance of collisions to gyromotion in determining the overall motion can be expressed by the ratio $\frac{\Omega_B}{\nu}$. The vertical variation of this ratio for both electrons and ions is illustrated in Figure 2.3. If the ratio is much greater than one, then many gyro orbits occur between each collision and so the collisions are negligible. On the other hand, if it is much less than one, then collisions are expected to occur within a fraction of a gyro orbit and the Lorentz force becomes negligible.

When collisions are dominant, the motions induced by electric fields tend to be suppressed, while the motions induced by winds tend to be stronger as the force applied by the wind is the effect of these collisions. This creates unique behaviour in the altitude range between roughly 80–120 km where the electrons experience few collisions




Figure 2.3. A typical vertical profile of the ratio between gyrofrequency (denoted here by a lower case ω) and collision frequency for electrons and ions. θ refers to the angle between the force \vec{F}_{ext} and the induced \vec{v}_d . Figure reproduced from Hargreaves (1992).

while the ion motion is still collision dominated. In this region, electrons and ions essentially move in response to different external forcing, although their motions are still coupled together by the electrons' response to polarisation electric fields caused by charge separation.

2.3 Ionospheric Vertical Structure

The bottomside ionosphere is the region below the main peak in electron density (generally $\sim 250 - 350$ km in altitude), and is typically described as consisting of up to four distinct layers as illustrated in Figure 2.4, depending on the conditions. For example, the D and F1 layers are typically absent during the night as shown. The layers form as a result of the different balances between ionisation, recombination and vertical transport




Figure 2.4. The typical altitude profile of the ionosphere at day and night, and under solar maximum and minimum conditions. Figure reproduced from Hargreaves (1992), after W. Swider, Wallchart *Aerospace Environment*, US Air Force Geophysics Laboratory.

at different altitudes, resulting from the variation in neutral composition (and density). The variations in these processes are determined by a combination of the spectrum of incoming solar radiation, the wavelength dependence of absorption cross-sections and the different recombination reactions available to different ion species.

The profile of an individual layer can often be approximated by a Chapman layer, given by

$$N = N_0 \exp \left(\frac{1}{2} \left(1 - \frac{z - z_m}{H} - \sec(\chi) \exp \left(-\frac{z - z_m}{H} \right) \right) \right) \quad (2.27)$$

if recombination is an alpha-type process, or

$$N = N_0 \exp \left(1 - \frac{z - z_m}{H} - \sec(\chi) \exp \left(-\frac{z - z_m}{H} \right) \right) \quad (2.28)$$

if recombination is a beta-type process, where N_0 is the peak plasma density, and z_m

is the altitude of peak plasma density (Chapman, 1931).

2.3.1 D-Region

The D-region does not include a maximum of electron density, as shown in Figure 2.4. The distinction between it and the E-region above is instead made because the processes that explain the E-region cannot explain the extension of the ionosphere below around 95 km (Hargreaves, 1992). The D-region requires minor neutral species (particularly NO and hydrated ions) and multiple sources of ionisation such as galactic cosmic rays to be included in order to explain the higher plasma densities than would be expected from a simpler model. Another complicating factor is that the higher neutral density allows more negative ions to form by electrons reacting directly with neutrals. Because of the high neutral densities compared to the other regions, the D-region disappears quickly and almost completely at night (Kelley, 2009).

2.3.2 E-Region

The E-region is the most ‘classical’ ionospheric layer, as it is extremely well described as an alpha Chapman layer. The dominant ions in this region are NO^+ and O_2^+ , although production of O_2^+ , N_2^+ , O^+ , He^+ and N^+ are all significant, with these forming NO^+ and O_2^+ through subsequent charge exchange reactions. The E-region is mostly formed by the longer wavelength portions of the EUV spectrum ($>\sim 80$ nm) as shorter wavelength radiation is largely absorbed in the F-region, but X-ray ionisation is also significant, particularly at solar maximum (Hargreaves, 1992).

The decay in density of the E-region at higher altitudes is largely due to the decrease in density of neutral molecular species such as N_2 and O_2 , causing a reduction in production rate. This occurs rapidly above the turbopause at around 100 km as the relative density of different species becomes uncoupled. This means that each species’

density decays with its own scale height, and so the heavier molecular gases rapidly become less significant than the atomic gases (particularly O).

One notable phenomenon which occurs in the E-region are so called ‘sporadic-E’ (Es) layers. These are thin, dense layers of plasma at E-region altitudes, mainly comprised of long-lived metal ions of meteoric origin (Whitehead, 1989). These form in the mid-latitudes due to the interaction of a shear in the neutral wind with the inclined magnetic field, which drives plasma towards the altitude of the shear node (Whitehead, 1961, 1989; Haldoupis, 2012). For the dominant molecular ions in the E-region, this process is unable to drive significant density enhancements due to the substantial loss rates, but the minority metal ions have much longer lifetimes (Whitehead, 1989). The occurrence of sporadic-E in the mid-latitudes has significant seasonal and diurnal variations, with occurrence maximising in summer afternoons (e.g. Hodos et al., 2022; Yu et al., 2019).

2.3.3 F-Region

The F-region contains the overall peak in ionospheric density, and is typically subdivided into two regions, F1 and F2. The F2 layer almost always contains the overall peak in density while the F1 layer exists below it as distinct enhancement in density above what would be expected for a single Chapman layer. One distinction between the two layers is in the nature of the loss process: in the F1 layer the loss rate is alpha-type whereas in the F2 layer it is beta-type due to the decreasing neutral density (Hargreaves, 1992). Another distinction is that the F1 layer contains a maximum of production rate, whereas no such maximum exists within the F2 layer.

F1 Layer

The F1 layer, like the E-region below it, is well approximated by an alpha Chapman layer. As for the E-region, the dominant ionic species are NO^+ and O_2^+ , but production

of O^+ , He^+ and N^+ is much more significant than O_2^+ or N_2^+ due to the higher relative concentrations of atomic species in the neutral atmosphere (Hargreaves, 1992). The ionisation is driven primarily by wavelengths between roughly 20 – 90 nm.

In order for the F1 layer to be clearly present, the altitude at which the recombination process becomes beta-type must be higher than the altitude of maximum ionisation rate. This means that above the altitude of maximum ionisation rate the loss rate will not decay particularly quickly with altitude as for an alpha-type process it is independent of neutral density. Once the loss process becomes beta-type, the loss rate begins to decrease roughly exponentially with altitude along with the neutral density, meaning that the equilibrium density starts to increase with altitude towards the F2 peak.

F2 Layer

The F2 layer, despite having the highest plasma density, does not occupy the region of peak ionisation rate. Instead, the increasing plasma density with height is due to the decreasing recombination rate as a result of the reduction in neutral density, which more than compensates for the decrease in production rate. Recombination in this region is beta-type, but considering only local ion production and loss is not sufficient to explain the peak in the F2 layer, as the recombination rate will always decrease faster with altitude than the production rate.

The F2 peak is instead only explained by including vertical diffusion of plasma as an additional loss term. The peak then occurs where the (local) plasma loss rate becomes dominated by vertical diffusion rather than recombination, typically around 250 – 350 km (Hargreaves, 1992). This means that the loss rate no longer decreases faster with altitude than the production rate as it is not dependent on the neutral density, and so the equilibrium density starts to decrease.

2.4 Ionospheric Latitudinal Structure

As well as its vertical structure, the ionosphere also displays prominent large scale structure in latitude. This is primarily driven by the differences in solar radiation exposure as a function of latitude and the Earth’s magnetic field structure. For example, the conditions around the magnetic equator are characterised by high intensity solar radiation and a quasi-horizontal magnetic field constraining vertical plasma motion. The equatorial ionosphere is not considered in any more detail here as the observations used in this thesis are exclusively at higher latitudes.

Moving poleward from the equator into the mid-latitude ionosphere, electron densities start to decrease due to the higher solar zenith angles and hence lower intensities of solar radiation available to drive ionisation. The mid-latitude ionosphere is typically considered to be relatively quiescent compared to the low- and high-latitude regions, and for many applications (e.g. GNSS) disturbances in the mid-latitudes can largely be ignored (Basu et al., 2002). Indeed, the need for a relatively undisturbed ionosphere is one of the primary motivations for building radio telescopes in the mid-latitudes (e.g. Dewdney et al., 2009), although observations using LOFAR have shown that the mid-latitude ionosphere still has substantial impacts on radio wave propagation at the frequencies of interest (e.g. Dorrian et al., 2023; Fallows et al., 2020; Flisek et al., 2023).

Moving poleward, the mid-latitude ionosphere transitions to the auroral region. This is characterised by the presence of significant energetic particle precipitation from the magnetosphere, which provides an additional driver of ionisation to counteract the decreasing contribution of solar radiation. As a result, electron densities typically start to increase with latitude at the equatorward boundary of this region, especially at night where solar photo-ionisation is absent. There is a clear minimum of electron density with latitude just equatorward of the auroral oval during the night, known as the mid-latitude trough (Moffett and Quegan, 1983; Rodger et al., 1992). This arises due to

plasma spending a long time in a region with little to no ionisation while recombination continues to lower the density. The steep electron density gradients either side of the trough are favourable for the development of several plasma irregularities (e.g. Eltrass et al., 2016; Liu et al., 2021, see Section 2.5.2).

Within the auroral region, electron densities are also often highly structured due to temporal and spatial variability in precipitation from the magnetosphere, in contrast to the smoothly varying solar radiation which dominates the production rate at lower latitudes. This is especially pronounced in the E-region as the ion lifetimes at these lower altitudes are significantly shorter than in the F-region, preventing diffusion and transport from smoothing the signatures of the precipitation distribution as effectively (Pryse et al., 1996).




Figure 2.5. The typical form of the high-latitude convection pattern as viewed from a co-rotating frame (A, left) and a sun fixed frame (B, right), shown as a function of magnetic latitude and magnetic local time. Figure reproduced from Spiro et al. (1978).

The mid-latitude trough and auroral region are also where the bulk plasma motion starts to be dominated by the high-latitude convection pattern. This is driven by electric fields which are mapped down magnetic field lines from the magnetosphere

into the ionosphere (Dungey, 1961). This mapping is due to the high field-aligned conductivity in the magnetosphere and topside ionosphere, which means that magnetic field lines are effectively electric equipotentials (Mozer, 1970). The typical pattern consists of two large convection cells centred around dawn and dusk, with antisunward flow across the pole and return flows at lower latitudes in both the dawn and dusk sectors, as shown in Figure 2.5. The stagnation point which leads to the formation of the mid-latitude trough is clearly visible at 18 MLT in the right panel, leading plasma towards the equatorward boundary of the convection pattern to spend extended periods in darkness.

Further poleward still, the auroral region gives way to the polar cap. This region is defined by existing on open magnetic field lines, field lines which connect out into the solar wind rather than to the other hemisphere (Nishimura et al., 2021). There is typically less precipitation here than in the auroral region, meaning that local ionisation rates in the polar cap are minimal. The plasma density is largely sustained by the high-latitude convection pattern, which draws plasma from mid-latitudes on the dayside over the polar cap. This plasma can then be structured as a result of variations in the convection pattern driven by changing solar wind conditions, creating structures such as polar cap patches (e.g. Carlson Jr et al., 2004; Tsunoda, 1988).

2.5 Ionospheric Variability

Besides the large scale structure of the ionosphere described in the previous sections, there are many types of localised transient structures that can cause significant deviations from the expected background state. These include externally driven processes both from below in the neutral atmosphere and from above in the magnetosphere and solar wind, and purely internal ionospheric processes driven by plasma instabilities.

Understanding of these structures is crucial as they have significant impacts on propagating radio signals used for navigation and communications (e.g. Banville and Langley, 2013; Jodalen et al., 2001; Priyadarshi, 2015). Due to their relatively small scale sizes, they cannot be directly simulated by global models of the ionosphere as they fall below the grid resolution (typically of order $2.5 - 5^\circ$ in latitude and longitude, e.g. Ridley et al., 2006; Qian et al., 2014).

2.5.1 Travelling Ionospheric Disturbances (TIDs)

One commonly observed type of ionospheric irregularity are TIDs, which generically refers to any wave-like ionospheric disturbance which propagates horizontally through the ionosphere. TIDs are generally accepted to be primarily generated by AGWs in the neutral atmosphere (Hocke and Schlegel, 1996; Hunsucker, 1982, see Section 2.1.1). Some other TID formation mechanisms such as electrodynamic instabilities in the night-time ionosphere (e.g. Kelley et al., 2003; Otsuka et al., 2013; Tsunoda and Cosgrove, 2001) have also been proposed to explain observations which do not conform to the predictions of the AGW based theory.

TID Theory

The theory of how an AGW generates the ionospheric perturbation referred to as a TID was developed by Hooke (1968), working from first principles to consider the effect of the gravity wave perturbation on all three terms of the continuity equation (2.3). The full inclusion of photo-chemical effects was a major advance on previous work. The analysis focused on effects in the F-region as this is where the largest perturbations to electron density occur.

In terms of transport effects, the main result of Hooke (1968) is that the AGW is ineffective in imparting its motion normal to \vec{B} to the plasma (at least in the F-region

due to the low collision frequency, see Section 2.2.3). As a result, it is a reasonable approximation that the perturbation to plasma velocity is entirely directed along the magnetic field. This perturbation in velocity contributes to perturbations in electron density by both direct compression and rarefaction of the plasma, and by moving plasma along density gradients.

The effect on ion production rates is a combination of two factors: the change in the local neutral number density n_{neut} and the change in the local ionising radiation flux due to changes in the neutral number density at higher altitudes. The change in the ionising radiation flux I is complicated and determined by the relative orientation of the AGW wave vector and the direction of incident solar radiation. If they are normal to one another, then the variations in I are maximised as the radiation experiences constant phase (Hooke, 1968). However, in a later work (Hooke, 1970c), the same author pointed out that this simple criterion ignores the effects of variations in neutral scale height with altitude and of viscous dissipation of the AGW. Both of these effects would refract the AGW, distorting the phase fronts and hence preventing perfect alignment of \vec{k} with the incident radiation at all altitudes as assumed in Hooke (1968).

Loss rates are the most difficult quantity to estimate perturbations to, as in principle each reaction rate may be dependent on the temperature of the reactants which is itself perturbed by the AGW. However, Hooke (1968) neglects to consider the temperature dependence as there was insufficient observational evidence for any particular temperature dependence. With this assumption, the passage of an AGW does not affect the rate coefficient of alpha-type processes, but perturbs beta-type processes in proportion to the change in neutral density.

These variations can then be inserted into equation (2.3) to obtain the expected variation in electron density. For the case of the F2 layer, photo-chemical effects can generally be neglected as the ion lifetimes are significantly longer than the wave period

(Hooke, 1970b). In this case, the density perturbation N' is given by

$$N' = iN_0 u_b \omega^{-1} \left[\left(\frac{1}{N_{e0}} \frac{\partial N_{e0}}{\partial z} + k_{zi} \right) \sin(\theta_B) - i k_{br} \right], \quad (2.29)$$

where N_0 is the unperturbed electron density, u_b is the projection of the neutral wind perturbation onto the magnetic field, ω is the AGW angular frequency, k_{zi} is the imaginary component of the AGW wavevector (i.e. corresponding to changes in amplitude with height), θ_B is the geomagnetic dip angle and k_{br} is the real component of the AGW wavevector projected onto the geomagnetic field (Hooke, 1968). As expected, the variation in electron density is proportional to the unperturbed electron density and to the amplitude of the AGW.

The dependence on the orientation of the AGW wave vector relative both to the geomagnetic field for dynamical effects and to the Sun for photo-chemical effects indicate that AGWs with certain orientations should be favoured to produce TIDs and this should manifest as a directional bias in the observations. As dynamical effects are dominant at higher altitudes, while photo-chemical effects are more significant lower down, the directional bias will in principle be different at different altitudes. The lower altitudes are expected to experience diurnal variations due to changes in solar irradiance whereas the bias at higher altitudes should be relatively constant (Hooke, 1970a). Another difference between the different altitudes is that there should be a phase difference such that the response due to dynamical effects lags that due to photo-chemical effects by one quarter period (Hooke, 1970c).

While this description of the ionospheric response to AGWs predicts certain biases in the direction of observed disturbances, this is not the only factor or necessarily the dominant factor that may determine this in practice. The spectrum of AGWs present will also play a controlling role, which is dependent on various properties of the neutral

atmosphere as discussed in Section 2.1.1.

TID Observations

TIDs have been observed with a wide range of instruments, including ionosondes (e.g. Negrea et al., 2016, 2018), incoherent scatter radar (e.g. Galushko et al., 1998; Oliver et al., 1997; Panasenکو et al., 2018), HF radars (e.g. Frissell et al., 2014, 2016; Oinats et al., 2016), HF Doppler sounders (e.g. Chum et al., 2010; Crowley et al., 1987; Fišer et al., 2017; Waldock and Jones, 1986), satellite measurements (e.g. Yin et al., 2019), GNSS Total Electron Content (TEC, integrated electron density along a given path) measurements (e.g. Ding et al., 2011; Otsuka et al., 2013; Sivakandan et al., 2021), radio interferometers (e.g. Jacobson et al., 1995; Kirkland and Jacobson, 1998) and radio telescopes (e.g. Dorrian et al., 2023; Fallows et al., 2020; Koval et al., 2017). Each technique offers different benefits, such as height resolved measurements, broad spatial coverage, or high sensitivity. As a result, a full picture of TID behaviour benefits from considering the full range of available observations in order to avoid any bias that may be inherent to a given technique. The following section provides a summary of TID observations, with a particular focus on measurements made in the European mid-latitude region as these are of most direct relevance to LOFAR.

TID Type	Wavelength (km)	Period (mins)	Velocity (m s^{-1})
LSTID	300-3000	30-300	300-1000
MSTID	100-300	10-30	50-300
SSTID	< 100	< 10	< 50

Table 2.1. The classifications of TIDs and their associated parameter ranges. Values given by Bianchi et al. (2013).

As TIDs cover a wide range of temporal and spatial scales, they are typically subdivided into different types depending on their parameters (Hunsucker, 1982). These categories are: large-scale (LS), medium- or meso-scale (MS) and small-scale (SS), and

the typical parameters are given in Table 2.1. LS and MSTIDs are the most studied due to the relative ease of observing them, while SSTIDs have only been addressed by a small number of authors (e.g. Baskaradas et al., 2014; Lan et al., 2018). LSTIDs are most commonly associated with geomagnetic storms (e.g. Ding et al., 2014; Habarulema et al., 2013) or large impulsive events such as volcanic eruptions (e.g. Themens et al., 2022) whereas MSTIDs are observed under all conditions (e.g. Frissell et al., 2016). Given that LSTIDs are therefore observed less frequently, and that MSTIDs in particular have been shown to have strong effects on observations made using radio telescopes similar to LOFAR (e.g. Dorrian et al., 2023; Helmboldt et al., 2012a; Loi et al., 2015a), the remainder of this section will focus primarily on observations of MSTIDs and SSTIDs.

The vertical structuring of TIDs is broadly consistent across many observations. The amplitude of variations in electron density typically maximises close to the F2 peak (e.g. Fedorenko et al., 2011), as predicted by equation (2.29). However, for short period variations this is not necessarily the case (e.g. Lan et al., 2018), as the Brunt-Väisälä frequency decreases with altitude and filters out short period AGWs before they reach the F2 peak. Another characteristic property is that TID phase fronts are typically significantly inclined, with the front at higher altitudes ahead of the front at lower altitudes. This manifests in many observations as a downward propagation of the TID structure (e.g. Friedman, 1966).

A critical parameter for understanding TIDs is their propagation direction, as this can provide information about potential TID sources and also characteristics of AGW and TID propagation from sources to the observer. TID propagation directions have therefore been widely studied (e.g. Chum et al., 2010; Frissell et al., 2014; Jacobson et al., 1995; Oinats et al., 2016; Perwitasari et al., 2022). The main trend in observed TID propagation directions is a correlation with the direction of the neutral wind in

the thermosphere. Specifically, the observed TIDs preferentially propagate in the opposite direction to the wind, consistent with the filtering of AGWs by the wind (e.g. Crowley et al., 1987; Fišer et al., 2017; Oinats et al., 2016; Waldock and Jones, 1986; Zalizovski et al., 2021). However, some authors report that the observed TID directions lag the wind direction by up to 90° (e.g. Crowley and Rodrigues, 2012; Waldock and Jones, 1986). Some offset between wind directions and observed wave directions may be expected given that the wind filtering is not necessarily most significant at the same altitude as the TIDs are detected, and it could take some hours for an AGW to propagate between these altitudes (Cowling et al., 1971; Vadas, 2007). However, it is unclear why different authors report different results, including some using the same measurement techniques.

One possible explanation for the variation in reported lag between wind and TID directions is based on the fact that these comparisons typically use climatological neutral winds in the absence of direct measurements and sometimes have relatively few TID observations included (e.g. Crowley and Rodrigues, 2012). In this case, either a systematic error in the model climatology or unusually large deviations from climatological behaviour at the times of observations could account for these observations. One key prediction of the wind filtering hypothesis for explaining observed TID propagation directions is that AGWs with higher phase velocities should be less affected by the winds and therefore TIDs with higher phase velocities should correlate less strongly with the wind direction. This was tested by Jacobson et al. (1995), who found that the apparent filtering was equally strong regardless of TID phase velocity, which they argued necessitated an improved explanation for the observed TID behaviour. The impact of wind filtering on TID propagation directions observed by LOFAR including this question of the impact of TID velocity is addressed in Chapter 7 of this thesis.

Wind filtering does not explain all features of observed TID propagation directions

however. For example, during the night in the mid-latitudes many authors report a clear preference for southwestward propagation (in the northern hemisphere) (e.g. Garcia et al., 2000; Otsuka et al., 2013). This directional preference cannot be explained by AGW driven TIDs, but authors have pointed to an electrodynamic process as a likely driver (e.g. Shiokawa et al., 2003). Specifically, the directional preference matches that of the Perkins instability (Perkins, 1973, see Section 2.5.2). However, the growth rate of this instability is insufficient to explain the observed MSTIDs, leading to the hypothesis of a coupling between the Perkins instability in the F-region and an instability of a sporadic-E layer beneath (e.g. Kelley et al., 2003; Tsunoda and Cosgrove, 2001). Other observations have shown that in certain regions a strong AGW source such as orography or tropospheric convection may dominate the observed TID directions even for times when the wind direction does not favour their propagation (e.g. Frisell et al., 2014).

One limitation of the results discussed above is that the methods are either limited to measuring a vertical profile at a single location (e.g. Negrea et al., 2016) or are only able to detect horizontal variations (e.g. Jacobson et al., 1995; Otsuka et al., 2013; Shiokawa et al., 2003). However, some studies using incoherent scatter radar (e.g. Medvedev et al., 2015, 2017; Nicolls and Heinselman, 2007) and multi-frequency HF Doppler sounding (e.g. Chum et al., 2021) have succeeded in resolving the full three-dimensional wavevector of the TID. By comparing these observations to theoretical TID/AGW dispersion relations such as those developed by Hines (1960) and Vadas and Fritts (2005), this enables direct verification of the causative relationship between TIDs and AGWs for individual cases. Results from Irkutsk Incoherent Scatter Radar indicate that the majority of TIDs are consistent with AGW driving, and that AGW related TIDs are particularly dominant during the daytime (Medvedev et al., 2017). These studies typically use climatological models for the atmospheric density and temperature profiles but allow the neutral wind to be a free parameter when fitting the AGW dispersion

relation. The inferred neutral winds are consistent with climatologies when taken as monthly medians (e.g. Medvedev et al., 2015, 2017) but vary significantly in individual cases (e.g. Nicolls and Heinselman, 2007).

The comparison of other MSTID parameters between authors is complicated by the range of measurement techniques and approaches to analysis that are used. All techniques, whether explicitly stated or not, have some bias in the range of wave parameters that they are most sensitive to, meaning differences between authors may simply reflect these biases rather than any underlying differences in the observed TID populations. These biases can arise directly from the measurement technique used, such as due to limited spatial ($\sim 50 - 100$ km e.g. Frisell et al., 2014; Otsuka et al., 2013) or temporal (5 min, e.g. Galushko et al., 1998) resolution, or from the analysis applied, such as requiring a sufficient amplitude of disturbance (e.g. Chum et al., 2010; Otsuka et al., 2013) or a single dominant wave signature (e.g. Jacobson et al., 1995). Given all these caveats, the most confident statement that can reasonably be made on the statistics of observed MSTID periods, wavelengths and velocities is that they are consistent with the broad range given in Table 2.1.

Despite the limitations imposed by the range of instrumentation and analysis techniques used to study MSTIDs, there is a general consistency on the relative occurrence rate of MSTIDs under different conditions. Most authors have reported no relationship between MSTID occurrence rates in the mid-latitudes and geomagnetic activity (e.g. Waldock and Jones, 1986; Frisell et al., 2016) although some studies focusing on data taken closer to the auroral oval have found a small increase in occurrence rate with auroral activity (e.g. Frisell et al., 2014; Oinats et al., 2016). Some authors have also reported an increase in daytime occurrence rate with solar activity (e.g. Oinats et al., 2016) although others have found no such relationship (e.g. Otsuka et al., 2021). One possible explanation for these differences was provided by Katamzi et al. (2012) who

reported that TID amplitudes increased with increasing solar activity in line with the increase in background density, as expected given equation (2.29). This suggests that the increase in occurrence rate with solar activity reported by Oinats et al. (2016) may in fact reflect an increase in amplitude allowing more TIDs to be detected.

While any relationships of mid-latitude MSTIDs with geomagnetic or solar activity are tentative at best, there are clear and consistent dependencies of MSTID occurrence rate on local time and season reported by a wide range of authors. The main dependency is seasonal, with a maximum of occurrence rate reported in winter (e.g. Ding et al., 2011; Fišer et al., 2017; Katamzi et al., 2012; Otsuka et al., 2013; Sivakandan et al., 2021; Waldock and Jones, 1986). Some authors also report a secondary peak of occurrence in summer (e.g. Fišer et al., 2017; Otsuka et al., 2013). A possible explanation for the winter peak in occurrence was suggested by Sivakandan et al. (2021), who highlighted that the thermospheric winds favour poleward AGW propagation in summer and equatorward in winter. Given that equatorward AGWs will have wind perturbations closer to parallel to the magnetic field, they are more able to transport the ionospheric plasma and hence create more significant TIDs which are more easily detected. The sub-seasonal variation of MSTID occurrence rate in winter was also studied by Frissell et al. (2016) over three years, who showed that variations in TID occurrence correlated with variations in the strength of the stratospheric polar vortex with a time delay of 1-2 weeks.

In terms of diurnal variations, many authors report a single daytime peak of MSTID occurrence around midday-early afternoon (e.g. Ding et al., 2011; Otsuka et al., 2013). However, when the diurnal behaviour is broken down by season, a more complicated picture emerges, with a single activity peak in winter and two in summer, one in the morning and one in the later afternoon (e.g. Fišer et al., 2017; Jacobson et al., 1995). Other authors have reported distinct populations of TIDs associated with the dawn and

dusk terminators in various seasons (e.g. Galushko et al., 1998; Panasenko et al., 2018). For nighttime TIDs, most authors report a single fairly broad maximum of occurrence rate with local time, with some seasonal variation in the timing of the maximum (e.g. Otsuka et al., 2013).

One complication that must be accounted for when comparing the seasonal, diurnal, geomagnetic and solar activity dependences of TIDs is that some authors describe a TID occurrence rate based on a binary classification of observations as either containing a TID or not (e.g. Otsuka et al., 2013) whereas others define an activity level (e.g. Frissell et al., 2016; Negrea et al., 2016) which can also incorporate other TID parameters such as amplitude and wavelength. However, even the binary classification is not independent of variations in TID amplitude, as an increase in amplitude will in general lead to more TIDs exceeding the threshold for detection and therefore an increase in the measured occurrence rate. As a result, while general comparisons of seasonal and diurnal variations are reasonable, care must be taken when addressing more subtle features of the data or attempting detailed comparisons between the results of different authors. The occurrence rates of TIDs observed using LOFAR are discussed in Chapter 7 of this thesis.

In comparison to MS and LSTIDs, little is known about SSTIDs. This is largely due to the fact that many widely used methods of studying TIDs, such as GNSS TEC maps and ionosonde measurements, do not typically have the resolution to reliably identify disturbances with such short spatial and temporal scales. Their existence at frequencies close to or below 10 min also presents a theoretical difficulty, as this is roughly the Brunt-Väisälä frequency in the thermosphere, suggesting that AGWs with shorter periods than this should not typically be able to exist at F-region altitudes.

Some dedicated studies of small time- and spatial-scale irregularities have been carried out; for example Baskaradas et al. (2014) operated an ionosonde at high cadence

with a fixed frequency and identified many irregularities with time scales (~ 10 s). However, these were not observed to repeat periodically, and may have been turbulent structures rather than TIDs. Ivanova et al. (2011) reported SSTIDs using oblique incidence sounding over Russia, and report that they often occurred simultaneously with LSTIDs, but do not specify any parameters of these TIDs. Lan et al. (2018) observed oscillations with periods < 10 min using a single ionosonde operated at high cadence (~ 1 min) and showed that these were restricted to altitudes below ~ 200 km. However, as these were only observed from a single location they were unable to determine wavelengths or phase speeds for these structures. An example of an SSTID observed using LOFAR is presented in Chapter 5 of this thesis.

2.5.2 Ionospheric Instabilities and Irregularities

Another important class of ionospheric variability is simply known as ionospheric irregularities. These are quasi-random fluctuations in ionospheric parameters driven by plasma instabilities within the ionosphere. A wide range of instability processes can occur in the ionosphere, with different processes being more significant in different regions. In this section, only those instabilities which are relevant to the mid-latitude and auroral regions will be discussed, as these are the regions of the ionosphere which are observable using LOFAR.

Ionospheric Instability Processes

Instability processes in general are mechanisms by which small perturbations to a background state are amplified. This requires a positive feedback loop drawing on a source of free energy such as bulk fluid flow, gradients in density and temperature, or electromagnetic fields. They can be characterised by a growth rate (typically denoted as γ) which depends on the background conditions and the wavevector of the initial perturbation.

For the instability to be active requires $\gamma > 0$, so the mathematical expression of the growth rate determines both the required background conditions for a given instability mechanism and the range of orientations and length scales of irregularity that can be generated.

The derivation of growth rates typically relies on the assumption of small amplitudes such that the perturbation does not appreciably impact the pre-existing background conditions. As the irregularity grows in amplitude, this simplification will break down and the linear growth described by γ will no longer be accurate. At this point, non-linear effects become significant, such as the generation of smaller scale irregularities. This process of irregularities generating smaller irregularities is referred to as a turbulent cascade, and once an instability has been active for a sufficient amount of time this generates a spectrum of irregularities. This is characterised by an outer scale determined by the instability mechanism and seeding process for the initial irregularities, an inner scale determined by the dissipation mechanism(s) such as viscosity which convert the energy of the smallest irregularities back to thermal energy, and a spectral index describing how the spectrum varies between these scales.

One commonly observed instability mechanism in the ionosphere is the gradient-drift instability (GDI: Linson and Workman, 1970). This occurs when the convective plasma motion (i.e. $\vec{E} \times \vec{B}$ drift) is aligned with the plasma density gradient perpendicular to the magnetic field. If a wavelike perturbation is seeded with a wavevector \vec{k} perpendicular to the density gradient, the Pedersen current (i.e. current parallel to \vec{E} in the plane perpendicular to \vec{B}) causes the ions to move relative to the electrons. The resulting displacement of the ion and electron density contours creates polarisation electric fields as illustrated in Figure 2.6. These polarisation fields amplify the initial perturbation by enhancing or suppressing the background $\vec{E} \times \vec{B}$ convection velocity in regions of reduced or enhanced density respectively. The growth rate of the instability in the

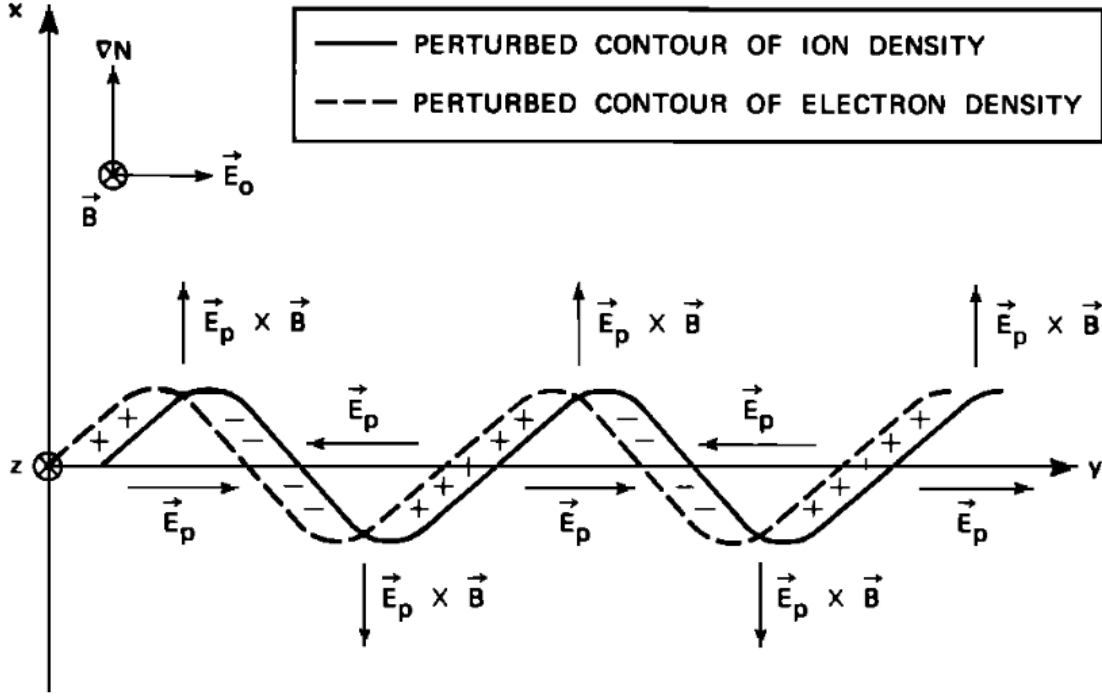


Figure 2.6. A schematic illustrating the gradient-drift instability. ∇N denotes the density gradient, \vec{E}_0 denotes the background electric field, \vec{B} denotes the magnetic field and \vec{E}_p denotes the polarisation fields. Reproduced from Tsunoda (1988).

short wavelength limit (i.e. $|\vec{k}|L \gg 1$ where L is the density gradient scale length) and neglecting electron collisions can be expressed as

$$\gamma = \begin{cases} \frac{k_y}{|\vec{k}|} \frac{\vec{k} \cdot \vec{E}}{|\vec{k}| |\vec{B}| L}, & \text{if } \omega \ll \nu_{in} \\ \frac{k_y}{|\vec{k}|} \left(\frac{\vec{k} \cdot \vec{E} \nu_{in}}{|\vec{k}| |\vec{B}| L} \right)^{\frac{1}{2}}, & \text{if } \omega \gg \nu_{in}, \end{cases} \quad (2.30)$$

where k_y is the y component of \vec{k} (co-ordinates are as in Figure 2.6), ν_{in} is the ion-neutral collision frequency and ω is the wave angular frequency (Huba et al., 1983; Keskinen, 1984). In both of the cases considered in equation (2.30), the condition for instability can therefore be expressed as $k_y \vec{k} \cdot \vec{E} > 0$.

Another ionospheric instability is the Temperature-Gradient Instability (TGI: Hudson and Kelley, 1976), which occurs when gradients of electron density and electron




Figure 2.7. A schematic illustrating the temperature-gradient instability. B denotes the magnetic field, κ_T denotes the temperature gradient scale length, κ_n denotes the density gradient scale length and E denotes the polarisation electric field. Reproduced from Eltrass and Scales (2014).

temperature in the plane perpendicular to the magnetic field are opposed. This instability arises from a resonance between a wave-like perturbation and the diamagnetic drift of the electrons (Kadomtsev, 1965). The diamagnetic drift is a drift velocity that arises from a pressure gradient within a magnetised plasma, and is mathematically analogous to the guiding centre drifts described by equation (2.24) although in this case the individual particles experience no guiding centre drift. Assuming a density gradient ∇N_e and an electron temperature gradient ∇T_e aligned with the positive x -direction and a magnetic field B in the z -direction, the diamagnetic drift of the electrons in the positive y -direction v_e can be expressed as

$$v_e = \frac{k_B T_e \nabla N_e}{e B N_e} + \frac{k_B \nabla T_e}{e B}, \quad (2.31)$$

where k_B is Boltzmann's constant and T_e is the electron temperature. If the electron temperature gradient is opposed to the density gradient (i.e. $\nabla N_e \nabla T_e < 0$), the drift velocity will increase in the lower density regions relative to the higher density regions

(as compared to the case with no temperature gradient). In the event of a perturbation to the boundary between hot and cold plasma this will lead to charge build-up as illustrated in Figure 2.7. The polarisation electric fields this creates will then lead to growth of the perturbation. The growth rate of this instability for a general orientation of the wavevector \vec{k} can be expressed as

$$\gamma = \begin{cases} \sqrt{\pi}\omega_n \left[\frac{U}{a_e} - \frac{\omega_T}{2k_z a_e} \right], & \text{if } k_z \lambda_e \gg 1 \\ \frac{1.03}{k_z \lambda_e} \omega_n \left[\frac{U}{a_e} - \frac{3\omega_T}{2k_z a_e} \right], & \text{if } k_z \lambda_e \ll 1, \end{cases} \quad (2.32)$$

where $\omega_n = k_\perp \frac{k_B T_e}{eB} \nabla(\ln N_e)$ is the frequency of the wave, U is the relative velocity of ions and electrons along the magnetic field (i.e. a measure of field-aligned currents), a_e is the electron thermal velocity, $\omega_T = k_\perp \frac{k_B T_e}{eB} \nabla(\ln T_e)$ is the temperature gradient drift frequency and λ_e is the electron mean free path (Hudson and Kelley, 1976; Kadomtsev, 1965). The pure TGI corresponds to the case with $U = 0$ and hence instability requires $\omega_n \omega_T < 0$, corresponding to opposed density and temperature gradients. The two expressions in equation (2.32) are derived under kinetic and fluid approaches respectively.

In the case of a significant velocity shear in the plasma, it can become unstable to the Kelvin-Helmholtz instability (KHI: Miura and Pritchett, 1982). This instability is driven by the free energy associated with the non-uniform plasma flow, and is characterised by the development of vortex structures (e.g. Keskinen et al., 1988). The behaviour of the KHI is strongly influenced by the relative orientation of the velocity and the magnetic field. As a result, it is common to distinguish between two different configurations, one in which the velocity is parallel to the magnetic field (parallel flow) and another in which the velocity and magnetic field are perpendicular (transverse flow) (e.g. Miura and Pritchett, 1982). As significant velocity shears in the ionosphere are most commonly associated with changes in the $\vec{E} \times \vec{B}$ drift (e.g. Keskinen and Ossakow,

F REGION $E \times B$ GEOMETRY WITH FIELD
ALIGNED CURRENT
(CURRENT CONVECTIVE INSTABILITY)




Figure 2.8. A schematic illustrating the current convective instability. ∇n_0 denotes the density gradient, \underline{E}_0 denotes the applied electric field, \underline{B}_0 denotes the magnetic field, $\underline{j}_{0\parallel}$ denotes the field-aligned current, \underline{E}'' denotes the polarisation fields induced by the field-aligned current and \underline{E}' denotes the total polarisation field arising from both $\underline{j}_{0\parallel}$ and $\underline{E}_{0\perp}$ (in this configuration $\underline{E}_{0\perp}$ has a stabilising influence but is assumed to be insufficient to fully counterbalance the destabilising influence of $\underline{j}_{0\parallel}$). Adapted from Ossakow and Chaturvedi (1979).

45

field-aligned currents can cause a configuration which is stable to the pure GDI to become unstable. This occurs when the relative drift of the electrons and ions due to the field aligned current is sufficient to exceed the relative drift caused by the ion Pedersen current which drives the pure GDI, as illustrated in Figure 2.8. If the directions of \underline{E}_0 and $\underline{j}_{0\parallel}$ were reversed, the same effect could cause the stabilisation of a situation which would be unstable to the pure GDI. In the long wavelength limit and assuming the plasma density gradient is in the positive y -direction, the magnetic field is in the positive z -direction, and the electric field in the positive x -direction, the growth rate of this instability can be expressed as

$$\gamma = \frac{-\frac{1}{N_e} \frac{\partial N_e}{\partial y} \left[\frac{E}{B} \frac{\nu_i}{\Omega_i} + v_d \frac{k_z}{k_x} \right]}{\left[\frac{\nu_i}{\Omega_i} + \frac{\nu_e}{\Omega_e} \right] \left(\frac{k_z}{k_x} \right)^2 + \frac{\nu_i}{\Omega_i}}, \quad (2.33)$$

where ν is the collision frequency, Ω is the gyrofrequency, the subscripts e and i refer to electrons and ions respectively, and v_d is the field-aligned drift velocity of ions relative to electrons due to the field-aligned current (Ossakow and Chaturvedi, 1979). The condition for instability can then be expressed as $\frac{E}{B} \frac{\nu_i}{\Omega_i} + v_d \frac{k_z}{k_x} < 0$ (Ossakow and Chaturvedi, 1979).

In the presence of currents flowing perpendicular to the magnetic field, the Farley-Buneman instability (also known as the two-stream instability) can develop (Buneman, 1963; Farley, 1963; Schmidt and Gary, 1973). This requires the relative velocity between ions and electrons to exceed the ion thermal speed (assuming a perturbation exactly perpendicular to the magnetic field: Fejer and Kelley, 1980; Schmidt and Gary, 1973). In this case a resonance exists between ion acoustic waves and the electrons with a

growth rate given by

$$\begin{aligned}\omega &= \frac{1}{1+\psi} \vec{k} \cdot \vec{v}_d, \\ \gamma &= \frac{1}{1+\psi} \left[\frac{\psi}{\nu_i} (\omega^2 - k^2 C_s^2) - k_z^2 D_{ez}^2 + \frac{\omega \nu_i}{k L_N \Omega_i} \right],\end{aligned}\tag{2.34}$$

where ω is the wave angular frequency, $\psi = \frac{\nu_i \nu_e}{\Omega_i \Omega_e} \ll 1$ is the anisotropy parameter derived from the ratios of collision frequencies ν and gyro frequencies Ω for the ions and electrons, \vec{v}_d is the drift velocity of electrons relative to the ions, C_s is the ion acoustic speed, k_z is the component of the wavevector along \vec{B} , D_{ez} is the field-aligned diffusion coefficient for electrons and $L_N = \frac{N_e |\vec{k} \times \vec{B}|}{k \cdot \vec{\nabla}_{N_e \times B}}$ is the length scale of the large scale density gradient perpendicular to \vec{B} projected onto \vec{k} (assuming wavelengths from a few metres to a few hundred metres within the E-region, Sahr and Fejer, 1996). Considering only the first term in the brackets (i.e. assuming no density gradients and \vec{k} perpendicular to \vec{B}), it is clear that instability requires $\omega^2 - k^2 C_s^2 > 0$, which implies that $\hat{k} \cdot \vec{v}_d > C_s$.

In the F-region of the night time mid-latitude ionosphere, where production is negligible but recombination is also slow enough for significant ionisation to survive, the inclination of the magnetic field means that gravity acts to drag the plasma downwards. This downward plasma motion is stabilised by an eastward directed electric field, which leads to $\vec{E} \times \vec{B}$ drifts which balance the gravity-driven motion (in general there will also be a contribution from meridional neutral wind) (Hamza, 1999; Zhou and Mathews, 2006). However, in the case where the electric field has a non-zero north-south component, this configuration is unstable to displacement of plasma along the field lines due to changes in the field-line integrated Pedersen conductivity (Perkins, 1973). The instability occurs if the change in current due to downward plasma motion acts to reduce the eastward current which is holding the plasma up against gravity. The

growth rate can be approximated (assuming the ion-neutral collision frequency is much less than the ion gyrofrequency) as

$$\gamma = \frac{E \cos(\theta_B)}{BH_n} \sin(\alpha) \sin(\theta - \alpha), \quad (2.35)$$

where θ_B is the dip angle between the magnetic field and the horizontal, H_n is the neutral atmospheric scale height, α is the angle between the wavevector k and magnetic east, and θ is the angle between the background electric field \vec{E} and magnetic east (Perkins, 1973). The $\sin(\alpha) \sin(\theta - \alpha)$ term gives rise to a directional selection illustrated in Figure 3 of Perkins (1973), where only wavevectors oriented between the electric field \vec{E} and magnetic east are unstable, and the growth rate maximises for $\alpha = \frac{\theta}{2}$.

One common feature of many of the instability mechanisms discussed here is that they involve charge separation to set up polarisation electric fields which drive the instability. Because of the extremely high field-aligned conductivity within the ionosphere, the field lines act as equipotentials and so these polarisation fields must be mapped along the field lines to higher and lower altitudes (Mozer, 1970). This is reflected in the preference for irregularity wavevectors perpendicular to the magnetic field and consequent observed tendency for ionospheric irregularities to be aligned with the magnetic field (e.g. Afraimovich et al., 2011; Berngardt et al., 2015; Herman, 1966). A further consequence is that an instability which may have a positive growth rate at one point along a field line may be suppressed by a sufficiently high Pedersen conductivity at another point along the same field line which shorts out the polarisation electric fields (Vickrey and Kelley, 1982). This is most commonly observed for instability mechanisms in the F-region in the case of significant E-region ionisation. The higher neutral densities at E-region altitudes lead to higher neutral collision frequencies and hence higher Pedersen conductivities than for the same plasma density in the F-region.

Irregularities in the Mid-Latitude Ionosphere

The mid-latitude ionosphere is typically considered to be the most quiescent region of the ionosphere (Basu et al., 2002). Certainly, the mid-latitude ionosphere contains the fewest sources of free energy for instabilities, lacking the intense current systems of the equatorial and auroral regions and the large convective flows of the high-latitudes. However, this does not mean that irregularities in this region are never observed (e.g. de Larquier et al., 2014; Fallows et al., 2020; Keskinen et al., 2004). The Perkins instability in the nighttime ionosphere has been proposed to explain the directional preference of MSTIDs (e.g. Otsuka et al., 2013), mid-latitude spread-F (e.g. Huang et al., 1994; Miller et al., 1997) and turbulent upwelling in coherent scatter radar measurements (e.g. Fukao et al., 1991; Kelley and Fukao, 1991), typically in combination with initial perturbations introduced by AGWs. The density gradients produced by TIDs (e.g. Dorrian et al., 2023; Fallows et al., 2020) and those present in the mid-latitude trough (e.g. Eltrass et al., 2016) are also associated with irregularity generation.

The mid-latitude trough in particular deserves special attention as regards mid-latitude irregularities, being a regular feature of the ionosphere that can produce significant irregularities in specific consistent regions. The density gradients on both the poleward and equatorward boundaries can be unstable to the GDI depending on the state of the high-latitude convection pattern. These density gradients are also consistently unstable to the TGI, as the minimum of plasma density in the trough is co-located with a maximum of electron temperature (Rodger et al., 1992). Observations using coherent-scatter radar (e.g. de Larquier et al., 2014; Greenwald et al., 2006), GNSS scintillation receivers (e.g. Eltrass et al., 2016) and in-situ satellite measurements (e.g. Liu et al., 2021) have confirmed the presence of irregularities in the trough. These have been shown to be concentrated in the regions of steepest temperature and density gradient (Liu et al., 2021) and modelling work has shown they are consistent with a

combination of the GDI and TGI (Eltrass et al., 2016). In the poleward wall of the trough, which also marks the transition from the mid-latitude ionosphere to the auroral oval, the presence of particle precipitation and associated field-aligned currents permits the development of the current convective instability (e.g. Siddle et al., 2004; Xiong et al., 2020).

Irregularities in the mid-latitude F-region are most commonly observed at night (e.g. Greenwald et al., 2006; Kane et al., 2012; Nishitani et al., 2019). This can be primarily explained by two factors, the first being that the mid-latitude trough, where many of the irregularities are generated, is only present during the night (Rodger et al., 1992). Another factor is the presence of significant E-region plasma densities during the day. As mentioned previously, the high Pedersen conductivity of the E-region can short out any polarisation electric fields associated with perturbations in either the E- or F-regions, suppressing the instabilities. This mechanism has been shown observationally to explain the observed variation of irregularity occurrence rates in the mid-latitudes by Kane et al. (2012). Another factor is that during the day, ionisation rates in the F-region can be substantial relative to the growth rates of instabilities in the mid-latitudes, particularly the Perkins instability (Perkins, 1973). This could lead to solar produced ionisation smoothing out perturbations in the plasma density faster than the instability mechanism is able to amplify it and thereby suppressing the growth of irregularities.

In addition to F-region irregularities, irregularities in the mid-latitude ionosphere are frequently observed associated with sporadic-E layers. These are most commonly observed with coherent backscatter radars (e.g. Ecklund et al., 1981; Yamamoto et al., 1991) which can detect the irregularities but are not able to confirm the presence of sporadic-E. However, simultaneous measurements using other techniques such as ionosondes and sounding rockets have been able to confirm the presence of sporadic-E layers when these irregularities are observed (e.g. Larsen et al., 2005; Maruyama et al.,

2006). Various authors have proposed instability mechanisms which may explain the occurrence of these irregularities, such as the GDI (e.g. Ecklund et al., 1981; Woodman et al., 1991) and the KHI in the neutral atmosphere (i.e. turbulence in the neutral atmosphere generated as a result of the wind shear which generates the sporadic-E layer, e.g. Larsen, 2000). Another instability mechanism specifically in night-time sporadic-E layers was proposed by Cosgrove and Tsunoda (2002), who showed that a sporadic-E layer is unstable to plane wave perturbations with phase fronts aligned north-west to south-east (in the northern hemisphere) due to the ions being more collisional than the electrons. Given the fact that this alignment is the same as that required for the Perkins instability (see above) and that electric fields can map efficiently between the E- and F-regions, it has been shown that these two instabilities can couple together, and that the growth rate of this coupled instability is higher than that of either instability process acting alone (Cosgrove and Tsunoda, 2004).

Irregularities in the Auroral Ionosphere

The auroral ionosphere, in contrast to the mid-latitudes, has long been known to contain intense irregularities (e.g. Aarons, 1973; Basu et al., 1983; Oksman and Tauriainen, 1971; Tsunoda, 1988). On a climatological basis, the irregularities in the auroral region are most common along the poleward and equatorward boundaries of the diffuse aurora (Tsunoda, 1988, and references therein) which correspond to steep latitudinal gradients in electron density. However, due to the dynamic nature of the auroral ionosphere, at a given moment in time irregularities may be produced anywhere within the auroral oval, such as those associated with discrete auroral forms (e.g. Ivarsen et al., 2024; Sahr and Fejer, 1996) or with large scale F-region structure carried by the high-latitude convection pattern such as auroral and boundary blobs (e.g. Jin et al., 2016; Weber et al., 1985).

Irregularities in the auroral ionosphere are associated with a range of instabilities which can act individually or in combination with one another. Naturally, given the presence of particle precipitation and hence field-aligned currents, the current convective instability is frequently active in this region (e.g. Jin et al., 2021; Ossakow and Chaturvedi, 1979). The GDI is also commonly active due to the large convection velocities associated with the high-latitude convection pattern (e.g. Basu et al., 1990), particularly on the trailing edge of auroral and boundary blobs where the necessary alignment of density gradients and $\vec{E} \times \vec{B}$ drifts occurs (e.g. Tsunoda, 1988; Weber et al., 1985).

As discussed by Keskinen et al. (1988), the irregularities generated by the GDI can then themselves become unstable to the KHI, creating additional irregularities. In general, the $\vec{E} \times \vec{B}$ drifts which drive the GDI may also themselves contain a shear component which could drive the KHI, meaning that both instability processes may be active simultaneously. Modelling work by Rathod et al. (2021) showed that the KHI becomes more dominant as the plasma becomes less collisional (i.e. generally becomes more dominant with increasing altitude). They highlighted that this may impose a limitation on comparisons of data from in-situ satellite density measurements and ground based scintillation measurements of irregularities associated with the same large scale structure, as the irregularities at different altitudes may be driven by fundamentally different physical mechanisms. The KHI is also expected to be preferentially active on the large-scale structure of the poleward boundary of the auroral oval compared to the equatorward boundary (Keskinen et al., 1988).

In the auroral E-region, the Farley-Buneman instability occurs due to intense horizontal electron currents associated with auroral precipitation, known as auroral electrojets (Boström, 1964). These irregularities have been shown to preferentially occur outside the regions of greatest ionisation and current density, as the increased plasma

density causes increased conductivity which can suppress irregularity growth by shorting out any polarisation electric fields which are generated (Ivarsen et al., 2024).

One counter-intuitive feature of auroral irregularities is that although many of the instability mechanisms draw their free energy from auroral precipitation, the presence of substantial auroral precipitation can actually suppress instability growth due to the enhanced E-region densities it creates. This effect has been observed by Takahashi et al. (2022) using a combination in in-situ rocket measurements and optical measurements to obtain simultaneous measurements of irregularity spectra and auroral precipitation. They confirmed that irregularity levels were lower in regions for which auroral precipitation caused significant E-region ionisation, consistent with the explanation proposed by Vickrey and Kelley (1982). The observations of Ivarsen et al. (2024) mentioned above which localised the irregularities generated by the Farley-Buneman instability to the edges of the observed auroral arcs rather than the centre provide further evidence for this phenomenon.

Chapter 3

Ionospheric Radio Propagation

3.1 Physics of Ionospheric Radio Propagation

The ionosphere has a significant impact on the propagation of radio waves, due to variations in refractive index n as a function of plasma density, magnetic field and ion-neutral collision frequency. The refractive index is dependent on the ratio of radio (angular) frequency ω to the plasma (angular) frequency ω_p . The plasma frequency is the characteristic frequency of bulk oscillations of the electrons in the plasma (there is a separate characteristic frequency for the ions but as this is lower by a factor of $\frac{m_i}{m_e} \sim 10^4$ it can be ignored here). For a cold plasma (i.e. where the thermal motion of the electrons is negligible relative to the oscillatory motion, which applies throughout the ionosphere), the plasma frequency can be expressed as

$$\omega_p^2 = \frac{Ne^2}{m_e\epsilon_0}, \quad (3.1)$$

where e is the charge of the electron, m_e is the electron mass and ϵ_0 is the permittivity of free space. In the mid-latitude ionosphere, the plasma frequency is typically $\sim 3 - 15$ MHz at the F2 peak. The refractive index at a radio (angular) frequency ω is

given by the Appleton-Hartree equation

$$n^2 = 1 - \frac{X}{1 - iZ - \left[\frac{Y_T^2}{2(1-X-iZ)} \right] \pm \sqrt{\frac{Y_T^4}{4(1-X-iZ)^2} + Y_L^2}}, \quad (3.2)$$

where $X = \frac{\omega_p^2}{\omega^2}$ represents the dispersive delay; $Y_T = \frac{\omega_T}{\omega}$ and $Y_L = \frac{\omega_L}{\omega}$ represent the effect of the magnetic field on the motion of the electrons; and $Z = \frac{\nu}{\omega}$ represents the effect of collisions between the plasma and the neutrals (Hargreaves, 1992). $\omega_T = \Omega_B \sin(\theta)$ and $\omega_L = \Omega_B \cos(\theta)$ are the transverse and longitudinal components of the electron gyrofrequency Ω_B where θ is the angle between the wave propagation direction and the magnetic field, and ν is the electron-neutral collision frequency. The interested reader is referred to Ratcliffe (1959) for a complete derivation of this equation.

Because the electrons gyrate around the magnetic field in a characteristic direction, the plasma becomes birefringent as represented by the \pm term in equation (3.2). The positive sign corresponds to the ordinary wave and the negative to the extraordinary wave, which are the two possible circular polarisations (left- or right-handed).

For many applications, it is possible to neglect the absorption and magnetic field terms to simplify the equation, giving

$$n = \sqrt{1 - X} \approx 1 - \frac{X}{2}, \quad (3.3)$$

where the approximation holds provided $X \ll 1$, that is that the radio frequency is much greater than the plasma frequency. This approximation allows n at a given radio frequency to be expressed as a linear function of plasma density N . It is important to note that as $n < 1$, the phase velocity of radio waves in the ionosphere $v_p = \frac{c}{n}$ is always

greater than c . The group velocity v_g on the other hand, can be expressed as

$$\begin{aligned} v_g = \frac{d\omega}{dk} &= \left(\frac{dk}{d\omega} \right)^{-1} = \left(\frac{d}{d\omega} \left[\frac{\omega n}{c} \right] \right)^{-1} = \left(\frac{d}{d\omega} \left[\frac{\sqrt{\omega^2 - \omega_p^2}}{c} \right] \right)^{-1} \\ &= \left(\frac{\omega}{c \sqrt{\omega^2 - \omega_p^2}} \right)^{-1} = cn, \end{aligned} \quad (3.4)$$

where it is assumed that absorption and magnetic field terms can be neglected as in equation (3.3). Given that $n < 1$ it is therefore clear that $v_g < c$, and that as ω approaches ω_p the group velocity will tend to zero.

Radio signals that pass through the ionosphere can therefore be used to infer the properties of the ionosphere along their propagation path. A widely used example of this is the use of GNSS signals to derive TEC (e.g. Coster et al., 2017; Habarulema et al., 2013; Otsuka et al., 2013). The details of how this is derived are explained in Section 3.3.2. These measurements have been used to study a wide range of ionospheric phenomena, including TIDs (e.g. Otsuka et al., 2013; van de Kamp et al., 2014) and the effects of solar eclipses (e.g. Coster et al., 2017). They are also key data source for ionospheric models which use them both for validation (e.g. Chen et al., 2020; Perlono et al., 2018) and as inputs to more accurately reflect the real ionospheric conditions (e.g. Elvidge and Angling, 2019; McNamara et al., 2013).

As well as refractive effects, small-scale plasma density structures can cause diffractive effects on radio waves. The characteristic scale of irregularities capable of this is known as the Fresnel scale D_F , given by

$$D_F = \sqrt{2\lambda L}, \quad (3.5)$$

where λ is the radio wavelength and L is the distance along the line of sight to the irregularity (Fallows et al., 2020, note this expression assumes an infinitely distant

source, appropriate for astronomical sources but not satellites). Diffractive effects on trans-ionospheric radio signals primarily arise from structures with sizes ranging from approximately the Fresnel scale to a tenth of that value (Basu et al., 1998). The Fresnel scale is defined by considering waves which are incident on a screen and then propagate a distance L to an observer. These waves will experience a phase difference due to the different path lengths from various points on the screen to the observer, and the Fresnel scale is the scale at which these phase differences reach a certain value. In the case of equation (3.5), this value is 2π radians, but other authors use definitions of the Fresnel scale corresponding to other values such as 1 radian, or which define it in terms of diameter rather than radius (e.g. Forte et al., 2022; Narayan, 1992). Therefore, when comparing results between authors it is vital to identify the particular definition used in each case.

For frequencies below the plasma frequency, the ionosphere is opaque and signals cannot pass through it (the refractive index becomes purely imaginary in equation (3.3)). Instead they are reflected from the layer at which the plasma frequency is equal to the radio frequency (assuming propagation parallel to the gradient of electron density i.e. typically \sim vertical in the mid-latitudes). This can be used to probe the bottomside ionosphere using an ionosonde, which transmits pulses at a range of frequencies and measures the delay of their echoes (e.g. Bibl and Reinisch, 1978; Verhulst et al., 2022). This reflection (up to higher frequencies than the plasma frequency due to non-normal incidence) is also exploited for long-range HF communications (e.g. Goodman et al., 1997; Wang et al., 2022) and over-the-horizon radar systems (e.g. Headrick and Skolnik, 1974; Thayaparan et al., 2018), enabling signals to be sent over the curvature of the Earth where no direct line of sight is available.

3.2 Modelling Ionospheric Radio Propagation

While the physics of radio propagation in plasmas is well understood, the complex structure and dynamics of the ionosphere mean that predicting the propagation characteristics in any realistic scenario requires a computational approach. Doing this requires compromises between computational speed, accuracy, and which effects are included within the model. It is important to understand what these compromises are for a given model to ensure that the approach used is appropriate for a given problem.

3.2.1 Raytracing

Raytracing is a very widely used and powerful tool for modelling propagation in the ionosphere (e.g. Koval et al., 2018; Siddle et al., 2004). It takes a purely geometric optics approach, meaning that it is unsuitable for situations where wave phenomena such as interference are significant, but captures refraction and reflection accurately. A key advantage of this approach is that rather than just determining the observed intensity at a given point, it is possible to track the path taken by an individual ray. This path provides information on angle of arrival and propagation time, both invaluable for many applications such as over-the-horizon radar (e.g. Pederick and Cervera, 2016).

3.2.2 Phase Screen Modelling

Phase screen modelling approaches start from a different physical model to raytracing, in that they use a full wave optics approach. This means that phase dependent effects such as interference can be fully represented. However, the trade off is that propagation through a continuous medium cannot be represented, and so the medium must be represented by one or more discrete ‘phase screens’. These impart a position dependent phase change on the incident wave, which then propagates on to the next screen or to

the observer plane.

To express this idea mathematically, the phase change $\Delta\phi$ induced at a given point on the screen is expressed as

$$\Delta\phi = k \int n dl \approx k \int \left(1 - \frac{X}{2}\right) dl = kL - \frac{k}{2} \int \frac{Ne^2}{m_e \epsilon_0 \omega^2} dl = kL - \frac{e^2}{2m_e \epsilon_0 \omega c} \text{TEC}, \quad (3.6)$$

where k is the free space wavenumber of the radio wave, $\int dl$ refers to integration along the line of sight, L is the distance propagated, c is the speed of light and $\text{TEC} = \int N dl$ is the TEC along the line of sight. The approximation in equation (3.6) is the same as the approximation made in equation (3.3), that magnetic field effects and absorption are negligible and $\omega \gg \omega_p$. The key insight of equation (3.6) is that the phase change induced by an ionospheric phase screen is proportional to the TEC perturbation divided by the radio frequency.

The numerical approach to phase screen modelling has been described by several authors (e.g. Ding et al., 2021; Hocke and Igarashi, 2003; Sokolovskiy, 2001). It exploits the fact that the effect of the phase screen is easily represented in spatial co-ordinates, while propagation of the wave is easily represented in wavevector co-ordinates. By Fourier transforming back and forth at the appropriate steps, the effects can be represented by simple multiplications. For certain specific perturbations, it can also be possible to derive an analytic solution for the observed wave after interaction with the screen (e.g. Meyer-Vernet, 1980).

Considering a wave with complex amplitudes $E(x, z_0^-)$ incident on a phase screen at z_0 (infinite in the x -direction), the amplitude of the wave after the phase screen $E(x, z_0^+)$ is given by

$$E(x, z_0^+) = E(x, z_0^-) \exp(i\Delta\phi(x)), \quad (3.7)$$

where $\Delta\phi(x)$ is the phase change imparted by the screen. This can then be Fourier

transformed in the x -direction to give the wave spectrum $\tilde{E}(k_x, z_0^+)$ where k_x is the Fourier conjugate of x (i.e. the wavevector component in the x -direction). The propagation to a screen at z_1 is then given by

$$\tilde{E}(k_x, z_1^-) = \tilde{E}(k_x, z_0^+) \exp \left(i(z_1 - z_0) \sqrt{k^2 - k_x^2} \right), \quad (3.8)$$

where k is the magnitude of the wavevector. This reflects the different amounts of phase accumulated by each angular component as it propagates a set distance in the z -direction. This spectrum can then be inverse Fourier transformed to give $E(x, z_1^-)$ and the calculation can then repeat from equation (3.7) as many times as necessary.

The phase screen approach is also used as an approximation to derive statistical quantities such as the S_4 index for a given spectrum of ionospheric irregularities (e.g. Forte et al., 2022; Mercier, 1962). This approximation is particularly useful in cases where the available measurements provide no information on where along the line of sight the irregularities exist (e.g. Helmboldt and Zaboltn, 2022). As well as determining statistical quantities for stochastic scintillation, phase screen models have been used to simulate propagation of astrophysical and man-made radio signals through a variety of specific ionospheric structures (e.g. Carrano et al., 2020; Ding et al., 2021; Hocke and Igarashi, 2003; Ludwig-Barbosa et al., 2019; Meyer-Vernet, 1980; Trigg et al., 2024; Wang et al., 2014; Wood et al., 2024).

3.3 Ionospheric Effects on Radio Signals

While the previous sections dealt with the underlying physics of radio propagation in the ionospheric plasma, and methods of modelling the signal propagation in realistic ionospheric scenarios, it is also useful to summarise some of the common effects observed on trans-ionospheric radio signals such as those observed using LOFAR. The physical

mechanism for the observations is that the ionospheric structures exist, and hence cause distortions to the observed radio signals, whereas the question for the researcher starts with the distortions to the radio signals and requires them to ‘reverse engineer’ the answer in terms of what structures were responsible. Hence, an understanding of the types of effects which may be observed and what structures cause them is fundamental to this work.

This section is not intended as an exhaustive description of all the effects that may be observed, but rather focuses on those relevant to observations made using LOFAR and to other observing techniques which are used within this thesis. For a more complete summary of ionospheric effects on radio propagation, including magnetic field effects and absorption, the interested reader is referred to Lawrence et al. (1964), Millman (1967), Flock (1987) and references therein.

3.3.1 Scintillation

One commonly observed effect of the ionosphere across all radio wavelengths is scintillation, which refers to rapid variation in received phase, amplitude, or both. This is frequently observed in GNSS data at high latitudes and in equatorial regions, but much less often in mid-latitudes (Basu et al., 2002). Scintillation in amplitude is typically quantified by the S_4 index (Briggs and Parkin, 1963), given by

$$S_4 = \frac{\sigma_I}{\langle I \rangle}, \quad (3.9)$$

where σ_I denotes the standard deviation of intensity I , and $\langle I \rangle$ denotes the mean intensity. For the purposes of the following discussion, the focus will be on amplitude scintillation rather than phase scintillation, as phase scintillation of natural radio sources cannot be observed with the LOFAR observations considered in this thesis.

Scintillation arises from irregularities in the propagation medium which impose distortions to the radio wave, and the fact that these distortions will evolve as the signal propagates on from the region containing the irregularities to the observer. Scintillation can be described by two distinct conceptual frameworks (e.g. Briggs and Parkin, 1963). One of these is the phase screen model described in Section 3.2.2, wherein the irregularities are collapsed to a thin screen which imposes perturbations purely to the phase of the signal (i.e. absorption is completely neglected). Amplitude scintillation then arises due to the evolution of this perturbed wavefront as it propagates to the observer and different regions of the wavefront interfere constructively or destructively with one another. The alternative description considers the irregularities as a population of scatterers, which scatter some portion of the radio energy from a given point in space. The observed scintillations are then the result of the interference of these scattered waves and the initial unperturbed waves.

Both of these descriptions provide equivalent mathematical results, but have distinct advantages in providing understanding of the observed phenomena. The phase screen model is mathematically more flexible, whereas the scattering model becomes cumbersome when the scattering is strong (Briggs and Parkin, 1963). However, the scattering model provides a useful physical explanation of this distinction between ‘weak’ and ‘strong’ scattering and related changes in observed scintillation behaviour. Weak scattering refers to the case for which the interactions between scattered waves are negligible, and hence the observed variations can be explained purely in terms of the interaction of scattered waves with the initial unscattered wave. Strong scattering, on the other hand, occurs when the interactions of scattered waves with each other are significant or even dominant in determining the observed scintillation.

The distinction between weak and strong scintillation in the phase screen model is described by introducing an additional parameter, the diffractive scale r_{diff} (e.g. Mevius

et al., 2016; Narayan, 1992). This is the scale over which the root mean square phase difference between two points on the screen is 1 radian, analogous to the definition of the Fresnel scale but considering the phase perturbations imposed by the screen rather than the geometric term. Weak scattering occurs for the case where $D_F \ll r_{diff}$, meaning that phase fluctuations imposed by the screen are small compared to the geometric term and so the received phase fluctuations and resultant amplitude fluctuations are small (Narayan, 1992). Strong scattering, on the other hand, refers to the case where $D_F \gg r_{diff}$. In this case, the definition of the Fresnel scale as the region over which the waves are relatively coherent breaks down due to the perturbations imposed by the screen, and r_{diff} takes over as defining the scale over which the phase is relatively coherent (Narayan, 1992). This also makes it possible for multiple distinct regions of the screen to have coherent phase and therefore for the waves from those regions to interfere with one another, analogous to the idea of strong scattering occurring when interactions between the scattered waves become significant.

The frequency dependence of $D_F \propto \frac{1}{\sqrt{f}}$ is given through equation (3.5), while it can be inferred that $r_{diff} \propto f$ from equation (3.6). Given that D_F increases with decreasing frequency while r_{diff} decreases with decreasing frequency, this suggests that strong scintillation will more commonly be observed at lower frequencies, and that for broadband or multi-frequency observations it is possible that the same ionospheric irregularities may induce weak scintillation at the higher frequencies and strong scintillation at the lower frequencies. This is obviously relevant to an instrument such as LOFAR, which has a wide observing bandwidth, and observations of ionospheric scintillation have confirmed this transition from weak to strong scintillation with decreasing frequency (e.g. Fallows et al., 2014). A characteristic of the difference between weak and strong scintillation in observations of this type is that weak scintillation is typically coherent over a relatively wide bandwidth ($\Delta f \sim f$, e.g. Narayan, 1992) whereas strong scintillation

is often uncorrelated on even closely spaced frequencies.

Interest in ionospheric scintillation has two primary motivations: firstly, it provides a way to study ionospheric irregularities which is relatively inexpensive and hence enables monitoring over large areas (e.g. Jayachandran et al., 2009), and secondly, scintillation has significant negative impacts on engineering systems which rely on trans-ionospheric radio signals, such as GNSS (e.g. Banville and Langley, 2013; Priyadarshi, 2015). GNSS signals at L band frequencies ($\sim 1.2 - 1.6$ GHz) in particular have been a focus of recent research (e.g. Afraimovich et al., 2011; Eltrass et al., 2016; Flisek et al., 2023; Jin et al., 2016; Spogli et al., 2013) for both of these reasons, as GNSS receivers are relatively cheap and the large number of GNSS satellites means that there are many signals available to study. From the perspective of understanding and mitigating impacts on operational systems this makes sense, as GNSS signals are relied upon for a wide range of applications such as navigation and monitoring the motion of tectonic plates (e.g. Gebre-Egziabher and Gleason, 2009; Jin et al., 2022a; Kaplan and Hegarty, 2017) and therefore understanding specifically the irregularities that affect GNSS signals is important. However, from an ionospheric science perspective, it is important to consider other signals, particularly at lower frequencies where the Fresnel scale is larger (approximately a factor 10 difference between GNSS and LOFAR frequencies) and the sensitivity to a given change in TEC greater, to ensure that a more complete picture of ionospheric irregularities can be obtained.

One key parameter that can be studied using scintillation is the spectrum of ionospheric irregularities responsible for the scintillation (e.g. Carrano et al., 2012; Forte, 2008; Rino et al., 2016). While in the ideal case, in-situ measurements of the irregularity spectrum are preferable, the availability of these is extremely limited as they are only realistically obtainable using sounding rockets (e.g. Takahashi et al., 2022). Scintillation measurements do not provide a direct measurement of the irregularity spectrum,

but the power spectrum of the variations in intensity and/or phase can be determined. In the case of weak scintillation, the relationship between the power spectrum of the intensity scintillation and the irregularity spectrum is relatively intuitive: the power spectrum is the irregularity spectrum high pass filtered by the Fresnel scale, and with the spectral index steepened by 1 (assuming the irregularities do not evolve as they move across the line of sight (i.e. frozen-in flow): Forte, 2008). As a result, both the Fresnel scale and spectral index of the irregularities can in principle be determined by identifying the spectral index and spectral break frequency of the scintillation power spectrum. In a simplified mathematical form, this relationship can be expressed as

$$P(\nu) \propto F(\nu|\nu_F)E(\kappa) \times \nu^{-1}, \quad (3.10)$$

where $P(\nu)$ is the observed intensity power spectrum as a function of temporal frequency ν , $F(\nu|\nu_F)$ is the filter function given the Fresnel frequency ν_F and $E(\kappa)$ is the irregularity spectrum as a function of spatial frequency κ . The spatial frequency κ can be related to the temporal frequency ν via the irregularity scan velocity v as $\kappa v = \nu$, where the scan velocity is the relative velocity of the irregularities to the line of sight. Similarly, the Fresnel frequency ν_F is related to the Fresnel scale D_F by $\nu_F = \frac{v}{D_F}$. For a more rigorous mathematical derivation of the relationship between irregularity spectra and observed power spectra under weak scattering, the interested reader is referred to Forte et al. (2022) and references therein.

In the case of strong scattering, there is unfortunately no such simple relationship between the observed scintillation power spectrum and the spectrum of irregularities responsible. It is possible to derive the expected scintillation spectra for a given spectrum of irregularities, as was demonstrated by Carrano and Rino (2016). As is clearly shown by their Figure 1, there is no longer a clear feature in the scintillation power

spectrum corresponding to the Fresnel frequency in the case of strong scintillation, and the behaviour deviates from the power law over a broad range of temporal frequencies. While Carrano and Rino (2016) do demonstrate a method for estimating the irregularity parameters which is applicable in the case of strong scintillation, this has not been widely adopted by other authors studying ionospheric scintillation.

Another quantity that can be estimated from scintillation measurements, by leveraging a network of receivers monitoring the same source, is the scan velocity v of the irregularities. This is achieved by cross-correlating the signals from different receivers to measure the time delay between the signals as a function of the baseline between the receiver pair. This is commonly used to simply determine the time delay on each baseline based on the peak of the cross-correlation, and then calculate a best fit velocity corresponding to the time delays on all baselines (e.g. Fallows et al., 2020). However, Grzesiak et al. (2022) highlight that this approach will provide unreliable estimates of the velocity under certain conditions. The estimation of velocities from scintillation data, including the improved method proposed by Grzesiak et al. (2022), is discussed in more detail in Chapter 8 and Appendix H of this thesis.

3.3.2 Refractive Delays and Determination of Total Electron Content

While scintillation arises from small-scale structure in the ionosphere, typically of a few kilometres or less, there are also significant effects of large-scale ionospheric structure on radio propagation. The most basic of these is the refractive delay of a signal caused by the non-unity refractive index of the ionosphere and which occurs even in the case of a perfectly horizontally stratified ionosphere. This effect forms the basis of some of the most common and foundational techniques for measuring the ionosphere, such as ionosondes and measurements of TEC from GNSS signals.

As mentioned previously, an ionosonde operates by transmitting pulses at a range of frequencies (typically between $\sim 1 - 20$ MHz) and measuring the time delay for an echo to be received from the ionosphere. The result is a plot of echoes as a function of frequency and virtual height, known as an ionogram. The virtual height is defined simply as the height of reflection assuming that the signal travelled at the speed of light c up until the point of reflection, which is in general significantly different to the true height of reflection. In an idealised case of a horizontally stratified ionosphere and a precisely known magnetic field (magnetic field effects in equation (3.2) are non-negligible at these frequencies) it is possible to ‘invert’ the profile of reflection frequency (plasma frequency) vs virtual height to obtain the profile as a function of true height. In this idealised case, the inversion is limited by the unknown profile at lower altitudes (where the plasma frequency is below the minimum sounding frequency) and in any valleys such as that between the E- and F-regions. In reality, extra complications are introduced by off-vertical propagation and small-scale irregularities which cause the observed echoes to be spread in range and/or frequency relative to the idealised clean profile. For more detailed and rigorous discussion of ionogram analysis, the interested reader is referred to Piggott and Rawer (1972), Reinisch et al. (2005), Wright (1990) and references therein.

The derivation of TEC from GNSS signals exploits the fact that these satellites transmit on two (or more) frequencies simultaneously (e.g. GPS uses L1 at 1575.42 MHz and L2 at 1227.60 MHz). Given the frequency dependence of the ionospheric refractive index in equation (3.3), these different frequencies will experience different delays in terms of both range and phase due to the ionosphere. The total observed phase and range delay at each frequency will include other contributions such as those from the troposphere and the satellite-receiver geometry, but it is possible to define a linear combination of the two which depends purely on the ionospheric term (Arikan et al.,

2008). The phase delay measurements are typically much more precise than the range measurements, but suffer from an inherent 2π ambiguity. As a result, the most effective means of estimating the TEC is to use the measured range delay to resolve this ambiguity in the TEC inferred from the phase delay (e.g. Arikan et al., 2008). This calculation relies on the assumption that the approximation in equation (3.3) holds, which is reasonable for GNSS frequencies given that typical plasma frequencies in the ionosphere do not exceed a few 10s of MHz. Other major complications include multipath propagation when the satellite is at low elevation, which is typically solved by simply discarding all signals for which the satellite was below $\sim 20 - 40^\circ$ in elevation (e.g. Spogli et al., 2021; Themens et al., 2022). For a more in depth explanation of the derivation of TEC from GNSS measurements the interested reader is referred to Arikan et al. (2008) and Leick et al. (2015).

3.3.3 Refractive Position Shifts

The previous section considered the refractive effects that can be observed when the ionospheric refractive index has no gradients perpendicular to the line of sight. When such perpendicular gradients are present, in addition to the refractive delay the signal will deviate from the direct line of sight. While the refractive delays are insignificant for radio astronomy (with the exception of higher order magnetic field terms which affect the signal polarisation, e.g. Porayko et al., 2023) these deviations of the signal cause shifts in the apparent position of radio sources (e.g. Stewart and McLean, 1982; Wild et al., 1959). Under the simplifying assumption of the ionosphere as a thin phase screen, the angular position shift is proportional to the gradient of TEC divided by the radio frequency squared.

While these effects are useful for studying the ionosphere (e.g. Helmboldt and Zabolotin, 2022; Jordan et al., 2017; Waszewski et al., 2022) they are extremely prob-

lematic for radio astronomy. By shifting the different frequencies by different amounts, this effect can obscure structure within a resolved source (e.g. Wild et al., 1959). Given that ionospheric conditions can vary substantially over a timescale of minutes, these effects are especially problematic for observations which require long integration times to reach the necessary sensitivity (e.g. Yatawatta et al., 2013) as variations in the refractive shift will act to smear the resulting image. However, calibration techniques used in radio astronomy are able to correct for this and other ionospheric effects sufficiently to permit useful imaging to still be carried out (e.g. de Gasperin et al., 2020a; Hothi et al., 2021). These calibration techniques provide precise measurements of ionospheric conditions which will be used in Chapters 6 and 7 of this thesis.

3.3.4 Focusing and Defocusing

To further extend the refractive effects beyond that considered above, the next step is to consider the effect of curvature (i.e. non-zero second derivative) in the TEC. The effect of this is focusing or defocusing of the source depending on the sign of the curvature, and has frequently been reported in observations of astronomical sources (e.g. Dorrian et al., 2023; Koval et al., 2017; Trigg et al., 2024; Wild and Roberts, 1956b). A useful understanding of the relevant physical quantities can be derived from considering the simple model of a one-dimensional parabolic lens, in which the phase change $\Delta\phi$ can be expressed as

$$\Delta\phi = Cx^2 = -\frac{e^2}{2m_e\epsilon_0\omega C}\nabla^2\text{TEC}x^2, \quad (3.11)$$

where C describes the curvature of the lens, $\nabla^2\text{TEC}$ is the second derivative of TEC with respect to x (the co-ordinate perpendicular to the line of sight). The expression of C in terms of $\nabla^2\text{TEC}$ and ω is based on the expression for ionospheric phase change in equation (3.6). At a distance z from the screen along $x = 0$, the condition for focusing

in the limit of geometric optics is that waves coming from all points on the screen have the same phase, meaning that the phase introduced by the lens in equation (3.11) exactly cancels the geometric phase accumulated in propagating from the lens to the observer. In the small angle approximation, this geometric phase $\Delta\phi_g$ can be expressed as

$$\Delta\phi_g = \frac{\omega x^2}{2cz}. \quad (3.12)$$

Hence, the condition for focusing $\Delta\phi + \Delta\phi_g = 0$ can be expressed as

$$\begin{aligned} \frac{e^2}{2m_e\epsilon_0\omega c}\nabla^2\text{TEC} &= \frac{\omega}{2cz}, \\ \omega^2 &= \frac{ze^2}{m_e\epsilon_0}\nabla^2\text{TEC}. \end{aligned} \quad (3.13)$$

From equation (3.13) it is clear that lower frequencies will be focused at closer distances z from the lens, and that the focusing power of the lens is purely a function of the TEC curvature $\nabla^2\text{TEC}$. It is also clear that for focusing to occur at a positive distance z (i.e. for the lens to be a converging lens) requires a positive $\nabla^2\text{TEC}$, corresponding to a minimum of the TEC.

One important characteristic of lenses in general is the intensity distribution at distances greater than the focal length (or equivalently for frequencies below the focal frequency when observing at a fixed distance). In the geometric optics limit considered above, the lens produces two or more discontinuities in intensity with respect to transverse distance x at which the intensity tends to infinity (e.g. Clegg et al., 1998). These are known as caustics, and an example of how they arise is shown in Figure 3.1. While in reality wave effects prevent the existence of such discontinuities or infinite magnification there are still dramatic variations in the observed intensity (e.g. Shi and Xu, 2021). The term caustics generally refers to the structure of the intensity distribution in space at a single frequency, but radio astronomical observations are more generally made at




Figure 3.1. A schematic of the effect of a Gaussian concave (i.e. defocusing) lens on an incident plane wave signal in the geometric optics limit. The top panel shows the lens perturbation itself, the middle panel shows the rays traced from the lens plane to the observer plane, and the bottom panel shows the spatial intensity distribution that would be observed at a single frequency in the geometric optics limit. Reproduced from Clegg et al. (1998).

a single point as a function of frequency and time. In this case, assuming the lens is moving across the line of sight, the time dimension of the observation corresponds to the x co-ordinate, while the frequency dimension corresponds (albeit non-linearly) to the z co-ordinate. Observations of caustics in this frequency-time space are typically referred to as ‘spectral caustics’ (Koval et al., 2017) and examples of such features observed using LOFAR are presented in Chapter 5 of this thesis.

While the intensity distribution beyond the focus is spectacular and eye-catching, the distribution before the focus is possibly more significant, given that it is likely to

be observed substantially more frequently. In this case a given lens will induce an enhancement of intensity which is essentially coherent across frequency when observed at a given distance, albeit with increasing magnitude towards the lower frequencies for which the observer is closer to the focal distance. In the case of a sequence of such lenses with varying separation and focal strength moving across the line of sight, it is conceivable that this effect may be mistaken for scintillation, and care must be taken to avoid confusing the two effects. One key element distinguishing them is the degree of frequency coherence. As discussed in section 3.3.1, the coherence bandwidth of weak scintillation is of order the frequency (Narayan, 1992) and significantly lower for strong scintillation, meaning that variations would be expected to be coherent over approximately an octave of frequency at most. For refractive lensing on the other hand, the variations would be coherent over many octaves. Unfortunately, observations spanning such a wide bandwidth are often not available. In this case, distinguishing refractive lensing from scintillation requires an estimate of the velocity of the ionospheric structures with respect to the line of sight (which could be obtained from cross-correlating the signals of multiple receivers as discussed in Section 3.3.1) and an estimate of the line of sight distance to the structures (which could be obtained from an ionosonde or estimated based on climatological values for the heights of the E- and F-regions). With these parameters one can estimate the Fresnel frequency, and can thereby distinguish scintillation which occurs on timescales shorter than the Fresnel frequency from refractive lensing which occurs on longer timescales.

Chapter 4

LOFAR: The LOw Frequency ARray



Figure 4.1. The distribution of LOFAR stations across Europe, including the planned stations in Italy and Bulgaria. Figure reproduced from ASTRON.

The LOw Frequency ARray (LOFAR: van Haarlem et al., 2013) is a radio telescope centred in the North-East of the Netherlands, with international stations spread between Ireland and Latvia as shown in Figure 4.1. It is primarily intended for radio astronomy: surveying radio sources (e.g. de Gasperin et al., 2021; Heald et al., 2015; Shimwell et al., 2017), searching for transient signals (e.g. Carbone et al., 2016; Kuiack et al., 2021a) and observing highly redshifted signals from the early universe (e.g. Hothi et al., 2021; Yatawatta et al., 2013). However, it can also provide detailed information on the ionosphere due to the distortions it causes to astronomical radio signals. LOFAR’s wide geographic spread allows it to observe the ionospheric conditions across a large area, while the dense ‘core’ of stations in the Netherlands can also provide measurements with high spatial resolution. A major strength of LOFAR for observing the ionosphere is the broad frequency coverage available. For example, whereas GNSS satellites transmit on only a few discrete frequencies, LOFAR can operate in a broad band between 10 MHz and 240 MHz, allowing much more detailed interpretation of the ionospheric structures with a single receiver.

The LOFAR stations are divided into three types, each of which has a slightly different configuration as described in van Haarlem et al. (2013). These are: core stations, 24 stations located within a ~ 2 km radius in the north-east of the Netherlands; remote stations, 14 stations spread further across the Netherlands with a maximum baseline of ~ 65 km to the core; and international stations, located across several countries as shown in Figure 4.1 with baselines of up to 1,000s of km.

Each LOFAR station consists of two separate antenna arrays, the low band array (LBA) covering 10 – 90 MHz and the high band array (HBA) covering 110 – 240 MHz (core stations have two separate HBAs). Each of these consists of roughly 50-100 individual dipole antennas, which are combined by a digital beam former to direct the antenna to any pointing direction without any moving parts. The individual station

beams have a full width at half maximum of roughly $1 - 3^\circ$ for the HBA and $3 - 20^\circ$ for the LBA depending on observing frequency and type of station (van Haarlem et al., 2013). It is important to highlight that at the time of writing, LOFAR is undergoing an upgrade to ‘LOFAR 2.0’ (e.g. ASTRON, 2023; Edler et al., 2021) which will provide improved sensitivity and the possibility of simultaneous LBA and HBA observations. All results presented and discussed in this thesis were obtained using the original LOFAR system as described in van Haarlem et al. (2013).

4.1 Ionospheric Results from Observations of Astronomical Radio Sources

The initial discovery of extraterrestrial radio emissions was made by Karl Jansky in 1932 (Jansky, 1933). Despite the obvious significance of these observations, development of the field of radio astronomy was delayed by the Great Depression and subsequently the Second World War (Swarup, 2015). After the war, aided by the developments in radio technology that had occurred in the interim, there was a renewed interest in radio astronomy. With the discovery of discrete radio sources (e.g. Bolton and Stanley, 1948; Ryle and Smith, 1948) it quickly became apparent that these sources exhibited rapid fluctuations in intensity (e.g. Hey et al., 1946; Bolton and Stanley, 1948; Ryle and Smith, 1948), which were shown to be due to variations in the ionosphere (Ryle and Hewish, 1950). Hence, radio astronomy became a means of studying the ionosphere, in addition to the astronomical results that had initially motivated its development. The following sections provide an overview of the ionospheric results obtained from these observations, using a variety of instruments and techniques.

4.1.1 Early Observations

Observations of astronomical radio sources for the purposes of studying the ionosphere were initially carried out in Cambridge and Jodrell Bank in the late 1940s (e.g. Little and Lovell, 1950; Ryle and Hewish, 1950; Smith, 1950). Much of the initial motivation for this work was as much astronomical as it was ionospheric, with the focus being on distinguishing fluctuations in the signal resulting from the ionosphere from any inherent variability in the sources themselves (e.g. Smith, 1950). This and other early work was able to establish several characteristics of scintillation in the mid-latitudes, such as a pronounced maximum occurrence at night (Ryle and Hewish, 1950) and correlations with both spread-F (Mills and Thomas, 1951; Singleton, 1964) and sporadic-E (Dagg, 1957; Wild and Roberts, 1956a). It also highlighted systematic effects that arise in analysing scintillation data from astronomical radio sources, such as elevation dependent effects (e.g. Ryle and Hewish, 1950) and the difficulty of separating seasonal and diurnal variations using observations of a single source (e.g. Wild and Roberts, 1956b).

As well as scintillation, some observations showed characteristics of focusing from larger scale convergent-lens-like structures due to their coherence over a wide bandwidth, which often repeated quasi-periodically (e.g. Warwick, 1964; Wild and Roberts, 1956b). These were subsequently identified as the signatures of TIDs (see Section 2.5.1) and similar observations made using LOFAR will be discussed in detail in Chapter 5 of this thesis. Observations by Dagg (1957) were found to be better explained by divergent-lens-like structures associated with sporadic-E layers.

Subsequently, with the advent of satellite beacons and the development of radar techniques for ionospheric research, interest in astronomical radio sources for studying the ionosphere waned. This occurred in parallel to a shift in radio astronomy due to the presence of the 21 cm neutral hydrogen (HI) line (van de Hulst, 1945) which motivated the development of higher frequency instruments (e.g. Baars and Kärcher, 2017; Lovell,

1957; Norris, 2017). At these higher frequencies ionospheric effects are less significant, which allowed for more accurate astronomical observations but meant there were fewer incidental opportunities for ionospheric research using instruments intended for radio astronomy.

More recently, interest within the radio astronomy community has returned to observing at lower frequencies. A major motivation for this is the desire to study the early universe at and beyond the ‘Epoch of Reionisation’ (the point at which neutral hydrogen was ionised by UV radiation, estimated to occur at a redshift of around $z \sim 6$ or earlier pushing the HI line to frequencies below ~ 250 MHz, e.g. Zaroubi et al., 2012). This requires low frequency telescopes to be built to observe this emission (e.g. van Haarlem et al., 2013). Much of the hardware for observing at these frequencies was already well developed, and indeed these were the frequencies at which the majority of early radio astronomy was conducted (e.g. Jansky, 1933; Ryle and Hewish, 1950; Wild and Roberts, 1956b). However, a central challenge of producing high-resolution images at these frequencies are the distortions introduced by the ionosphere, in addition to instrumental effects which are a problem across all frequencies (van Haarlem et al., 2013). The ionospheric effects are increased in significance both by the lower frequencies themselves, and the longer interferometer baselines required to achieve the desired angular resolution at these frequencies, which allow much greater variations in ionospheric conditions between interferometer elements.

Various strategies for accounting for these distortions have been developed, generally referred to as interferometric calibration. These techniques initially only accounted for direction independent effects (i.e. those that are constant across the field of view of the telescope, e.g. Cornwell and Wilkinson, 1981; Intema et al., 2009), but more recently have been extended to include direction dependent effects (e.g. de Gasperin et al., 2019, 2020a; Smirnov, 2011b). The inclusion of direction dependent effects in the calibration

is especially significant from an ionospheric perspective as it requires separating the contributions to the calibration solution based on their cause. Some instrumental effects are direction independent and therefore only require a single value for the entire field of view, whereas other effects, including ionospheric distortions, require solutions which vary across the field of view (Smirnov, 2011b). This means that modern interferometric calibrations for radio astronomy explicitly calculate ionospheric parameters and provide a unique dataset for ionospheric research (e.g. Mevius et al., 2016; de Gasperin et al., 2018; Beser et al., 2022).

4.1.2 Observations Using LOFAR

Using data from the Kilpisjärvi Atmospheric Imaging Receiver Array (KAIRA, an instrument in Arctic Finland built with LOFAR hardware: McKay-Bukowski et al., 2014), Fallows et al. (2014) showed that ionospheric scintillations can give rise to the characteristic ‘parabolic arc’ in the secondary spectrum (the 2D Fourier transform of the dynamic spectrum, also commonly referred to as the ‘Delay-Doppler spectrum’). Parabolic arcs were initially observed in scintillation from the interstellar medium (e.g. Stinebring et al., 2001), and arise from scattering from a thin screen containing a spectrum of irregularities (Cordes et al., 2006). The curvature of the arc is purely determined by the distance to the screen and the velocity of the irregularities, making this a powerful method of investigating ionospheric structure.

It is difficult to disentangle the velocity and distance contributions to the curvature when measuring from a single station. However, Fallows et al. (2020) used multi-station observations of scintillation with LOFAR to calculate the velocity from cross-correlation analysis. This allows the curvature of the arc to be used to calculate the distance to the screen, and hence the corresponding altitude within the ionosphere. In this case, they were able to isolate two distinct plasma populations, with differing velocities,

which gave rise to two distinct parabolic arcs in the secondary spectrum. These were shown to lie in the D- and F-regions respectively, and are the first reported observation of two simultaneous independent TIDs propagating at different altitudes in different directions.

Using observations from a single LOFAR station in Poland, Błaszkiwicz et al. (2020) found a negative correlation between the observed pulsar flux density and the rate of TEC index (ROTI) calculated from GNSS data. They ascribe this to ionospheric refraction, which moves the apparent position of the pulsar away from its expected position. This means it moves away from the LOFAR beam centre, to a position at which the sensitivity is lower. They showed that this correlation was localised to lines of sight passing through the ionosphere within ~ 20 km of one another, indicating the scale size of the irregularities responsible.

Given the widespread use of GNSS signals to study scintillation and irregularities in the ionosphere (e.g. Jin et al., 2016; Spogli et al., 2013), it is useful to compare results from LOFAR to those from GNSS to understand the differences in sensitivity between the instruments. Such a comparison was carried out by Flisek et al. (2023), who showed that LOFAR frequently observed scintillation between $\sim 22 - 76$ MHz on magnetically disturbed days, but there was no corresponding increase in scintillation on GNSS frequencies. There was, however, an increase in ROTI measured by the GNSS receivers in some of these cases. They interpreted these observations as showing that irregularities were generated but above the Fresnel scale for GNSS (~ 400 m at F-region altitudes) which were still able to induce scintillation in LOFAR observations due to the larger Fresnel scale associated with the lower frequencies.

The ionospheric signatures in many observations such as those discussed above are characteristic of scintillation. However, some observations show variations which suggest distinct coherent ionospheric structures are responsible. For example, Dorrian

et al. (2023) presented an observation in which an undisturbed signal suddenly transitioned to a regular pattern of alternating signal enhancements and depletions, before returning to its quiet state. This occurred in multiple sources and was observed from two international stations separated by ~ 500 km. By comparison to GNSS TEC data, they were able to show that this feature in the LOFAR observations corresponded to an MSTID, and that the timescale of the signal enhancements corresponded to refractive focusing from substructure within the TID on scales of ~ 20 km. These features in the LOFAR data also persisted for up to an hour after the TID structure in the GNSS data had broken up.

Other observations of coherent ionospheric structures have been reported by Trigg et al. (2024), who showed multiple cases in which a series of quasi-periodic signal fades were observed across the whole observing band with interference like fringing preceding and following them. These features repeated with periods of $\sim 1 - 2$ min and a sequence of 17 such fades was observed in one case. Although parabolic arcs were present in the secondary spectra, this did not allow direct estimation of the altitude of the structures from the LOFAR data as the features were only observed by individual international stations and so no velocity estimate was available. Comparison to nearby ionosondes showed that sporadic-E layers were present in both cases, but there was also significant F-region ionisation which prevented any definitive conclusions being drawn from these observations. The data were replicated using a phase screen model (see Chapter 5 for details) containing a series of Gaussian phase perturbations. These perturbations had horizontal scales of ~ 1 km or less and amplitudes of ~ 1 mTECu, demonstrating the sensitivity of LOFAR to even very small perturbations in the ionosphere.

A similar observation to those in Trigg et al. (2024) was reported by Wood et al. (2024), although in this case there was only a single signal fade of much longer duration (~ 10 min) and the fringing was much more developed with a pronounced asymmetry

between the fringing before and after the signal fade. As this feature was observed by stations in the Netherlands, the velocity could be estimated and it was found to be quasi-stationary, with its apparent motion dominated by the changing position of the radio source in the sky. The structure was shown to evolve significantly over timescales of tens of minutes and spatial scales of tens of kilometres, and to have an anisotropic structure with a much greater extent in the east-west than north-south direction. In this case, ionosonde measurements showed a sporadic-E layer which varied substantially with time, and the amplitude of those variations was shown to be consistent with the structure causing the variations observed by LOFAR.

As well as analysis of intensity spectra from individual stations, the combined interferometric observations of the LOFAR network can be used to identify ionospheric effects. For example, Kuiack et al. (2021b) observed brief (~ 10 s) radio flashes coinciding with known faint (below their detection threshold) radio sources in narrow-band LOFAR observations around 60 MHz. By comparing delay across frequencies and other considerations, they were able to show that these likely originate from focusing within the ionosphere, such as from the electron density minima of a TID. This shows that even without using the wide bandwidth, LOFAR observations can reveal clear ionospheric signatures.

Dedicated ionospheric studies using LOFAR interferometric data have focused on the output of the calibration process mentioned above. The primary output of interest from the calibration is differential TEC (dTEC), the difference in TEC between the lines of sight from two different stations. Mevius et al. (2016) used the dTEC variance on different baselines to investigate the anisotropy of ionospheric structure and its variation with spatial scale, based on observations using the HBA. They found power law behaviour with a typical spectral index of -1.89, which they interpret as resulting from a mixture of turbulent behaviour (expected spectral index -1.67) and coherent

wave-like structures (expected spectral index -2). They also showed that the majority of observations exhibited significant anisotropy, with the preferred orientation of the anisotropy along the geomagnetic field.

de Gasperin et al. (2018) were able to extend this approach to data from the LBA, finding consistent results to Mevius et al. (2016). They also demonstrated a method for estimating absolute TEC from the calibration solutions, by combining the dTEC and differential Faraday rotation (dFR) measurements with a model of the geomagnetic field. Comparison with global TEC maps derived from GNSS data showed that the LOFAR values were consistent, but the necessary simplifying assumptions used to derive it may mean that LOFAR is not able to achieve any higher precision than GNSS in absolute TEC.

To focus more on individual transient structures that are present in the dTEC solutions, Beser et al. (2022) used wavelet analysis to identify dominant directions of disturbances across a range of timescales. They were able to show that these dominant directions tracked variations in the geomagnetic field over the course of the observation, and variations in F-region drift velocity. However, this analysis did not attempt to determine any other properties of the disturbances, such as wavelength or amplitude. A complementary approach that makes certain simplifying assumptions to enable the estimation of wavelengths and amplitude is presented in Chapters 6 and 7 of this thesis.

A further avenue for studying the ionosphere with LOFAR is the measurement of Faraday rotation, which occurs when polarised radio waves pass through a magnetised plasma. As this is an integrated effect along the line of sight, it is dominated by the contribution of the interstellar medium, but on yearly and shorter timescales the variability is dominated by the near-Earth plasma (ionosphere and plasmasphere, the approximately fully ionised plasma extending above the ionosphere: Porayko et al., 2019). By comparing the Faraday rotation in pulsars measured by LOFAR to that

predicted based on global maps derived from GNSS TEC measurements and a model of the Earth’s magnetic field, Porayko et al. (2023) were able to show that a single layer model was inadequate. Instead, some information on the vertical structure of the ionosphere and plasmasphere was necessary to explain the observations.

4.1.3 Observations Using Other Radio Telescopes

The Very Large Array

Many studies of the ionosphere have been carried out using interferometric data from the Very Large Array (VLA: Thompson et al., 1980) in New Mexico. The derivation of dTEC from the VLA uses a significantly different methodology than for LOFAR, due to the much narrower bandwidth available ($\sim 1.5 - 3$ MHz depending on observing frequency). As different effects cannot be separated based on their frequency dependence, the dTEC is derived by subtracting a slowly varying baseline value from the residual phases, based on the assumption that instrumental effects will vary over timescales of hours and anything more rapid than that must be ionospheric in origin (Helmboldt et al., 2012b). By comparing the dTEC estimates from different frequency bands and polarisations, the results can be shown to be consistent with one another, but there may still be some residual contribution from instrumental effects. It also inevitably limits analysis to only variations with timescales $< \sim 1$ hour.

An example of what has been achieved using data from the VLA is the work of Jacobson and Erickson (1992), who used Fourier analysis of dTEC data derived from 46 hours of VLA observations to identify various wave signatures. This included the previously unreported phenomenon of Magnetic East Directed (MED) waves, which appeared to be short period, fast moving waves closely aligned with magnetic East. Later work (Hoogeveen and Jacobson, 1997; Jacobson and Erickson, 1993) showed that

these structures were actually located in the plasmasphere. The high apparent speed of these waves was a result of their approximate co-rotation with the Earth relative to the stationary radio source. Similar structures observed by LOFAR are described in Appendix F of this thesis.

Further work using dTEC derived from the VLA was carried out by Helmboldt et al. (2012a,b). They utilised the layout of the VLA, which consists of 3 distinct straight ‘arms’, to define both an overall TEC screen across the whole array capturing larger scale variations, and individual 1-dimensional variations along each arm to capture smaller scale features. They also derived an improved approach to identifying waves (Helmboldt et al., 2012a) which relaxed constraints that had been imposed by Jacobson and Erickson (1992). Their results were shown to correlate with both MSTIDs and quasi-periodic echoes in the E-region on a climatological basis. These techniques were further used by Helmboldt and Intema (2012) to identify disturbances which they associated with plasma flows from the plasmasphere into the nighttime ionosphere.

Coker et al. (2009) used dTEC from the VLA in combination with optical measurements from an all-sky camera to confirm that the presence of wave-like structure in the VLA observations coincided with observed gravity wave activity in the mesosphere. They also considered the effects of sporadic-E layers, and showed using (non-simultaneous) observations from an incoherent scatter radar that such layers could induce significant perturbations to dTEC on timescales ranging from 1 – 20 min. This was significant as previous work (e.g. Jacobson and Erickson, 1992) had assumed that structures with periods significantly less than 10 min could not be ionospheric as this is below the Brunt-Väisälä period for the F-region, and neglected to consider E-region structures.

A different approach to studying the ionosphere using measurements from the VLA was taken by Cohen and Röttgering (2009), who used measurements of the differential

refractive shifts between sources detected within the field of view to study TEC gradients across an area larger than the VLA itself. They showed that elevation effects in the data could be statistically removed by considering the ionosphere as a thin shell at an altitude of 400 km indicating that their observations were dominated by structures in the F-region. They found that the magnitude of the TEC gradients was higher during the day than at night, which is expected given the increased plasma densities, and also that the diurnal variations were most significant for the largest spatial scales.

These refractive shift measurements were used in combination with single source dTEC measurements by Helmboldt et al. (2012c) to infer a climatology of MSTIDs. This included many associated with orographic gravity waves and some nighttime waves with a clear preference for south-westward or north-eastward propagation. These nighttime waves were further investigated by Helmboldt (2012), who showed using ionosonde measurements that all such waves were observed in the presence of sporadic-E layers and typically under conditions of bulk downward motion in the F-region. This provided direct observational evidence to support the hypothesis that this nighttime directional preference was associated with the Perkins instability (see Section 2.5.2), with its growth rate amplified by a coupling with an instability in a sporadic-E layer beneath (Cosgrove et al., 2004; Kelley et al., 2003; Tsunoda and Cosgrove, 2001).

The Murchison Widefield Array

In addition to the VLA, ionospheric work has been carried out using observations from the Murchison Widefield Array (MWA: Tingay et al., 2013) in Western Australia. This is distinct from other instruments like the VLA and LOFAR as it uses much shorter baselines (maximum baseline ~ 3 km, similar to the LOFAR core but with more antennae). Naturally, this limits the information that can be obtained from comparing information from the same source from different locations. However, the

field of view available with the MWA is much larger than other radio interferometers ($\sim 25^\circ$ beam width (Tingay et al., 2013) compared to $\sim 2.5^\circ$ for LOFAR at similar frequencies (van Haarlem et al., 2013)). This wide field of view allows studies to use the apparent motion of sources in various parts of the field of view to characterise ionospheric variations across the field of view rather than across the size of the array itself.

Using this approach of considering the apparent motion of many (~ 1000) sources across the field of view, Jordan et al. (2017) found a moderate correlation between magnitude of apparent motions and their direction anisotropy (i.e. larger source displacements tend to occur due to more coherent structures). They also defined a metric for the ionospheric quality of astronomical observations. The metrics derived by Jordan et al. (2017) were used by Waszewski et al. (2022) in conjunction with intensity scintillation measurements. These measurements are sensitive to different scales (baseline ~ 2 km and Fresnel scale ~ 300 m respectively) meaning that their relationship can indicate the transfer of energy between different scales in the ionosphere. The correlation between refractive shifts and scintillation was shown to be strong and broadly consistent with the relationship expected for pure Kolmogorov turbulence, although it was not possible to uniquely determine the spectral index from this data.

Another study (Loi et al., 2015a) found that, under quiet geomagnetic conditions, night-time position offsets occurred favourably along the NW-SE direction, consistent with the observations from the VLA discussed above (Helmboldt et al., 2012c; Helmboldt, 2012) and from GNSS data in Europe (Otsuka et al., 2013). Similar analysis of a larger set of observations by Helmboldt and Hurley-Walker (2020) also found this preference for NW-SE propagation, and showed that these disturbances were correlated with sporadic-E detected at a nearby ionosonde as reported by Helmboldt (2012). Two other recurring populations were detected, one with fronts aligned N-S (i.e. essentially

field aligned and consistent with the MED waves of Jacobson and Erickson (1992)) and one with fronts aligned E-W and large spatial scales which was shown to correlate with variability in the subtropical jet stream.

Using data from a single observation, Loi et al. (2015b, 2016b) were able to show the presence of density ducts in the upper ionosphere/lower plasmasphere aligned with the geomagnetic field. These are likely the same structures reported by Jacobson and Erickson (1992) using the VLA, but in this case a direct measurement of the altitude of the structures was available. This was achieved by dividing the MWA into two halves (east and west) and then comparing the parallax between the position offsets measured from the two sub-arrays to determine the altitude of the structures responsible. A further case study (Loi et al., 2016a) showed these structures developing following the passage of a large amplitude TID during undisturbed geomagnetic conditions, suggesting a possible causal link.

The Deployable Low-Band Ionosphere and Transient Experiment

As a development of the work done with dedicated astronomical instruments to study the ionosphere, Helmboldt et al. (2021) developed the Deployable Low-Band Ionosphere and Transient Experiment (DLITE). This is a compact four element radio interferometer designed to minimise cost and maximise ease of deployment for use in studying the ionosphere. It probes the ionosphere by measuring both intensity scintillation and apparent displacement of bright sources such as Cassiopeia A and Cygnus A. This apparent motion, especially over such a compact array, can be reliably converted to a gradient in TEC across the field of view. In conjunction with ionosonde data, this system was able to establish a strong correlation between TID activity as measured by source displacement and km-scale ionospheric irregularities responsible for intensity scintillation (Helmboldt and Zabotin, 2022). Scintillation measurements using this sys-

tem were also shown to correlate with variations in coherence of radio signals reflected over the horizon by the ionosphere (Helmboldt, 2023).

4.2 LOFAR Ionospheric Scintillation Observations

Dedicated ionospheric observations with LOFAR are taken in ‘beamformed’ observing mode, in which every station observes a given source independently and records the received intensity. While LOFAR is capable of observing across a very wide frequency range, to ensure adequate resolution in frequency it is necessary to use only a subset of the full coverage for a given observation. For ionospheric observations this is typically between roughly 25 – 65 MHz, with 200 frequency channels giving ~ 200 kHz resolution. This is selected to avoid regions with significant radio frequency interference (RFI), which tends to dominate above ~ 80 MHz due to the FM band and below ~ 25 MHz. The intensities are integrated over ~ 10 ms, sufficiently short that ionospheric scintillation will not be smoothed out (Fallows et al., 2014) and that RFI can be distinguished.

As well as the frequency and time resolutions, selection of suitable sources is critical to observing ionospheric structures with LOFAR. The primary consideration is that the source must be bright and compact. This is necessary because the beam of a single station at these low frequencies has a large angular size and so records significant intensity from background sky temperature emission which is not affected by ionospheric variability (van Haarlem et al., 2013). The source must be compact so that it is of a smaller angular size than typical ionospheric structures. The two sources that best fit this criteria are Cygnus A (3C405) and Cassiopeia A (3C461), with the added benefit that they are circumpolar (always above the horizon) from almost all LOFAR stations. This means that observing windows are not artificially limited by the choice of sources,

although the variation in source elevation with time and season is a significant factor that must be accounted for in any statistical analysis (see e.g. Chapter 8).

This observing approach provides dynamic spectra of intensity as a function of frequency and time for each source at each station. However, these contain information besides just the impact of the ionosphere, such as: RFI, variation between channels due to the antenna bandpass and source spectrum, and slow time variations due to the elevation dependence of the antenna sensitivity. A procedure to remove these effects is outlined in Fallows et al. (2020). First, to identify RFI, the data is flattened by dividing by a median filter with a footprint smaller than typical time and frequency scale of ionospheric variability (e.g. $19.5 \text{ MHz} \times 4.2 \text{ s}$ in Fallows et al., 2020). Any point in this flattened array which exceeds the median by more than 5 standard deviations is then flagged as RFI. If a channel contains 20% or more RFI it is removed, and remaining isolated RFI points are replaced by linear interpolation from the surrounding data in both time and frequency. Once the RFI has been removed, the variation between channels and across long time periods is removed by dividing each channel by a fitted 3rd order polynomial.

Due to the (comparatively) low computational resources required, ionospheric observations with LOFAR are typically carried out as ‘filler time’, meaning that they are carried out whenever no astronomical observations are scheduled. This leads to a stochastic distribution of observations with time, rather than any systematic selection for specific observing conditions. This is beneficial for characterising the range of possible ionospheric signatures that can be captured by LOFAR, but with the downside that extreme conditions cannot be characterised due to a lack of observations.

Unlike techniques for probing the ionosphere with artificial radio signals such as GNSS and radar, the natural signals LOFAR uses are affected by plasma besides the ionosphere. For example, scintillation can be observed from the solar wind (e.g. Fallows

et al., 2016, 2023; Tokumaru et al., 2019) and interplanetary medium (e.g. Briske et al., 2009; Cordes et al., 2006; Wu et al., 2022). As a result, care must be taken to ensure that the region responsible for the observed signal is identified correctly. This is often possible from a single receiver due to the inferred properties of the scattering medium being unphysical (e.g. too high density / velocity / size) for a given medium (e.g. Fallows et al., 2016). With an array of receivers such as LOFAR, cross-correlations can be used to estimate the velocity of the structures as discussed in Section 3.3.1. Ionospheric effects have significant travel time even between stations in the Netherlands, and are often localised to a subset of the LOFAR array (e.g. Dorrian et al., 2023; Trigg et al., 2024), whereas interplanetary or interstellar scintillations will appear almost simultaneously across the whole array. Combining the observations in this way gives another means of distinguishing variations originating from the ionosphere from those originating at greater distances.

These LOFAR observations can complement GNSS data as they cover different frequency ranges and therefore observe different plasma effects on propagation due to the difference in Fresnel scales (\sim a factor of 10) and LOFAR frequencies being much closer to typical plasma frequencies. This means that together they are able to probe a wider range of ionospheric structures than either can individually. With the ongoing (at time of writing) upgrade to LOFAR 2.0 (e.g. ASTRON, 2023; Edler et al., 2021) LOFAR will become an even more powerful instrument for observing the ionosphere, as this will increase the effective collecting area for the LBA by a factor of two leading to increased sensitivity and enable simultaneous observation with the LBA and HBA among other improvements.

Chapter 5

Case Study: Inferring Small-Scale TIDs with LOFAR

This chapter describes a case study of a LOFAR observation primarily using data from the UK station, which has been published in the Journal of Space Weather and Space Climate (Boyde et al., 2022). Features are observed which repeat pseudo-periodically, and using a single phase screen model it is shown that they can be explained by focusing from an SSTID (wavelength ~ 20 km, period ~ 10 mins). For some of the features, the modelling indicates that a linear (i.e. single sine wave) model for the TID is insufficient, showing that LOFAR is capable of resolving structure of such TIDs even on scales below the wavelength. Model parameters are compared to complementary observations from the Dutch LOFAR stations, GNSS receivers and an ionosonde to confirm that the selected values are physically reasonable. The observed features appear to be a subset of a more general class of features referred to as ‘spectral caustics’ (Koval et al., 2017), but the short wavelength of the TID leads to interference effects that are not present in previous observations.

5.1 LOFAR Observations

On the 15th of September 2018, between 10:15 and 11:48 UT (observation ID L667596 under project code LT10_001), LOFAR was used to observe the radio sources Cygnus A (3C 405: RA 19h59m28s, Dec. 40.73°) and Cassiopeia A (3C 461: RA 23h23m24s, Dec. 58.82°) as described in Section 4.2. The data was processed to remove RFI and other non-ionospheric effects as described in Section 4.2, using a filter size of $1.95 \text{ MHz} \times 0.5 \text{ s}$.

The geomagnetic conditions at the time of the observation were quiet, with a Kp index of 2-. The observation was taken at solar minimum, with an F10.7 solar flux index of 69 sfu. No solar flares were reported in the preceding days based on X-ray flux measurements from the GOES satellites (accessed via www.solarmonitor.org). This absence of obvious external drivers suggests that the observations are most likely associated with purely terrestrial processes.

Figure 5.1 shows the dynamic spectrum of Cygnus A observed from the UK LOFAR station (UK608: 51.1°N , 1.4°W , 176 m above sea level). Cygnus A was at a very low elevation, rising from 4.8° to 11.7° over the course of the observation and moving from 19.9° to 36.3° in azimuth clockwise from North. The features are pseudo-periodic, with one appearing roughly every 10 minutes over approximately an hour. These repeated intensity enhancements suggest focusing from a succession of electron density minima, such as would be associated with a TID and has been reported in solar radio observations by Koval et al. (2019). The features consistently display a broadening towards lower frequencies with clear interference fringes, as well as a time asymmetry with lower frequencies leading the higher frequencies.

Some individual features are shown in Figure 5.2, illustrating the fine structure that is present. Although the fringes at lower frequencies are consistent in all features, there are significant qualitative differences. The type of feature appears to alternate, as features 2 and 4 both show approximately non-dispersive (i.e. simultaneous across




Figure 5.1. The filtered dynamic spectrum of Cygnus A from the LOFAR station in the UK (station UK608: 51.1°N, 1.4°W, 176 m above sea level). The frequency scale is inverted, with lower frequencies at the top of the plot. The normalised intensity is the intensity after the filtering process, with a normalised intensity of 1 being the expected value for an undisturbed ionosphere. Individual features are numbered to distinguish them in further discussion. The horizontal white lines are the channels removed due to RFI contamination.

a range of frequencies) intensity enhancements at higher frequencies, with a sharp transition to dispersive behaviour below ~ 40 MHz, similar to a transition between weak and strong scattering behaviour (e.g. Rino et al., 1981), and the presence of interference fringes in the lower frequencies. Conversely, features 1, 3 and 5 do not share this sharp transition, and they display a much more symmetrical form in the low frequencies. Features 1 and 5 also show a clear splitting into a ‘doublet’ at frequencies above the apparent focal frequency (lower on Figure 5.2), but it is unclear if this is also the case for feature 3 due to the proximity of the focal frequency to the upper frequency limit of the observation.

Besides the lower frequencies leading the higher frequencies in time, another more




Figure 5.2. The dynamic spectra of the 5 features highlighted in Figure 5.1. Note that each panel uses a different colour scale for intensity to provide maximum contrast in each panel, and that each has a different time scale.

subtle asymmetry in the features is the fringe spacing on the leading and trailing edges of the envelope in features 1, 3 and 5. In features 3 and 5, the fringes are more closely spaced in frequency along the trailing edge of the feature (later in time) than the leading edge (earlier in time) whereas for feature 1 the opposite is true. This indicates some asymmetry in the ionospheric structures responsible for the features, which is not considered further within this case study.

To the best of the author’s knowledge, features such as these in dynamic spectra have not been reported elsewhere in such detail. The ‘spectral caustics’ described by Koval et al. (2017) in solar radio observations have a similar overall form, but do not include interference fringing. In their classification, feature 3 would be described as ‘inverted V like’, features 1 and 5 as ‘X like’ and features 2 and 4 appear similar to

their ‘fibre like’ caustics. Many of their observations have been shown to correlate with the minima of passing MSTIDs (Koval et al., 2019) and they were able to explain at least the ‘inverted V like’ caustics as arising from focusing from the minima of such TIDs (Koval et al., 2018). The absence of fringes in their observations is likely due to the larger angular size of the Sun compared to Cygnus A as will be explored in subsequent modelling.

The observations of Kuiack et al. (2021b) of source magnification between 57.6 – 62.5 MHz using LOFAR show an example which may reflect the same phenomenon observed here, although it is difficult to be sure in such a narrow band observation. They reported one case with a main peak of intensity with secondary peaks to either side, described well by a first order Bessel function of the first kind. Most of their observations do not display the secondary peaks, consistent with the observations of Koval et al. (2017), suggesting that the features shown in Figure 5.2 are a special case of a more general phenomenon. Kuiack et al. (2021b) argue that these magnifications were most likely a result of focusing from electron density minima in the ionosphere, such as TIDs. Another possible example of similar features in early radio astronomical observations is provided by Wild and Roberts (1956b) in what they refer to as ‘ridges with fine structure’, although in their observations the lower frequency cutoff of 40 MHz prevented the full structure at low frequencies from being apparent.

5.2 Modelling

5.2.1 Analytic Modelling

Some of the features observed here using LOFAR are strongly similar to those simulated using an analytic single phase screen model by Meyer-Vernet (1980). This model replicates the envelope broadening towards lower frequencies and the pattern of over-

lapping fringes, by assuming an infinite one-dimensional thin phase screen containing a single sinusoidal variation. With some minor simplifying approximations, Meyer-Vernet (1980) showed that an analytic solution can be found for the observed intensity, meaning that the parameter space can be quickly and easily sampled to compare the model to the observed dynamic spectrum. The parameter space can be constrained by other observations such as ionosondes and GNSS TEC measurements, as will be explored in Section 5.2.2.

Table 5.1. The definitions of the variables used in the analytic phase screen model.

Variable	Definition	Unit
Φ_0	The amplitude of the phase variation	rad Hz
d	The length scale of the TID (wavelength Λ divided by 2π)	m
L	The distance from phase screen to observer	m
v	The TID propagation velocity	m s^{-1}
$\Delta\theta$	The source angular size	rad
α	The phase gradient of the background ionosphere	rad Hz m^{-1}

The model uses co-ordinates x and z , where x is aligned with the phase screen and z is normal to it along the line of sight, with the source at $z = -\infty$ and the screen at $z = 0$ as shown in Figure 5.3. The phase change $\Delta\phi$ imparted by a given point x on the screen at time t is given by

$$\Delta\phi(x, t) = \frac{\Phi_0}{f} \cos\left(\frac{x - vt}{d}\right) + \frac{\alpha x}{f} = \frac{\Phi_0}{f} \cos(X(x, t)) + \frac{\alpha x}{f}, \quad (5.1)$$

where f is the frequency of the radio wave and the other variables are defined in Table 5.1. The inverse frequency dependence of the phase screen amplitude arises from equation (3.6) and relies on the same basic assumptions. The variation in the background ionosphere (represented by α) is assumed to be linear as this is the most general form that can be included while retaining the analytic solution.

The amplitude of the phase perturbation Φ_0 and the phase gradient α can both

naturally be related to the TEC via equation (3.6). This gives relationships of

$$\begin{aligned}\Delta\text{TEC} &= \frac{4\pi m_e \epsilon_0 c}{e^2} \Phi_0, \\ <\nabla\text{TEC}> = \frac{-4\pi m_e \epsilon_0 c}{e^2} \alpha,\end{aligned}\tag{5.2}$$

where ΔTEC is the amplitude of the TEC perturbation of the wave and $<\nabla\text{TEC}>$ is the background TEC gradient along the wave propagation direction. The factor of -1 has been dropped from the relationship between ΔTEC and Φ_0 , whereas it is retained in the relationship between $<\nabla\text{TEC}>$ and α as the direction of the gradient is significant. The additional factor of 2π with respect to equation (3.6) arises because of the change from angular frequency ω to frequency f .




Figure 5.3. A schematic diagram of the model, showing the co-ordinate system and several parameters. Dashed arrows represent raypaths for a given frequency.

Rather than directly inputting a value for L , it is more physically meaningful to define the altitude of the phase screen h and use the known elevation of the source θ . The value of L can then be calculated from these and the radius of the Earth R_E as

$$L = -R_E \sin(\theta) + \sqrt{(R_E \sin(\theta))^2 + (R_E + h)^2 - R_E^2},\tag{5.3}$$

assuming that the Earth is a perfect sphere and the observer is at a negligible height above sea level. The source is assumed to have a Gaussian brightness distribution (intensity as a function of angle) of

$$I(\theta_k) \propto \exp\left(\frac{\theta_k^2}{\Delta\theta^2}\right), \quad (5.4)$$

where θ_k is the angle between the radio wavevector and the z -axis, and for Cygnus A the angular size $\Delta\theta$ is assumed to be 0.5 arcmin (Carilli et al., 1991; Skrutskie et al., 2006). While the actual structure of Cygnus A is not Gaussian, consisting of two distinct lobes, this simple model is adopted here to allow the analytic modelling.

The normalised intensity I that would be observed at the point (x, z) and time t is calculated in the Fresnel approximation. By assuming that the dominant contributions come from the regions of the screen close to the line of sight (i.e. small angles of deflection), the analytic expression for I can be shown to be

$$I = 1 + 2 \sum_{p=1}^{\infty} J_p \left(\frac{2\Phi_0}{f} \sin \left(\frac{pZ}{2} \right) \right) \cos(pX(x, t')) \exp(-(pu)^2), \quad (5.5)$$

where J_p is a Bessel function of the first kind, $Z = \frac{Lc}{2\pi f d^2}$ is the distance to the screen normalised by the radio wavelength and TID length scale, $t' = t + \frac{Lc\alpha}{2\pi f^2 v}$ is the time including a correction for the deviation of the raypath from the z -axis arising from refraction in the background ionosphere (Meyer-Vernet et al., 1981, see Section 3.3.3 for an explanation of this effect), and $u = \frac{\Delta\theta L}{2d}$ is the normalised angular size of the source (Meyer-Vernet, 1980). Given that the value of x chosen for the observer in equation (5.5) does nothing except shift the observed feature in time, $x = 0$ can be assumed without loss of generality. The infinite sum is in practice possible to compute due to the rapid decay of both the Bessel function term and the exponential for high p (Meyer-Vernet, 1980).

The two terms in the sum which both decay with high p provide the explanation for fringes appearing in certain cases and being absent in others. If the sum is truncated by the decay of the exponential term (i.e. by the finite size of the source) then the fringes are smoothed out. If, on the other hand, the Bessel function is primarily responsible for the truncation, then the fringes are observed. The characteristic p values at which the terms start to decay are $p \approx \frac{1}{u}$ and $p \approx \frac{e\Phi_0}{f}$ for the exponential and Bessel function terms respectively (e is Euler's number) (Meyer-Vernet, 1980). This means that the approximate condition for the presence of fringes can be expressed as

$$\frac{e\Phi_0}{f} < \frac{2d}{z\Delta\theta}, \quad (5.6)$$

meaning that for a given phase screen the presence of fringes requires the source size to be below a certain threshold.

In practice, the phase screen is not known a priori. Instead, the source size is known, and the focal frequency is the simplest physical observable. Based on the simple parabolic lens model introduced in Section 3.3.4, the focal frequency is expected to be dependent on $z\nabla^2\text{TEC} \propto \frac{z\Phi_0}{d^2}$ (see equation (3.13)), so the observed focal frequency will determine the quantity $\frac{z\Phi_0}{d^2}$. Equation (5.6) can be rearranged to provide an upper limit on $\frac{z\Phi_0}{d}$ for a given frequency. The combination of these two constraints indicates that the presence of fringes places an upper limit on the wavelength of the TID responsible, as to increase $\frac{z\Phi_0}{d}$ while holding $\frac{z\Phi_0}{d^2}$ constant requires an increase of length scale d .

In order to directly compare the model to observations, it is important to extend the model of Meyer-Vernet (1980) to account for the background intensity recorded by LOFAR. This is a combination of diffuse background emission and instrumental effects, and in the LBA corresponds to a total System Effective Flux Density (SEFD) of ~ 30 kJy (van Haarlem et al., 2013) which is unaffected by ionospheric irregularities. The

flux density of Cygnus A at the frequencies considered here is ~ 20 kJy (de Gasperin et al., 2020b). However, the received signal power from Cygnus A will be elevation dependent due to the difference in projected area of the receiver, whereas the power corresponding to the SEFD will be approximately elevation independent as the increase in field of view at low elevations counterbalances the reduced sensitivity. As a result, the observed intensity is given by

$$I = 1 + 2 \frac{I_{CygA} \sin(\theta)}{I_{CygA} \sin(\theta) + I_{sys}} \sum_{p=1}^{\infty} J_p \left(\frac{2\Phi_0}{f} \sin \left(\frac{pZ}{2} \right) \right) \cos(pX(x, t')) \exp(-(pu)^2), \quad (5.7)$$

where I_{CygA} is the flux density of Cygnus A, I_{sys} is the SEFD and θ is the elevation.

5.2.2 Applying the Model

In order to compare the observed features to this model, some parameters were fixed based on complementary observations discussed below, and others were adjusted to match the observed intensity distribution. Based upon these complementary observations, the screen altitude h , velocity v and source size $\Delta\theta$ were fixed to 200 km, 50 m s^{-1} and 0.5 arcmin respectively. For the remaining parameters: d and Φ_0 were manually adjusted to match the focal frequency and fringe spacing of the dynamic spectrum (lower values of d lead to wider fringe spacing for a given focal frequency), and α was set to match the time asymmetry that was observed. In order to give a quantitative comparison between the model and observations, several characteristic values were defined: the focal frequency, the frequency of the 5th fringe along each edge of the envelope (leading and trailing) and the time between the 5th fringe on the leading and trailing edges. The fringe location was defined as the maximum intensity within the fringe, and the values for each feature and their modelled replicas are given in Table 5.2.

The strength of this model in replicating the types of features observed by LOFAR




Figure 5.4. Feature 3 along with its modelled replica (note that different colour scales are used in each panel). The time axes cover the same duration, with $t = 0$ in the model corresponding to the minimum of electron density (i.e. the maximum of phase change) lying on the line of sight. The parameters used here are: $\Phi_0 = 8 \times 10^8 \text{ rad Hz}$, $\Lambda = 20 \text{ km}$, $\theta = 7^\circ$ and $\alpha = 1.89 \times 10^5 \text{ rad Hz m}^{-1}$.

is clearly illustrated by Figure 5.4. The overall envelope shape, pattern of overlapping fringes and time asymmetry are all present, and the required wavelength suggests that the features are caused by an SSTID ($\Lambda = 20 \text{ km}$). The modelled intensity variations also match very closely in magnitude to the observation. If the background is not accounted for, the modelled focal intensity is ~ 15 times the mean (not shown), demonstrating the importance of correcting for this effect. The periodicity of the model in this case is shorter than the observations, meaning that multiple features are shown, but the individual modelled features share their essential qualitative characteristics. In this case, the width of the modelled feature given in Table 5.2 is significantly lower than the observed feature, but this is due to the edges of the observed feature being approximately straight lines in time-frequency space, whereas the modelled feature curves outwards significantly, meaning that it is not possible to match the width across a range

Table 5.2. The characteristic values of observed and modelled features. f_{foc} denotes the focal frequency, f_{fringe} denotes the frequency of the fifth fringe (third for feature 5) on the leading and trailing edge, and Δt_{fringe} denotes the time between the fifth (third for feature 5) fringe on the leading and trailing edges.

Feature	f_{foc} (MHz)	f_{fringe} (Leading, Trailing) (MHz)	Δt_{fringe} (s)
Feature 1	40.81	35.15, 33.00	100.7
Feature 1 Replica	40.23	34.37, 34.37	93.8
Feature 3	54.88	34.56, 37.49	179.3
Feature 3 Replica	54.88	38.86, 38.67	115.3
Feature 5	38.67	28.70, 30.27 (3rd fringe)	89.1 (3rd fringe)
Feature 5 Replica	39.84	29.88, 29.88 (3rd fringe)	94.0 (3rd fringe)
CS001 Feature	43.35	30.07, 33.00	179.3
CS001 Feature Replica	41.99	31.44, 31.44	109.1

of frequencies. There is a small discrepancy in fringe spacing between observation and model as indicated in Table 5.2, however it is clear that features of this type can be reasonably qualitatively represented by this model.

Figure 5.5 shows several data sources that were used to constrain various parameters of the model. The ionosonde data in the left panel were used to constrain h , as TID amplitudes typically maximise around the altitude of maximum electron density (Fedorenko et al., 2011). This gave an approximate value of $h = (199 \pm 7)$ km based on the mean and variance of the manually scaled hmF2 values between 10:15 and 11:45 UT (the duration of the LOFAR observation), and so a value of $h = 200$ km was used for all modelling.

The TEC data from the MADRIGAL database (Rideout and Coster, 2006) were used to attempt to constrain α , as it provides in principle a maximum value if the TID propagation aligns perfectly with the TEC gradient. The gradients were calculated from the difference in TEC between two cells separated by 4° of latitude, selected to minimise the noise while ensuring the values are still sufficiently local (longitudinal gradients are neglected as they are approximately zero). All gradients between 10:15-




Figure 5.5. The complementary data used to constrain the model parameters. The leftmost panel shows the altitude of peak ionospheric electron density (hmF2) measured from the Juliusruh ionosonde (54.6°N, 13.6°E). The middle panel shows the distribution of North-South vTEC (vertical TEC) gradients (positive implies increasing to the North) calculated from MADRIGAL GNSS TEC based on the difference over 4° of latitude centred on points between 58°-61°N, 5°-7°E and 10:15-11:45UT. The third panel shows vTEC measured from the Southampton (SOTN) GNSS receiver (50.9°N, 1.5°W) for all azimuths between 0° and 60° clockwise from North.

11:45 UT centred within 58° – 61°N and 5° – 7°E are considered, corresponding to the observation of Cygnus A from UK608 assuming a thin shell ionosphere at an altitude of 200 km. The distribution of these gradients is shown in the middle panel of Figure 5.5, and were converted to α by converting from vertical TEC to the slant path observed by LOFAR (assuming an elevation of 7° corresponding to the third feature and a thick shell ionosphere between 150 – 500 km) and using equation (5.2). The mean and variance of the TEC gradients suggested that $\alpha = (1.20 \pm 0.39) \times 10^4 \text{ rad Hz m}^{-1}$ in the northward direction.

The GNSS TEC data from Southampton were used to constrain the amplitude of variations in the phase screen, as the TEC perturbation amplitude ΔTEC along the line of sight can be related to Φ_0 using equation (5.2). By inspection of Figure 5.5, the waves in the vTEC data have $\Delta\text{TEC} \approx 0.05 \text{ TECu}$, but the corresponding sTEC observed along the LOFAR line of sight cannot be precisely calculated as it is highly

sensitive to the relative orientation of the line of sight and TID phase front, especially for short wavelength TIDs. However, as an order of magnitude estimate this is still useful, suggesting $\Phi_0 \sim 5 \times 10^8 \text{ rad Hz}$.

As well as the data shown in Figure 5.5, the LOFAR observation itself can be used to constrain the velocity v of a structure if it is observed by several stations. The dense network of stations in the Netherlands provided observations of a single feature from all 24 stations in the LOFAR core, with an example shown in Figure 5.6. Following the approach described in Fallows et al. (2020), this allowed the propagation velocity of the structure to be estimated by considering the cross-correlations between stations to estimate the propagation time along each baseline (remote stations were not included as the feature evolved in frequency over these larger scales). Combining this gave an estimate for the velocity of the structure relative to the motion of the LOFAR lines of sight of $(57 \pm 5) \text{ m s}^{-1}$ at an azimuth of 104° clockwise from North. For all modelling the velocity was assumed to be $v = 50 \text{ m s}^{-1}$, although the velocity of the structures observed from the UK station is unlikely to be exactly the same as those observed from the Netherlands.

Correcting for the apparent source motion from each station made it possible to estimate the physical velocity of the ionospheric structure. Assuming an altitude of 200 km, this suggests a velocity of $(75 \pm 5) \text{ m s}^{-1}$ relative to the Earth. It is also necessary to project α onto the propagation direction, as this projected value is what is required by the model. In this case, this suggested a small negative value of $\alpha = (-2.9 \pm 1.0) \times 10^3 \text{ rad Hz m}^{-1}$ (neglecting any uncertainty contribution from uncertainty in the propagation direction).

The amplitude of the phase perturbation used in Figure 5.4 is consistent with that implied by the GNSS TEC measurements shown in Figure 5.5, but the value of the gradient term α is greater in magnitude and positive despite the expected small negative

value. This suggests that although the chosen value of α can approximately replicate the observed time asymmetries, the actual physical mechanism for the asymmetry is not a large-scale TEC gradient. One other possible explanation for the asymmetry is the low elevation of the source (7°), causing significant refraction due to the vertical gradients in electron density. Because this refraction is frequency-dependent, the different frequencies take different paths through the ionosphere and so ‘see’ the TID passing at different times. This would cause the lower frequencies to lead the higher frequencies if the structures propagate towards the observer. However, assuming a propagation azimuth of 104° is correct, given the azimuth of Cygnus A as viewed from the UK station the TID should be seen to propagate away from the observer, not towards them. An alternative explanation not explored further here is the alignment of the line of sight with the phase fronts of the TID, which was shown by Koval et al. (2018) to create significant asymmetries in their modelled dynamic spectra.

As well as the features observed from the UK station shown in Figure 5.2, one similar feature was observed in the spectrum of Cygnus A by many of the LOFAR stations in the Netherlands (which was used to estimate v above), with an example shown in Figure 5.6. This feature is the one which most closely matches the model in its form, with no doublet above the focal frequency and its asymmetry very closely replicated by the effect of the chosen value of α . However, this value of α is again inconsistent with the observed TEC gradients and velocity which imply a negative α , so one or more other physical mechanisms must be responsible. As in Figure 5.4, the feature can only be replicated in the model by an SSTID ($\Lambda = 25$ km). The absence of a doublet in both the model and this observation suggests that the doublet in features 1 and 5 may arise from some perturbation to the simple sine wave; this will be discussed in the following section. Unlike in Figure 5.4, the modelled intensity is noticeably higher than is observed, suggesting that the model may be overly idealised in many




Figure 5.6. The feature observed in the dynamic spectrum of Cygnus A by one of the stations in the Netherlands (station CS001: 52.9°N, 6.9°E), alongside the modelled replica (note that different colour scales are used in each panel). As in Figure 5.4, the time scale of both panels is the same, with $t = 0$ corresponding to the minimum of electron density (i.e. the maximum of phase change) lying on the line of sight. The parameters used here are: $\Phi_0 = 8 \times 10^8$ rad Hz, $\Lambda = 25$ km, $\theta = 8.8^\circ$ and $\alpha = 7.58 \times 10^4$ rad Hz m $^{-1}$.

cases. Also, despite the observed feature showing distinct curvature in the edges of its envelope unlike feature 3, the modelled replica is still significantly narrower as shown in Table 5.2. This could be corrected by increasing the TID length scale d , but this would cause the fringe spacing to decrease meaning that the model cannot precisely match the observation in this case (unless v is an overestimate).

5.2.3 Numerical Modelling

The analytic model derived by Meyer-Vernet (1980) provides a very clear replication of several of the observed features, but the most prominent aspect that is absent is the doublet above the focal frequency shown in Figure 5.2. This intensity distribution cannot be replicated by the single sine wave model, clearly suggesting that there is scope to extend the model. In order to investigate this, it was necessary to move from

the analytic model to a numerical phase screen approach.

The numerical phase screen approach was outlined in detail in Sokolovskiy (2001) and is described in Section 3.2.2 of this thesis, and allows for an arbitrary number of parallel screens to be considered. However, for the purposes of this work, the modelling was limited to a single phase screen, so as to minimise the number of free parameters. While multiple phase screens may provide a more detailed representation of the ionosphere, it will be shown that a single phase screen is nevertheless capable of replicating the features observed by LOFAR.

The use of numerical phase screen models in the ionosphere has previously focused on anthropogenic radio sources, such as ground- or space-based radar systems and GNSS signals (e.g. Carrano et al., 2020; Ding et al., 2021; Hocke and Igarashi, 2003; Ludwig-Barbosa et al., 2019; Wang et al., 2014), meaning that care must be taken in defining the input spectrum of a natural source as considered here. As the source is incoherent, each angular component must be treated separately, to avoid the spurious phase relationships between the components that occur if they are combined into a single spectrum. The intensity contributions from each angular component are then summed to give the total observed intensity distribution. As in the analytic model, the source is assumed to be Gaussian for simplicity, although any intensity distribution is possible in the numerical model.

The phase screen is sampled with 2 m resolution, to ensure that the full angular spectrum is available throughout the observing band of LOFAR (the upper frequency limit corresponds to a wavelength of 4.70 m). The screen contains 2^{18} samples, giving an extent of 524.288 km. This ensures that the wavelengths of the TIDs considered are far lower than the size of the screen, and provides a high angular resolution when the Fourier transform is applied. The numerical phase screen calculation gives intensity as a function of x , which is then converted to a function of t by assuming a constant

velocity v .

The similarity of the low frequency behaviour of features 1 and 5 to the analytic model suggests that the simple sinusoid is a good starting point. The doublet at high frequencies is indicative of focusing from two discrete points, corresponding to two minima of phase curvature, rather than the single minimum per period present with a single sine wave. The simplest way to modify the simple sinusoid to achieve this is to add a first harmonic of lower amplitude to equation (5.1), meaning that $\Delta\phi(x, t)$ becomes

$$\Delta\phi(x, t) = \frac{\Phi_0}{f} [\cos(X(x, t)) + A \cos(2X(x, t) + \phi)] + \frac{\alpha x}{f}, \quad (5.8)$$

where A and ϕ are the relative amplitude and phase offset of the harmonic respectively. As the observed features are approximately time symmetric (except for the distortion represented by α), the phase screen should retain its symmetry, corresponding to $\phi = 0$. This means that this modification to the phase screen introduces a single additional free parameter to the model, A . Although the comparison to the observed features in Section 5.2.2 suggests that the physical cause of the asymmetry was not the horizontal TEC gradients represented by α , it was retained as a simple and effective term to replicate the observations.

For the screen to contain two minima of phase curvature requires $A < -0.0625$. A further constraint is suggested by the fact that the intensity between the two peaks of the doublet is not suppressed (i.e. is ≥ 1), indicating that there is no de-focusing in this region. This lack of de-focusing indicates that the curvature is not positive at any point between its minima, corresponding to the constraint $A > -0.25$. An example of a perturbation satisfying these constraints is shown in the left panel of Figure 5.7 alongside the unperturbed phase variation, while the right panel compares the phase curvature in the two cases to illustrate the double minima.

The presence of significant harmonics in TID waves has been reported previously.




Figure 5.7. An illustration of the effect of the harmonic. The left panel shows the phase perturbations associated with the simple sine wave perturbation given by equation (5.1) (red dashed line) and the perturbation given by equation (5.8) (black solid line), with crosses marking the curvature minima. The right panel shows the phase curvature associated with each. Both perturbations are normalised to an amplitude and wavelength of 1 and have $\alpha = 0$. The harmonic has a relative amplitude of $A = -0.12$.

For example, van de Kamp et al. (2014) identified a TID using incoherent scatter radar and GNSS TEC maps, with the observed periods differing by a factor of two. They explained this as the two methods identifying different harmonics of the same underlying TID structure. Harmonic generation in the ionosphere has also been predicted in numerical modelling work by Kirchengast (1997), where it occurred even for a monochromatic driver due to non-linear processes such as frictional heating. While these works both focused on TIDs with periods on the order of an hour, theoretical work has shown harmonic generation, from both spectra of AGWs and non-linear ionospheric responses to monochromatic AGWs, without any assumptions about the period

of the waves considered (Chao-Song and Jun, 1991; Nekrasov et al., 1995; Nekrasov and Shalimov, 2002).




Figure 5.8. Feature 1 and its modelled replica. Both plots use the same time scale and $t = 0$ corresponds to the TID minimum lying on the line of sight (note that different colour scales are used in each panel). The parameters used here are: $\Phi_0 = 1.1 \times 10^9 \text{ rad Hz}$, $\Lambda = 30 \text{ km}$, $\theta = 5^\circ$, $\alpha = 1.45 \times 10^5 \text{ rad Hz m}^{-1}$ and $A = -0.12$.

The addition of this harmonic to the modelled phase screen is effective in replicating the doublet, as shown in Figure 5.8. The qualitative agreement between the model and observation remains strong, but as in Figure 5.6 the intensity is overestimated (the observed focal intensity is ~ 1.75 whereas the modelled value is ~ 2.8). Another discrepancy visible in Figure 5.8 is that the model predicts slightly higher intensities for the fringes along the centre of the envelope than on the boundaries, which is not apparent in the observation. Unlike Figures 5.4 & 5.6, the width of the observed feature matches very closely to the model in this case.

The addition of the harmonic is also able to replicate feature 5, as shown in Figure 5.9. The TID amplitude and wavelength are both lower than those used to replicate feature 1 (wavelength reduced by factor 2, amplitude by factor 3), which is required to




Figure 5.9. Feature 5 and its modelled replica. Again the time scale of the model and observation are the same, but the intensity scales used are different. The parameters used here are: $\Phi_0 = 3.5 \times 10^8 \text{ rad Hz}$, $\Lambda = 15 \text{ km}$, $\theta = 8.4^\circ$, $\alpha = 7.37 \times 10^4 \text{ rad Hz m}^{-1}$ and $A = -0.12$.

increase the fringe spacing while retaining approximately the same focal frequency. In this case, the qualitative agreement is not as strong as for feature 1 (Figure 5.8), as the modelled feature spreads out far more towards the low frequencies than towards the high frequencies, which is not the case for the observed feature. The overestimate in intensity for this feature is much smaller than Figures 5.6 & 5.8, with a model value of ~ 2.3 compared to an observed value of ~ 2 .

5.3 Discussion

The modelling described above is able to represent many of the features observed by LOFAR in this observation, despite being an extremely simplified representation of the actual ionosphere. Whereas the actual ionosphere varies in three spatial dimensions and in time, the phase screen model is able to replicate these observations despite its simpli-

fications. This modelling allows approximate TID parameters such as wavelength and TEC perturbation amplitude to be estimated, and indicates that LOFAR is picking up signatures of TIDs too small to be resolved by most common techniques of ionospheric monitoring.

The extension to numerical modelling makes it possible to build upon the initial model of Meyer-Vernet (1980), and suggests that LOFAR can infer the presence of non-linearities in the TID perturbation. This means that as well as identifying TIDs with shorter wavelengths than those observed by other techniques, these observations can provide information on the fine structure of such TIDs. The numerical modelling approach also enables non wave-like perturbations to be considered, meaning that other seemingly unrelated features in dynamic spectra can potentially be replicated with this approach (see e.g. Trigg et al., 2024; Wood et al., 2024). The main limitation on this is that the model, as currently defined, could not replicate a situation where multiple structures are present and propagate in different directions or with different velocities, as was observed by Fallows et al. (2020).

The observations from the stations in the Netherlands provide confirmation that these small-scale disturbances are propagating horizontally, supporting the arguments of both Ivanova et al. (2011) and Lan et al. (2018) that their observations of short period ionospheric variations at a single point were manifestations of SSTIDs. Although only one feature was observed from the Netherlands, and hence no period is directly observable in this case, the phase screen modelling indicates that it is caused by a structure with a length scale consistent with being an SSTID.

The introduction of non-linearities to the modelling also builds upon the explanation provided by Koval et al. (2018) for their ‘X-shaped’ caustics. They observed some intensity enhancement above the focal frequency when modelling focusing from a perfectly sinusoidal TID, and proposed that this focusing could therefore explain both

their ‘inverted V like’ and ‘X like’ caustics. However, their model does not replicate the two sharp intensity enhancements but rather a more diffuse intensity enhancement above the focal frequency, and also does not extend as far above the focal frequency as the observed ‘X like’ features do. If the ‘X like’ caustics are arising from non-linearities, such as the harmonic described in Section 5.2.3, this also explains why they occur in some cases and not others. However, there may be other non-linear perturbations that would lead to similar features in dynamic spectra, meaning that their association with the presence of a harmonic in the TID perturbation is not guaranteed.

The effectiveness of the phase screen model, despite clearly being an oversimplification, suggests that the TID perturbation may be highly localised in altitude. This would be consistent with the observations of Lan et al. (2018) that short period perturbations are restricted to lower altitudes, up to around the observed F2 peak in this case. If the AGW driving the observed TID can only just reach the altitude of the F2 peak the resulting plasma density perturbations would be expected to be extremely localised in altitude (i.e. not extending significantly into the topside). The low altitude of the F2 peak (~ 200 km) may also explain why these features were observed in this case but are not seen more regularly. If the F2 peak was at higher altitude, then any TID driven variations in electron density and hence TEC would be dominated by AGWs capable of reaching this higher altitude, thereby reducing the significance of the shorter wavelength waves observed here.

The phase screen model also provides an estimate of the TEC perturbations required to explain the observed features. In the cases considered here, the TEC perturbation amplitudes are $\sim 0.04 - 0.13$ TECu ($1 \text{ TECu} = 10^{16} \text{e}^- \text{m}^{-2}$) which is close to the theoretical lower limit of the sensitivity of GNSS TEC measurements (Otsuka et al., 2013). The inferred harmonics are a factor of ~ 10 lower in amplitude, suggesting that LOFAR ionospheric scintillation observations are capable of detecting the effect of TEC

variations on the order of mTECu, as has also been suggested by Trigg et al. (2024).

However, there are still significant discrepancies between the observed and modelled dynamic spectra, most notably the inadequacy of the horizontal TEC gradient to explain the asymmetry. Although the parameter α does not provide a physically reasonable explanation for the observed time asymmetry, it is possible to closely replicate the observations with an arbitrary value. This suggests that although the physical mechanism is not the large-scale horizontal gradients in the ionosphere identified from MADRIGAL GNSS TEC data, whatever mechanism is responsible causes similar frequency dependence. The asymmetries arising from the low elevation, and from the alignment between the line of sight and TID phase fronts demonstrated by Koval et al. (2018), are both candidates. However, this was not explored here as neither effect can be quantified in a way that allows it to be directly added to the simple single phase screen model.

As well as the discrepancies between observed and modelled spectra, there are two features in the UK observation of Cygnus A that the model does not replicate. Features 2 and 4 are distinct from the others largely due to the nature of their time asymmetry. Whereas the others remain approximately symmetric around a curve described by some value of α , these features show almost no dispersion in the higher frequencies and strong dispersion in the lower frequencies. This suggests that any perturbation capable of explaining these features must either be spatially asymmetric or change significantly in time as it passes through the line of sight. Therefore, in order to attempt to replicate these features, the model perturbation would need to be extended beyond what has been considered here, to allow for strongly asymmetric perturbations.

5.4 Conclusions

Features in dynamic spectra observed by LOFAR on the 15th of September 2018 have been shown to be caused by the passage of an SSTID across the line of sight. These features are characterised by a broadening towards the lower frequencies overlaid with interference fringes and a focal frequency typically around 40 – 60 MHz. Phase screen modelling is able to replicate the observed intensity distribution in time and frequency, and allows the wavelength and amplitude of the SSTID to be estimated given its altitude. The structures that explain the observed features are too small to show up in GNSS TEC maps (with typical pixel resolutions $\sim 0.15^\circ$ in Europe but data smoothed over ~ 5 pixels, e.g. Otsuka et al., 2013; Koval et al., 2019), but TEC time series from individual GNSS receivers do show waves with the expected amplitude and period present at the time. Favourable conditions for observing the features require a compact source to resolve the fringes and a low elevation to allow for a long focal length.

The simplicity of the model required to replicate the observations is striking because it implies that the ionospheric variations were dominated by a single structure with a well defined altitude and wavelength. There is no evidence of scattering from a spectrum of irregularities as has been reported by Fallows et al. (2020), suggesting that either conditions were not favourable for significant growth of irregularities or that such irregularities had not had time to develop. Further work using the numerical phase screen approach has shown that it is a powerful tool for explaining a range of features in LOFAR observations (e.g. Trigg et al., 2024; Wood et al., 2024).

These observations show that LOFAR is capable of resolving ionospheric structure on a scale other approaches tend to miss. GNSS TEC maps and ionosondes are typically too coarse and GNSS scintillations arise from much smaller structures (Basu et al., 1998). This suggests LOFAR observations are able to complement others to provide a more complete description of TIDs across Europe, both by observing shorter

wavelengths and by revealing the presence of non-linearities. The absolute sensitivity to perturbations of $\sim \text{mTEC}_u$ implied by the phase screen modelling is also a unique feature of LOFAR observations which is not possible with GNSS TEC measurements regardless of the density of the receiver network. However, a more systematic study would be required to determine whether these features are a regular occurrence or if the observation considered here happened under specific favourable conditions that are rarely repeated.

Chapter 6

A Method for Identifying Wave Signatures in Interferometric Calibration Solutions

This chapter describes work on identifying ionospheric waves using the differential TEC (dTEC) values derived as part of the calibration process for interferometric observations using LOFAR, which has been published in Radio Science (Boyde et al., 2024). These values are obtained at a precision of < 1 mTECu (Mevius et al., 2016) and this, coupled with the dense network of stations in the Netherlands, makes them potentially an extremely useful data source for studying ionospheric disturbances. The fact that the values are only obtainable as dTEC between two antennas rather than absolute TEC at a single location introduces some unique complications to any analysis.

6.1 Differential TEC and Absolute TEC

The calibration process used to derive the dTEC values is explained in Appendix A. The vital points for the following discussion are that the dTEC values are derived for the line of sight to a bright ‘calibrator’ source meaning that variations in ionospheric conditions across LOFAR’s field of view are negligible, and that the values are obtained at an extremely high precision (< 1 mTECu: Mevius et al., 2016) and with a typical time resolution of ~ 4 s.

The key difference between GNSS derived absolute TEC and LOFAR derived differential TEC is that the LOFAR values are not strictly local to the station for which they are derived. They include a combination of variation at that station and variation at the chosen reference station. This distorts the measured amplitude and timing of fluctuations compared to the actual structures passing across the line of sight.

In order to illustrate the difficulties this can cause, a simplified model of a 1-dimensional ionosphere with a single structure propagating through with constant velocity is useful. The expected TEC/dTEC values measured using GNSS receivers and LOFAR stations at identical positions can then be calculated (ignoring limitations on time sampling and measurement accuracy).

As an example, it is useful to consider wave-like perturbations such as would be expected from a TID. The expected observed values for a sinusoidal perturbation are shown in Figure 6.1. The GNSS TEC time series for the two measurement locations are identical except for a time shift corresponding to the propagation time from one station to the other, while the LOFAR differential TEC shows clear differences. The phase difference between the two stations does not correspond to the propagation time, and the amplitudes are both distorted, with the shorter baseline giving a slightly reduced amplitude and the longer baseline giving an amplitude almost twice the actual perturbation.

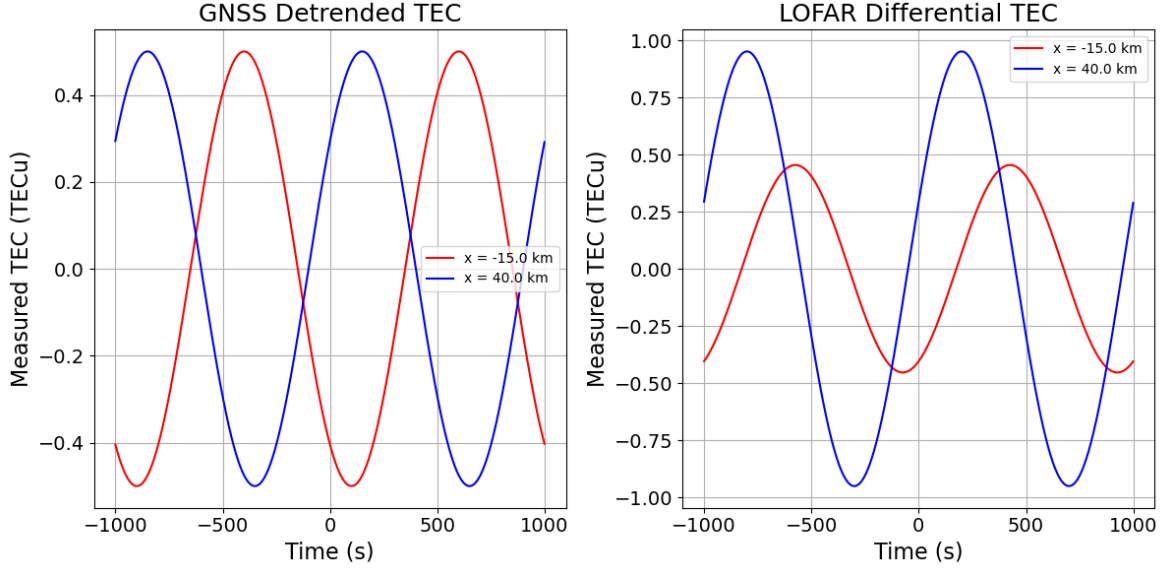


Figure 6.1. The modelled TEC time series for GNSS receivers (left panel) and LOFAR calibration solutions (right panel) for a sine wave perturbation. The reference station for LOFAR is assumed to lie at $x = 0$. The blue curve represents a receiver at $x = 40$ km and the red curve represents a receiver at $x = -15$ km. Note the difference in vertical axes. The sine wave has amplitude 0.5 TECu, wavelength 100 km and velocity 100 m s^{-1} .

For a sine wave perturbation, it is possible to analytically calculate the expected amplitude and phase observed in LOFAR calibration solutions. This is because the dTEC at a position \vec{r} and time t is given by

$$\begin{aligned} \text{dTEC}(\vec{r}, t) &= \Delta\text{TEC} \left(\cos(\vec{k} \cdot \vec{r} - \omega t) - \cos(-\omega t) \right) \\ &= A \left(\Delta\text{TEC}, \vec{k}, \vec{r} \right) \cos \left(\phi(\vec{k}, \vec{r}) - \omega t \right), \end{aligned} \quad (6.1)$$

where ΔTEC is the amplitude, \vec{k} is the wavevector, \vec{r} is the station-reference separation vector and ω is the angular frequency. The amplitude $A(\Delta\text{TEC}, \vec{k}, \vec{r})$ and phase $\phi(\vec{k}, \vec{r})$ of the observed wave are given by

$$A(\Delta\text{TEC}, \vec{k}, \vec{r}) = \Delta\text{TEC} \sqrt{2 - 2 \cos(\vec{k} \cdot \vec{r})}, \quad (6.2)$$

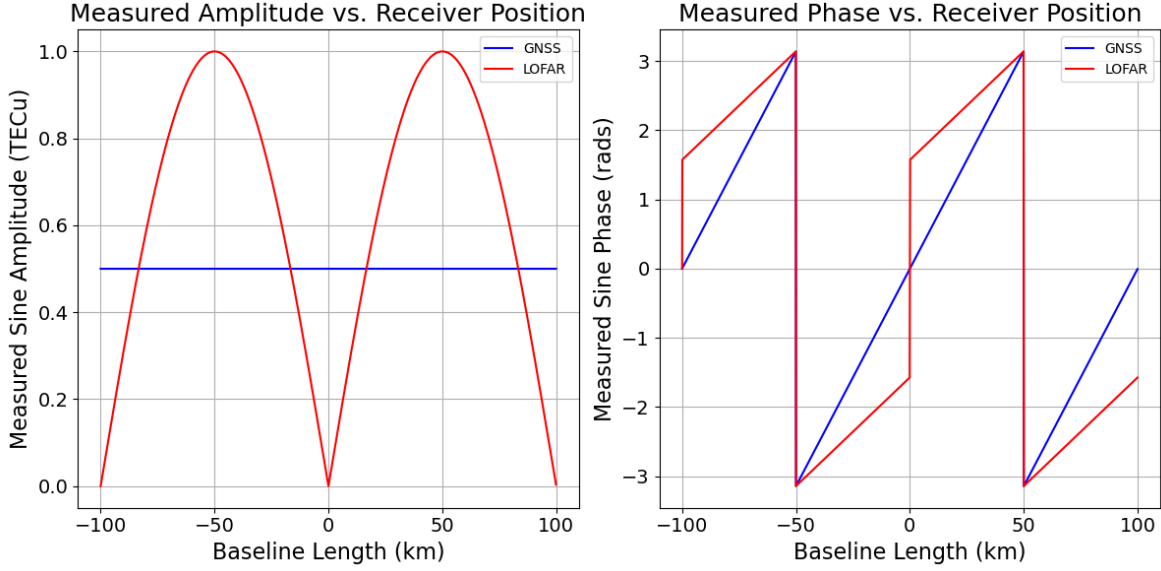


Figure 6.2. The measured amplitude (left) and phase (right) of the sine perturbation as a function of baseline for GNSS receivers (baseline independent) and LOFAR calibration solutions. The sine wave considered is the same as in Figure 6.1.

and

$$\begin{aligned}
 \phi(\vec{k}, \vec{r}) &= \arctan\left(\frac{\sin(\vec{k} \cdot \vec{r})}{1 - \cos(\vec{k} \cdot \vec{r})}\right) + \vec{k} \cdot \vec{r} = \arctan\left(\cot\left(\frac{\vec{k} \cdot \vec{r}}{2}\right)\right) + \vec{k} \cdot \vec{r} \\
 &= \arctan\left(-\tan\left(\frac{\vec{k} \cdot \vec{r} \pm \pi}{2}\right)\right) + \vec{k} \cdot \vec{r} = \frac{1}{2}(\vec{k} \cdot \vec{r} \bmod 2\pi).
 \end{aligned} \tag{6.3}$$

In practice, given that the ‘true’ $t = 0$ is unknown, the ωt term in equation (6.1) will be an additional unknown in any fitting using these equations.

The dependence of amplitude and phase on baseline length is illustrated in Figure 6.2, based on equations (6.2) and (6.3). Also shown are the equivalent values for a GNSS receiver observing the same wave. The right panel of Figure 6.2 shows the clear phase discontinuity across zero baseline length and at integer multiples of the wavelength (red line, the apparent discontinuities at $x = \pm 50$ km are simply a result of the phase wrapping). By estimating amplitude and phase for the observed wave-

like variations in dTEC, the wave parameters can therefore be estimated even if the wavelength is significantly larger than the array of stations (up to ~ 500 km, see Section 6.4). The phase discontinuity also provides a clear indicator of propagation direction in the 2D case of the actual LOFAR array, as the orientation of this discontinuity will be perpendicular to the propagation direction for an idealised wave.

In practice, while the observed wavelength should be accurate, the observed amplitude will be sensitive to the observing geometry. If the lines of sight from LOFAR to the calibrator source lie in planes of constant phase of the TID, then the observed amplitude will be maximised, whereas in less favourable geometries it will be suppressed. Another issue could arise due to the difference in orientation of the lines of sight from two different antenna in their local frames, due to the curvature of the Earth. This effect should be negligible provided only stations within the Netherlands are considered, as the differences in azimuth and elevation to a given source between these stations are significantly less than 1° .

6.2 Sample Data

To illustrate the analysis methods described in this chapter, the calibration solution for an 8-hour observation (ID L691726) made using the HBA between 120.2 – 187.4 MHz with a frequency resolution of ~ 48.8 kHz is used. This observation ran from 15:15 UT to 23:15 UT on 21 December 2018, and targeted the source 3C48 (RA 1h37m41s, Dec. 33.16°). As 3C48 is suitable for use as a calibrator, dTEC solutions were able to be derived directly for the target source in this case. While the international stations were included in this observation, in this chapter only the stations in the Netherlands where structures are more likely to span all lines of sight are considered. Over the course of the observation, the apparent position of 3C48 moved from due east ($\sim 90^\circ$ azimuth)

to due west ($\sim 270^\circ$ azimuth) as viewed from the LOFAR core, starting and ending at an elevation of $\sim 44^\circ$ and reaching a maximum elevation of 70° .

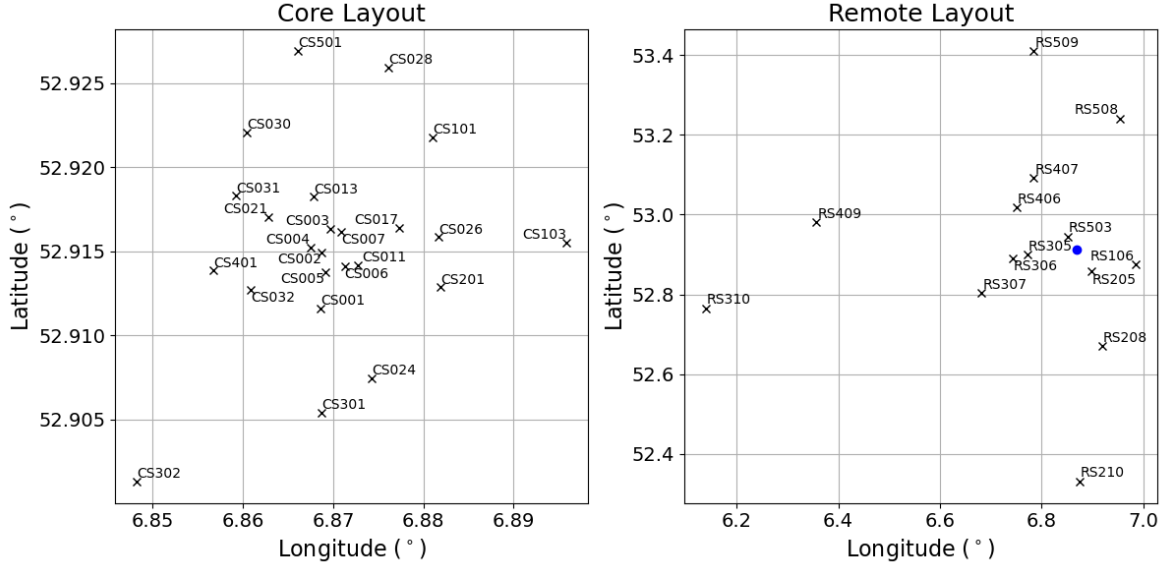


Figure 6.3. The layout of core (left) and remote (right) stations in the Netherlands. The aspect ratio is not set to match physical distances meaning distances and directions are slightly distorted. The blue dot in the right hand panel shows the location of CS001.

The layout of the LOFAR stations in the Netherlands, shown in Figure 6.3, is significant in determining its ability to resolve ionospheric structures. This is largely due to the layout having clear preferred directions, especially in the remote stations. There are long baselines from the core to the north, west and south but significant gaps in between, rather than the pseudo-uniform coverage that would be more suitable for ionospheric observation. The exact distribution of baselines corresponding to the dTEC values is not fixed however, due to the possibility of choosing a different reference station. It is likely that different distributions of baselines will be more or less sensitive to disturbances with different orientations, and comparing results with different choices of reference is shown in Section 6.4 to be beneficial.

The dTEC solutions for the full observing window for a range of stations are shown in Figure 6.4. These show activity across timescales ranging from minutes to hours.




Figure 6.4. A selection of dTEC time series for the whole of observation L691726 on 21 December 2018, referenced to CS001HBA0. The top panel shows the dTEC on various baselines within the core, while the bottom panel shows dTEC on baselines to remote stations with one core station baseline for comparison. Note the difference in magnitude of the vertical scales in each panel. The remotes shown are the furthest from the core along their respective directions, as shown in Figure 6.3.

The baselines within the core (top panel) are generally all very closely correlated with one another. Some remote station baselines (bottom panel) show appreciable correlation, such as those to RS310 and RS409, due to their similar orientation and length. However, there are also several features in the remote station baselines that are unique to a single baseline, such as the brief, sharp increase in dTEC on the baseline to RS210 just after 17:00 UT. Also evident is the large scale structure of the ionosphere, with certain stations retaining a consistent offset from zero. For example RS210 (south of the

core) observes a consistently positive dTEC whereas RS509 (north of the core) observes consistently negative dTEC. This is expected given the typical increase in TEC in the mid-latitudes with decreasing latitude. It is important to note that the magnitudes of the dTEC values here, even on the longest baselines to the remote stations, are significantly smaller than the magnitude of variations that have been published previously (see e.g. de Gasperin et al., 2018; Mevius et al., 2016). This is likely related to the present observation being made at solar minimum during winter solstice and largely at night, meaning ionospheric densities and therefore TEC and TEC variations are particularly small. As a result, these data should provide a good test case to demonstrate the sensitivity of these measurements to extremely small variations in TEC.

For plotting purposes, in order to isolate local ionospheric variability from the global structure and its diurnal variations it is useful to detrend the time series to remove long timescale variations. This is commonly done with GNSS TEC (e.g. Coster et al., 2017; Koval et al., 2019; Otsuka et al., 2013) by calculating the mean in a sliding window (typical width $\sim 30 - 60$ min) and subtracting this from the data. This means that the initial and final sections of the time series are lost but in this case with continuous 8 hour time series this is acceptable. This detrending can also distort the amplitude of certain signals, particularly those with periods close to the window size (Auchère et al., 2016). Therefore it is only used here for illustrative purposes and not for any of the subsequent analysis, with a detrending window of 60 min.

Within the 8-hour long observation, this chapter focuses on the period between approximately 17:20-17:50 UT during which there is clear wave activity across the array. A visualisation of a some of the detrended data during this time period is shown in Figure 6.5. The map shows the ionospheric pierce point (IPP) of each station colour coded based on the detrended dTEC value. The IPP is defined as the intersection between the line of sight and a hypothetical ionospheric ‘shell’ at some altitude, in this




Figure 6.5. The detrended TEC values for all the remote stations plus CS001HBA0, detrended relative to RS306 (black cross). Values are shown for four different time stamps. each two minutes apart. Points are mapped to the IPP location assuming a 350 km shell height. Again the aspect ratio is not correct causing the longitudinal distances to be exaggerated relative to the latitudinal distances.

case 350 km. The four time stamps shown correspond to a period in which wave-like structures propagate roughly west to east across the array. For this plot the reference station has been set to RS306 as this is closer to the midpoint of the array in the longitudinal direction, making it easier to discern the longitudinally propagating structure. The enhancement is beginning to appear in the westernmost stations in the first panel, passes across the reference at roughly the third panel (bottom left) and is clearly visible east of the reference in the final panel. The wave propagation is clearer when all timestamps are combined to form a video of the dTEC variations (see Supplementary Material of Boyde et al., 2024).

6.3 A Method for Characterising Waves using Wavelet Analysis

The data shown in the previous section illustrates that waves can be observed in the dTEC solutions, but the underlying wave parameters are not immediately apparent from the raw dTEC values. It is clear that the wavelength of the wave is larger than the longest baselines in the array as multiple wavefronts are not seen simultaneously, but the actual wavelength and amplitude of the wave are not clear. The simple plane wave model shown in Section 6.1 provides a way to estimate these parameters, as it predicts a dependence of the observed amplitude and phase on the baseline and wave parameters.

The observed amplitude and phase can be estimated using a wavelet transform. This takes the observed time series and convolves it with a ‘wavelet’ of varying time scale and time of peak amplitude. This transforms the 1D time series at each station into a complex 2D wavelet spectrum $W(T, t_0)$ as a function of period T and time t_0 . The wavelet used in this case is the Morlet wavelet, given by

$$\psi(t, s, t_0) = \sqrt{\frac{1}{s\sqrt{\pi}}} \exp\left(\frac{-(t - t_0)^2}{2s^2}\right) \exp\left(\frac{-5i(t - t_0)}{s}\right), \quad (6.4)$$

where t is time and s is the width parameter ($s(T) = \frac{5T}{2\pi}$). This represents an oscillation suppressed by a Gaussian such that only approximately one period of the oscillation is non-negligible. The wavelet coefficients calculated from the data are a numerical approximation to the continuous wavelet transform given by

$$W(T, t_0) = \int \text{dTEC}(t) \psi(t, s(T), t_0) dt. \quad (6.5)$$

This definition of the Morlet wavelet is the standard definition used by Torrence

and Compo (1998) and widely implemented in software packages. However, as shown by Liu et al. (2007), it is biased towards larger periods and the amplitude of the wavelet coefficient does not correspond to the amplitude of the wave. This can be corrected by dividing the wavelet coefficients by a factor $\sqrt{s f_s}$ where f_s is the sample frequency.




Figure 6.6. The wavelet power spectrum for RS310 with RS306 as the reference station. Black contours denote the 95% local significance level based on the estimated noise spectrum (see Section 6.3.1 for details). Red dashed lines indicate the time range over which the noise was estimated, meaning that significance levels outside this time range may be unreliable. Shaded regions at either side are the cone of influence, in which the values are unreliable due to edge effects.

An example wavelet power spectrum for station RS310 using RS306 as a reference is shown in Figure 6.6. The wave activity identified in Figure 6.5 corresponds to the enhancement of wavelet power at around 17:30 UT at a period of $\sim 8 - 12$ min. It should be noted that the contours for the confidence levels over most of the spectrum are unreliable as the noise spectrum was only estimated over the 1 hour range shown by

the red dashed lines (the method for estimating the noise spectrum will be explained in the following section). The noise as defined here includes turbulent variations of the ionosphere which vary with time, as well as instrumental noise contributions which vary with changing source elevation (e.g. Mevius et al., 2016). This means that it is likely that some of the areas identified as ‘significant’ outside of the region of noise estimation are spurious.




Figure 6.7. The wavelet power (left) and phase (right) calculated for a period of roughly 9.2 mins at 17:36:13 UT. The black x shows the location of the reference station (RS306). The sign of the phase discontinuity is arbitrary due to the phase wrapping and an unknown offset. Aspect ratio is not set to match physical distances.

To test whether the simple plane wave model is appropriate, the geographical distribution of power ($|W|^2$) and phase ($\arg(W)$) at the time and approximate period of the observed waves in Figure 6.5 is shown in Figure 6.7. The period was selected as the one with the highest amplitude at any of the Dutch stations at this t_0 . The east-west variability in both parameters is apparent, as is the discontinuity of phase and the near symmetric variation of power. This indicates that the plane wave model in this case provides a good approximation to the observed wave behaviour and may be useful to

estimate the wave parameters.

6.3.1 Noise Estimation

As well as estimating the wave parameters, it is useful to quantify the uncertainty in the values obtained from the wavelet transform. Doing this requires an estimate of the uncertainties in the wavelet coefficients (specifically in the wavelet power and phase). This ‘noise’ will include contributions from both instrumental/model errors in the measured values and real turbulent variations in the ionosphere, all of which can vary with time and/or observing geometry. To estimate the uncertainty in the wavelet parameters, it is necessary to estimate this noise power as a function of period and then compare it to the wavelet power.

As well as providing a basis for estimating uncertainty in wave power and wave phase, the noise spectrum gives a basis for specifying confidence levels. These are based on the fact that the noise power will be chi-squared distributed with two degrees of freedom (corresponding to the real and imaginary components of the wavelet coefficient) around the underlying noise spectrum. Confidence levels can be used to determine which baselines to consider in any fitting, excluding any which are likely to be noise dominated. The definition of these confidence levels is explained in detail in Auchère et al. (2016).

A common noise model across many areas of geophysics is the AR1 model (e.g. Torrence and Compo, 1998), which is also implemented into many common wavelet packages. This assumes that the noise can be described by a Markov process, where each value is only influenced by the preceding value. In other words, the noise contribution σ_i at timestamp i can be expressed as

$$\sigma_i = \alpha\sigma_{i-1} + z_i, \tag{6.6}$$

where α is the lag-one autocorrelation describing how ‘red’ (i.e. how biased towards lower frequencies) the noise is, and z_i is a sample from a Gaussian with mean zero and variance proportional to noise power. The advantage of this is that its parameters can be very easily estimated directly from the time series, but it can only describe a limited range of spectral shapes. The dTEC data does not match the spectral shape of the AR1 model, and so attempting to apply it would give extremely misleading results.

A more general approach to noise estimation is presented by Auchère et al. (2016). This involves taking the Fast Fourier Transform (FFT) to obtain the power spectrum of the data and fitting a parametrised spectrum to it. The fit is based on maximising the likelihood using the fact that the Fourier power spectrum of a pure noise time series will be distributed around the true underlying spectrum according to a chi-squared with two degrees of freedom. While this is an extremely statistically thorough approach, it still requires an assumed spectral shape for the noise, which has not been identified for the dTEC data. This was due to qualitatively different spectra being found for different baselines and different observations. For the models that were tested, the fitting also regularly found local rather than global maxima, and resolving this would have required very computationally expensive approaches.

The difficulty of identifying a suitable assumed noise spectral shape can be avoided by estimating the noise in a way that does not have an assumed functional form. For example, Robust Local Regression (RLR: Ruckstuhl et al., 2001, described in detail below) makes minimal assumptions about the spectral shape. All it requires is that the noise spectrum is locally linear. While this is not the case for the dTEC data, the spectrum does locally approximate a power law, meaning that it is linear in log-log space and so the fit can be carried out in this representation. Another advantage of RLR is that it directly accounts for the presence of signal by iteratively excluding spikes in the spectrum from the fitting, which is not achieved by the other methods described

previously.

To estimate the noise spectrum, RLR carries out weighted linear regression separately for each Fourier frequency. The weights are determined initially by the distance from each point to the point of interest in log-frequency space, with a functional form of

$$K_i = \max(0, 1 - |u_i|^3), \quad (6.7)$$

where u_i is the ratio of the separation in log-frequency space to the kernel width h . This kernel width is set for each point of interest to ensure 30 points with non-zero weight following the suggestions of Ruckstuhl et al. (2001). This provides the local component of RLR, but to ensure robustness this process must be iterated. On subsequent iterations, the local weights K_i are modified by multiplying by a robustness weight w_i given by

$$w_i = \begin{cases} \max\left(0, 1 - \left(\frac{\epsilon_i}{b\sigma}\right)^2\right), & \text{if } \epsilon_i > 0 \\ 1, & \text{otherwise} \end{cases} \quad (6.8)$$

where ϵ_i is the calculated (log) FFT power minus the fitted value from the previous iteration (i.e. the residual), b is the ‘robustness parameter’ set to 3 following Ruckstuhl et al. (2001), and σ is proportional to the median absolute residual. This then works to exclude sharp (i.e. narrower than h) peaks from the fitting procedure and thereby estimate the underlying noise without being biased upwards by the signal. In all test cases examined, the estimated spectrum was found to converge within 5 iterations, at which point the same weighting is used to interpolate onto the wavelet periods rather than the Fourier frequencies to give the estimated noise in the wavelet spectrum.

While fitting in log-log space is necessary to ensure local linearity, it also introduces a bias into the results. This is because the mean of the logarithm of a chi-squared distributed variable is lower than logarithm of its mean (Vaughan, 2005), but this

bias can be corrected by adding a constant to the fitted values before rescaling from log power. In practice, it is also necessary to first ‘prewhiten’ the time series before calculating the Fourier power spectrum. Otherwise, the higher frequencies will be dominated by the apparent discontinuity caused by assuming the signal periodically repeats rather than the true noise. After fitting, the spectrum can then be corrected for this by dividing by $2 - 2 \cos \left(\frac{2\pi f}{f_s} \right)$ (Percival and Walden, 1993). As well as the prewhitening, the signal is windowed using a Hamming window (Harris, 1978) in order to mitigate the spreading of signal power over multiple bins when the signal frequency does not match one of the Fourier frequencies. However, this also suppresses the noise power, creating a downward bias. To quantify this bias, the performance on artificial noise time series was tested, the results of which are described in Section 6.3.2.




Figure 6.8. An example of the noise fitting procedure using an hour (17:06-18:06 UT) of the data shown in Figure 6.4 from RS310 with RS306 as the reference. The black solid line represents the Fourier power spectrum while the red dashed line is the noise spectrum estimated by RLR.

An example of the noise fit results is shown in Figure 6.8, for an hour of the data centred on the waves shown in Figure 6.5 which peak in amplitude at approximately 17:36 UT. The RLR method clearly generates a smoothed version of the Fourier power spectrum as expected. There is also a distinct peak in the Fourier spectrum above the noise fit between $1 \times 10^{-3} - 2 \times 10^{-3}$ Hz, corresponding to periods of around 8 – 16 min. This is the signature of the waves visually identified in Figure 6.5, and the RLR fit has successfully fitted under this signal as intended (although it is probably still somewhat biased above the true noise in this case, see Section 6.3.4).

6.3.2 Noise Fit Validation

While the noise estimate in Figure 6.8 appears reasonable, because the true noise spectrum is unknown it is useful to test the method for bias using synthetic noise. For this purpose, 10,000 1-hour time series were generated with noise following the AR1 model of equation (6.6). Although as stated previously the noise spectra in the real data do not resemble the AR1 model, given that the fitting method chosen is in principle effective regardless of the precise spectral shape of the noise, this should provide a good test of the effectiveness of the noise estimation. Each of these time series had a sinusoidal signal added with variable frequency and signal-to-noise ratio (SNR) to test the ability to accurately exclude any signal from the noise fitting. The parameter α was varied between 0 and 0.99 and showed no impact on the performance of the noise estimates, as expected given no assumption is made about the exact spectral shape. The estimated noise was compared to the true value P_n given by

$$P_n = \frac{\sigma^2}{1 - 2\alpha \cos\left(\frac{2\pi t_s}{T}\right) + \alpha^2}, \quad (6.9)$$

where σ is the standard deviation of the sample z_i in equation (6.6), t_s is the time between samples (~ 4 s) and T is the period (Torrence and Compo, 1998). This comparison was made at both the signal period, and 10 times this value to capture the performance of the noise estimation in the absence of signal.

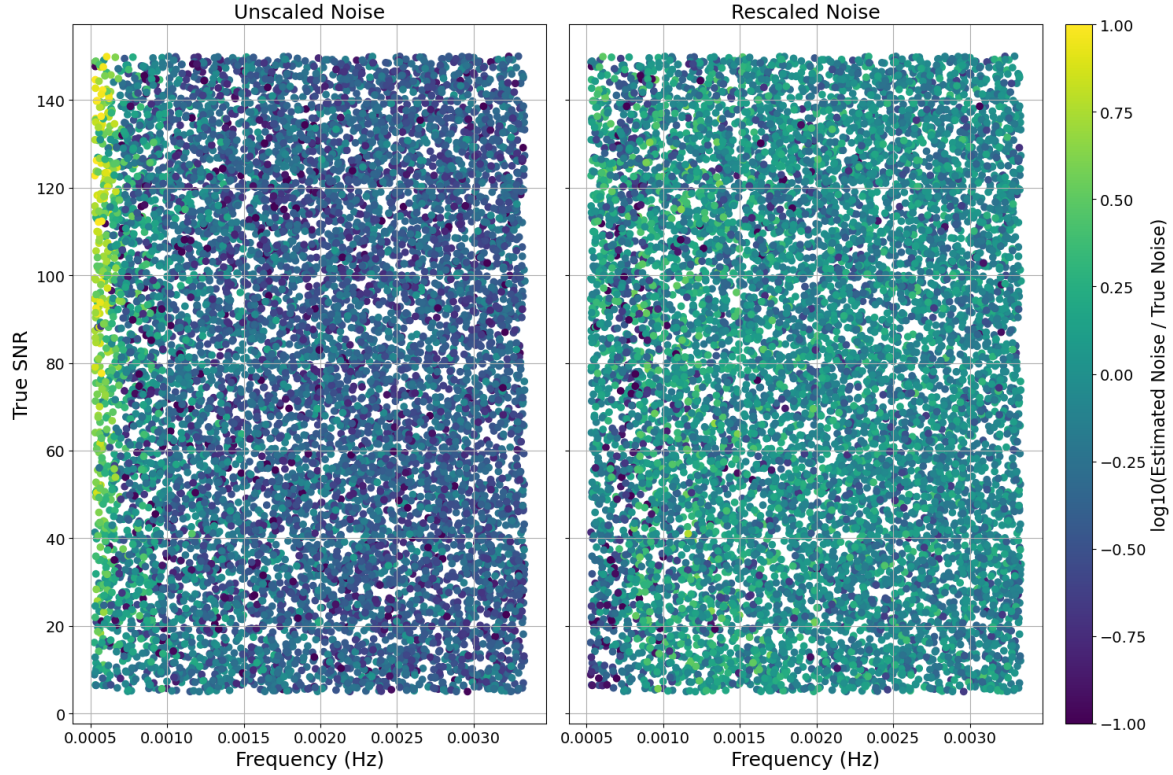


Figure 6.9. The errors in estimated noise power as a function of signal frequency and SNR before rescaling (left) and after rescaling by equation (6.10) (right). A logarithmic colour scale is used so that both overestimation and underestimation of the true noise are equally apparent.

The performance of the noise estimate at the signal frequency for a range of signal frequencies and SNRs is shown in the left panel of Figure 6.9. The performance is mostly uniform, with only the lowest signal frequencies showing any systematic variation. This is an unavoidable problem as for very low frequencies ($< \sim 0.8$ mHz in this case) relative to the frequency resolution of the FFT, there are insufficient Fourier bins at lower frequencies to quantify noise power, and so the increased power due to the presence of the signal cannot be easily distinguished from an increase in noise power. However,

for the vast majority of signals, this indicates a consistent performance as desired. To quantify the exact performance, the data shown in the left panel of Figure 6.9 were fitted by

$$\frac{P_{est}}{P_{true}} = C \left(B \exp \left(\frac{-T_0}{T} \right) + 1 \right), \quad (6.10)$$

where P_{est} is the estimated noise power, P_{true} is the true noise power given by equation (6.9) and C , B and T_0 are the fit parameters. These were found to be $C = 0.471 \pm 0.007$, $B = 570 \pm 50$ and $T_0 = 7700 \pm 150$ s. The SNR was not included in the fit as in practice this is unknown without an accurate noise estimate, and so cannot be used to obtain the accurate noise estimate in the first place.

The rescaled noise estimates are then shown in the right panel of Figure 6.9. Except for SNR dependence at the lowest frequencies, there is no remaining trend, showing that the RLR noise estimation method with the rescaling described by equation (6.10) is a generally consistent and unbiased means of estimating the noise spectrum. Applying this same correction to the estimated noise at 10 times the signal frequency leaves some residual underestimation of the noise in these regions by about 20-25%. This is expected as a result of the one-sided robustness weighting function given by equation (6.8), which inevitably underestimates noise in the absence of a signal (Ruckstuhl et al., 2001). While this is not ideal, accurate estimation of the noise in the presence of a signal is prioritised over accuracy in the absence of signal as there is no clear way to achieve both simultaneously.

6.3.3 Noise Correlation

Now that the noise power can be estimated, there is a second consideration in these data: the fact that the ‘noise’ includes real turbulent variations in the ionosphere. This means that the noise on nearby baselines such as those in the LOFAR core will have

appreciable correlation. Accounting for this correlation is vital in order to accurately identify both the optimal fit parameters and their uncertainty. However, measuring the spatial correlation is difficult as it requires the noise component of the signal to be isolated from genuine coherent signals, otherwise the correlation will likely be overestimated.

The correlation might also be expected to vary depending on the noise timescale considered. This is because larger timescales will generally correspond to structures with larger scale sizes which would be expected to correlate over a greater distance. Hence, it is sensible to estimate the correlation structure independently for each timescale to capture any variation resulting from this. However, larger timescales introduce the problem of the expanding ‘cone of influence’ - the region of time which the wavelet coefficient at a given point in time is sensitive to. This reduces the number of independent samples that can be used to generate the correlation, meaning that in practice only relatively short timescales (up to ~ 5 min for a 1 hour dataset) can be analysed in this way. The details of this analysis are described in Appendix B.

The correlation length r_c and scale factor A estimated based on equation (B.1) are shown in Figure 6.10. Figure 6.10 shows a consistent correlation length for periods over 1 minute. Below this, the instrumental (i.e. uncorrelated) noise represented by A starts to dominate and the correlation length values themselves are largely meaningless as the fit values are so low that they are essentially spurious. Averaging over the results for periods from 1 to 5 minutes gives a median value of 13.2 km and upper and lower quartiles of 9.2 km and 19.7 km respectively. Given that this value is apparently constant between 1 – 5 min, it is reasonable to extrapolate this value to higher periods as well in the absence of a direct means of measuring correlation for the longer periods.

Although the fits here suggest $A = 1$ as the optimal solution, in practice reducing this to $A = 0.95$ was necessary. This was done because using $A = 1$ was found to give




Figure 6.10. The median estimated correlation length (r_c , left) and correlation scale (A , right). The shaded region indicates the interquartile range. The approach used to deriving these quantities is described in Appendix B.

unreasonable solutions, such as predicted power exceeding all measured powers by more than 1σ . This is likely due to the waves not being described precisely by an ideal plane wave, which introduces another source of variation not captured by this correlation estimate and resulting estimated error covariance. The root mean square residual on the fit to estimate the correlation length was found to increase with increasing period, but this likely reflects the reduction in sample size at longer periods due to the increase in width of the cone of influence mentioned previously rather than a weakening of the underlying relationship. For all further analysis, r_c was taken to be the median value of 13.2 km.

6.3.4 Fit Method

Using the wavelet coefficients, estimated noise power and an estimate of the noise correlation, a covariance matrix corresponding to the data can be defined and used to define the best fit parameters. These are the parameters which minimise

$$\chi^2 = \vec{\epsilon} \cdot (\underline{\underline{C}}^{-1} \cdot \vec{\epsilon}), \quad (6.11)$$

where $\vec{\epsilon}$ is the vector of phase and power residuals after fitting to equations (6.2) and (6.3) and $\underline{\underline{C}}^{-1}$ is the inverse of their combined covariance matrix. The residual vector $\vec{\epsilon}$ is determined by the four free parameters of the fit: the wavevector magnitude $|\vec{k}|$, the wavevector azimuth θ , the phase offset ωt and the wave amplitude ΔTEC . The fit is carried out for both phase and power simultaneously as both relationships are dependent on wavevector, meaning this is necessary to obtain a fully accurate solution.

To ensure the fitting algorithm reliably finds the optimum solution, it must be provided with a reasonable initial estimate of the parameters. This relies on the assumption that the wavelength is significantly greater than the baseline lengths, so that there is only one phase discontinuity and the power variation is approximately quadratic. The first parameter to be estimated is the azimuth, using the wavelet powers and following the approach of Beser et al. (2022). The quantity $\frac{|W|^2}{|r|^2} \hat{r}$ is used, which is expected to depend only on baseline orientation relative to the wave propagation direction. Then Principal Component Analysis (PCA) is used to identify the dominant direction in this quantity, which ideally corresponds to the wave direction. Due to the uneven distribution of baseline directions provided by LOFAR (see Figure 6.3), this initial estimate has a bias which needs to be corrected. This is achieved by generating idealised data for waves propagating at a range of azimuths and using the PCA estimated azimuths to define a means of inverting to recover the true azimuth, described in Appendix C.




Figure 6.11. The phase (left) and power (right) data as a function of projected baseline length with the fitted values shown. The fit parameters were: $|\vec{k}| = 4.08 \pm 0.94 \times 10^{-5} \text{ m}^{-1}$, $\theta = 84.4 \pm 2.8^\circ$, and $\Delta\text{TEC} = 8.4 \pm 1.9 \text{ mTECu}$.

With an estimated azimuth, the baselines are projected onto this direction to estimate the other parameters. An example of this projection for the sample data from Figure 6.7 is shown in Figure 6.11, although in this case the data is projected onto the final fitted azimuth rather than the initial PCA estimate. The phases are sorted by projected baseline and then unwrapped, then the sign of the discontinuity is determined and the discontinuity is removed. In Figure 6.11, this would mean adding 180° to all phases with $x > 0$. After removing the discontinuity, linear regression is used to provide an estimate for $|\vec{k}|$ and ωt . With the estimate of $|\vec{k}|$, ΔTEC is estimated by averaging $\frac{|W|^2}{(kx)^2}$ on the longest positive and negative projected baselines, which is the long wavelength approximation to equation (6.2).

While in most cases the method as described above is effective, it can encounter problems in some cases due to the discontinuity in the phase relationship. This can cause very sharp changes in χ^2 with θ where the sign of one projected baseline x changes and hence moves to the other side of the phase discontinuity, which can be local minima

and hence confuse the fit. To mitigate this issue, after the fit, any baselines which are within $\max(0.1^\circ, \sigma_\theta)$ (σ_θ is the estimated azimuth uncertainty) of being perpendicular to \vec{k} are removed and the fit is repeated, using the previous ‘optimal’ solution as its initial guess. The floor of 0.1° is necessary as in these cases the sharp variation in χ^2 with θ can lead to estimated σ_θ values of $< 0.001^\circ$ which are obviously unreasonable. Once no baselines lie within $\max(0.1^\circ, \sigma_\theta)$ of perpendicularity to \vec{k} , the fit is accepted.

The result of this method of fitting projected onto the optimal azimuth are shown in Figure 6.11 for the sample data from Figure 6.7. The agreement between the simple plane wave model and the data is very clear, although the estimated uncertainties do significantly overestimate the spread of the data around the fit in this case. The fit suggests a wavelength of 154 ± 35 km which is towards the lower end of what is typically classified as an MSTID. The estimated amplitude of 8.4 ± 1.9 mTECu is around an order of magnitude below the level which is typically required for confident identification of TIDs in GNSS TEC maps (e.g. Otsuka et al. (2013) used a threshold of 0.2 TECu), demonstrating the value of LOFAR dTEC measurements’ low noise floor (Mevius et al., 2016).

6.4 Validation Using Synthetic Data

To ensure that the method is accurately able to recover the true wave parameters, synthetic data including noise was generated. By applying the fit method to this data, it was possible to compare the known true wave parameters to those estimated by the fit and characterise the performance of both the estimated parameters and their estimated uncertainties. The method used to generate the synthetic data is described in Appendix D. An important caveat to note for this analysis however is that the performance on synthetic data will inevitably be better than for real data as the approximation of a

plane wave is precisely true for the synthetic data.

To improve the fit performance, it proved to be necessary to attempt the fit using several different references to get good performance across all azimuths. A set of three references were found to be sufficient for this: RS205, RS306, RS406. These provide a range of orientations relative to the core, while also ensuring that there are a reasonable number of baselines across a wide range of directions (i.e. they are not on the edge of the array). This is important because if the selected reference is on the edge of the array, the phase discontinuity will not be discernible. The selected fit parameters were then those from the reference which gave the smallest value of χ^2 in equation (6.11).

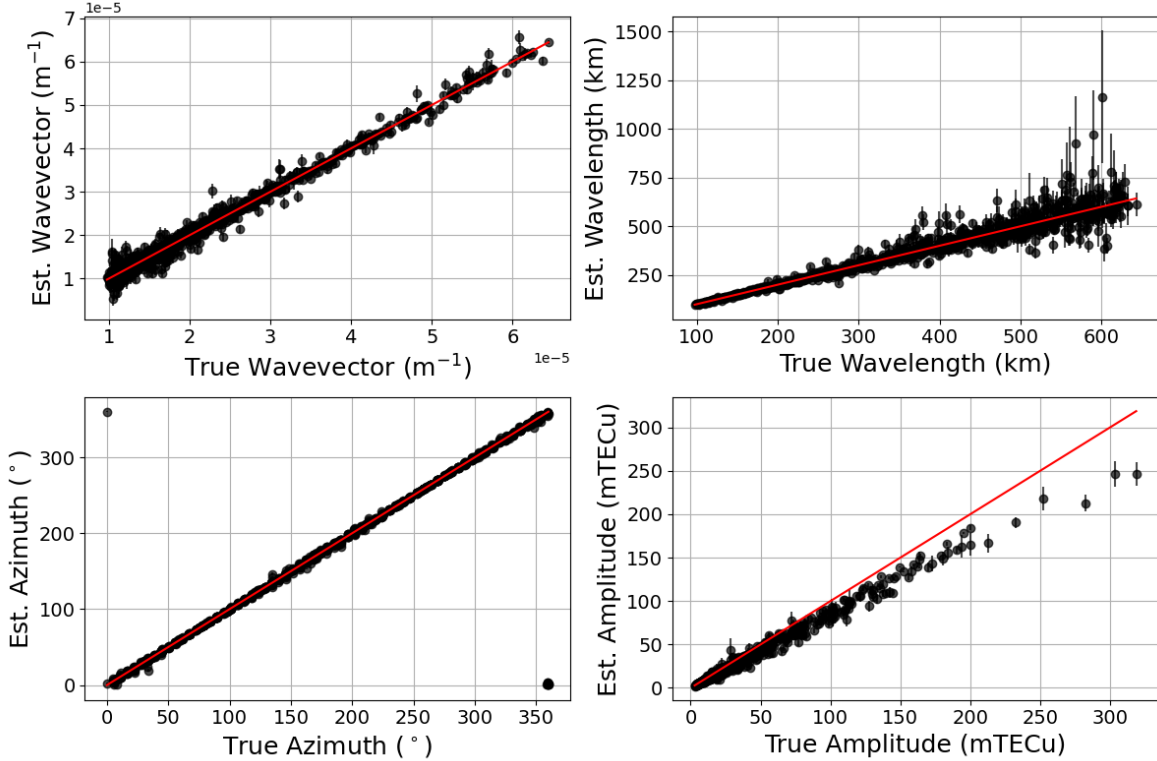


Figure 6.12. The relationship between true (horizontal axis) and estimated (vertical axis) parameters. Top left is wavevector $|\vec{k}|$, top right is wavelength Λ , bottom left is azimuth θ and bottom right is amplitude ΔTEC . Red lines represent the ideal (i.e. $x = y$) behaviour.

The correlation between true and estimated values of the TID parameters are shown in Figure 6.12. These plots exclude any fit for which the estimated uncertainty in $|\vec{k}|$ or

Parameter	RMSE	Bias	Correlation
$ \vec{k} $	$1.33 \times 10^{-6} \text{ m}^{-1}$	-0.03%	0.994
Λ	45.8 km	0.03%	0.959
ΔTEC	8.56 mTECu	-12.86%	0.994
θ	1.697°	-0.065°	1.000

Table 6.1. The performance of the fit in predicting the various wave parameters. RMSE is root mean square error. Bias is median percentage error except for θ where it is median error.

ΔTEC exceeds the fit value (i.e. relative uncertainty $> 100\%$), and those for which the estimated wavelength exceeded 5000 km (to ensure clarity of the plots by restricting the necessary range of the vertical axis in the top-right panel), which collectively accounted for only 0.4% of fits. Statistical measures of the fit performance are also given in Table 6.1, indicating the high correlation of all parameters. The RMSE in wavelength Λ is high, but this is largely dominated by the decrease in performance at long wavelengths shown in Figure 6.12. At shorter wavelengths, the errors will be much smaller than ~ 50 km (e.g. considering only $\Lambda < 250$ km gives an RMSE of 8.2 km). The other clear issue is a systematic underestimation of amplitude ΔTEC . This is a major contributor to the RMSE on ΔTEC , and, if the data is rescaled to remove the bias, the RMSE reduces to 4.25 mTECu. The azimuth is roughly uniformly accurate at all values, with only a slight decrease in accuracy for roughly North-South propagation ($\theta \sim 0^\circ$ or 180°) compared to East-West propagation visible in Figure 6.12. The reason for this is likely the shorter E-W baselines available compared to N-S baselines (see Figure 6.3), making the data less sensitive to changes in azimuth for N-S propagation.

The bias in amplitude is partially explained by a previously unnoticed bias in the wavelet transform itself. Applying the wavelet transform to a sine wave of unit amplitude, the amplitude of the wavelet coefficient at the corresponding frequency is roughly 0.93. This was not reported by Liu et al. (2007), likely as that work was focused on removing the bias in comparing amplitudes across scales rather than accurately deter-

mining the amplitude itself. This would give rise to a $\sim -7\%$ bias however, not the -12.86% found here. An extra contribution may be that at longer periods, even points in the wavelet spectrum which are outside the cone of influence (the region in which the effects of the finite length of the time series are generally defined as significant) still show noticeable decreases in amplitude beyond the $\sim -7\%$ general bias due to the proximity of the cone of influence in both directions (i.e. both forwards and backwards in time). The bias towards the longer periods in the sampling method used to generate the synthetic data is therefore likely responsible for this further underestimation of the true amplitude.

The other notable feature in Figure 6.12 is the increased spread around the true value at longer wavelengths / smaller wavevectors. This represents a fundamental limitation of LOFAR in identifying large scale waves. The wavelength is primarily determined from the gradient of the phase in equation (6.3), which decreases as wavelength increases. When the wavelength is significantly larger than the largest baselines in the array (~ 100 km), the uncertainty in the phase values will significantly degrade the estimated gradient and hence wavelength. The results here suggest that this decrease in performance is significant for wavelengths above roughly $400 - 500$ km for the moderate to high SNR values considered here. At low SNR, this effect would likely be even more severe.

As well as the performance of the estimated wave parameters themselves, it is useful to consider how well the uncertainties in these parameters were estimated. Figure 6.13 shows the normalised residuals (i.e. (estimated value - true value) / estimated uncertainty) for the various fit parameters. In order to accurately debias the amplitude residuals, it proved necessary to calculate the median percentage error independently in 200 s period bins from $300 - 1700$ s. The resulting values were consistent with the explanation proposed above for the bias, with the lower period bins showing biases of

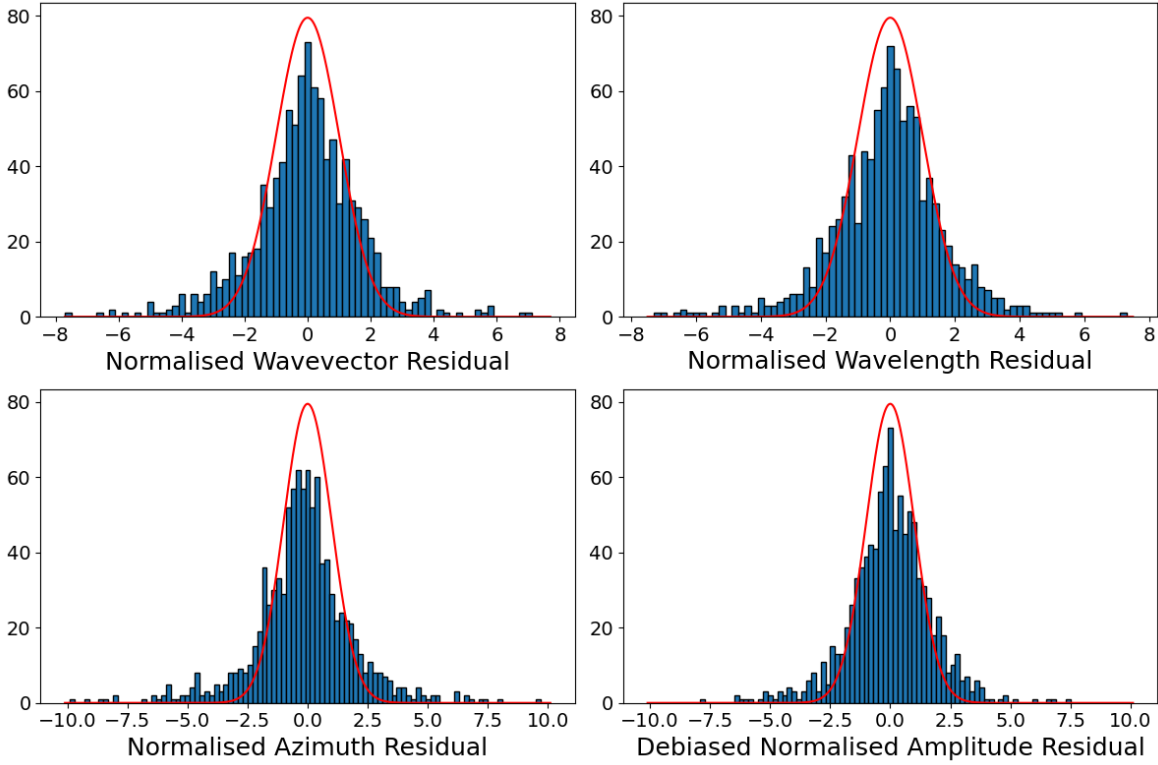


Figure 6.13. The normalised residuals of the fitted parameters relative to the true values. Amplitude estimates are first debiased by rescaling according to the median percentage error in each 200 s period bin. The vertical scale gives the number of samples in each bin. Normalisation consists of dividing the residual by the estimated uncertainty. Red curves show a unit Gaussian for reference.

$\sim -6.5\%$ before dropping significantly above ~ 900 s, reaching -21.8% in the highest bin. By comparison to the unit Gaussians (shown in red), all uncertainties seem to be slightly underestimated. The standard deviations of the normalised residuals in each parameter range from $\sim 1.68 - 2.15$, where a standard deviation of 1 would be expected for perfect uncertainty estimation.

Some of this underestimation of the errors may be related to mismatching of the true wave period to the wavelet period, as the fit relies on the closest wavelet period. Another factor is the imperfect nature of the noise estimation method. While it is unbiased (after rescaling), there is still appreciable variance of the estimated noise value around the true value. This will lead to the fit giving more weight than it should to certain

baselines and less to others, making the estimated parameters less accurate in a way which cannot be captured by the uncertainty estimates based on those noise estimates. A small improvement in the uncertainty estimates can be achieved by multiplying the uncertainty by the square root of the reduced chi-squared χ_r^2 of the final fit. This is defined as

$$\chi_r^2 = \frac{\chi^2}{2N_{base} - 4}, \quad (6.12)$$

where χ^2 is as defined in equation (6.11), N_{base} is the number of baselines included in the fit and 4 is the number of fit parameters. This rescaling is equivalent to ignoring the absolute value of the different uncertainty estimates and only considering their relative magnitudes. This rescaling provides standard deviations of $\sim 1.47 - 1.72$ for the normalised residuals. If this rescaling were applied to the example shown in Figure 6.11, the results would be: $\Lambda = 154 \pm 8 \text{ km}$, $\Delta \text{TEC} = 8.4 \pm 0.4 \text{ mTECu}$, $\theta = 84.4 \pm 0.7^\circ$. The significant difference in uncertainty estimates after rescaling in this case is due to the clear overestimation of the uncertainties in Figure 6.11 and is not representative of the impact of this change in the majority of cases.

6.5 Conclusions

The dTEC values calculated for calibration of astronomical observations with LOFAR represent a largely unexplored source of data on ionospheric dynamics over Europe. They are obtained with extremely high precision (Mevius et al., 2016), and there is already a large amount of data available thanks to the regular astronomical survey observations made using LOFAR (e.g. de Gasperin et al., 2021; Heald et al., 2015; Shimwell et al., 2017). These data can enable more detailed studies of TIDs over Europe, including lower amplitude and shorter wavelength waves, and therefore giving a more complete picture of the dynamics of the ionosphere and, ultimately, the neutral

atmosphere in which it is embedded. Such a statistical study of waves observed by LOFAR is the subject of the following chapter.

The method presented here for extracting information on waves present in these data is shown based on synthetic data to be capable of identifying waves with wavelengths up to ~ 500 km (roughly 5-10 times the size of the array). Larger wavelengths may also be identified but this is only possible for high amplitudes as the SNR required increases with increasing wavelength.

The approach developed here is also not highly specific to LOFAR, and could be adapted to identifying waves in data from other distributed networks. The specific fitted relationships would naturally differ in the case of absolute rather than differential measurements. However, the general method of quantifying the noise spectra in particular may be useful for a range of applications. This is especially true if the underestimate in parameter uncertainties shown in Figure 6.13 can be corrected.

As shown in Section 6.3.4, waves are detectable in LOFAR dTEC which would be below the detection threshold for GNSS TEC. This suggests that these data can complement existing work based on analysis of GNSS data (e.g. Otsuka et al., 2013) by extending the range of wave parameter space that can be sampled, particularly to lower amplitudes and possibly to shorter wavelengths as well. It also complements work using scintillation data from LOFAR which characterises TID structure on scales of ~ 20 km (e.g. Chapter 5 of this thesis; Dorrian et al., 2023).

Given the large number of very short (100s of metres to kilometres) baselines within the LOFAR core, a similar analysis method to that presented here but applied only to the core stations may be able to probe even shorter wavelengths than those $> \sim 100$ km considered here. However, this would be dependent on whether the correlation length shown in Figure 6.10 poses an obstacle to comparing signals between such similar baselines. In the following chapter it is shown that shorter wavelengths down to \sim

20–30 km can be identified by applying this method to the full Dutch LOFAR network, but it may still be the case that some additional short wavelength waves are being missed which may be identified by limiting the analysis to the core stations. Similarly, including the international stations may enable longer wavelength waves to be identified, extending above the ~ 500 km upper wavelength limit identified here. However, as the longest baselines are in the E-W direction and most large-scale TIDs propagate equatorwards from the auroral zone (e.g. Habarulema et al., 2013), this may not provide the most favourable baseline distribution.

Chapter 7

A Climatology of Waves Derived from Interferometric Calibration Solutions

This chapter describes a statistical study of the results obtained by applying the analysis method developed in the previous chapter to $\sim 2,700$ hours of LOFAR calibration solutions, and has been accepted for publication at the Journal of Space Weather and Space Climate (doi: 10.1051/swsc/2025002). This includes extending the method to allow automated identification and characterisation of waves. The high precision of the dTEC measurements obtained using LOFAR enables the study of TIDs in a high level of detail, investigating their occurrence rates, wavelengths, propagation directions and more, as well as detecting other disturbances in the plasmasphere.

7.1 Data and Processing

For this study dTEC solutions calculated from the observations made for the LOFAR LBA Survey (de Gasperin et al., 2021) were used. This data set comprises 2,723 one-hour observations using one of three calibrator sources (3C 196, 3C 295 or 3C 380; 43

additional observations did not have a dTEC solution successfully derived) based on which was closest in the sky to the target field for the observation. Almost all these observations were made during daytime as this has been found to reduce the severity of ionospheric scintillation observed with LOFAR (Fallows, private communication) and allow more accurate imaging (de Gasperin et al., 2021). This survey started in 2017, but the vast majority of observations ($> 90\%$) were obtained between June 2020 and March 2023.




Figure 7.1. The distribution of observing hours within the LBA survey as a function of Kp index (top left), F10.7 solar flux index (bottom left), and UT on the horizontal axis and month on the vertical (right). Local time is approximately $UT + 0.5$ hours. These distributions exclude observations for which the calibration solutions were too unreliable for analysis, see text below for details.

Figure 7.1 shows the distribution of observing time as a function of the Kp index, F10.7 index and UT / season. This clearly illustrates the fact that observations were made during day time (right panel), with good coverage across almost all seasons except

for the lack of observations during May. The left panels show that observations were predominantly made during low geomagnetic activity (top left) and low solar activity (bottom left). This was not due to deliberate scheduling during quiet conditions, but rather representative of the typical conditions during the later part of solar minimum and the start of the rising phase of solar cycle 25.

7.1.1 Data Flagging and Interpolation

Before using these dTEC solutions to derive wave parameters, the data was first processed to remove errors which would distort the results. Firstly, for some baselines at some time steps, the calibration solution failed to find a dTEC value and these points were flagged and filled with zeroes. Besides these flagged failed solution values, it was found that occasionally the dTEC value on a given baseline would change at a single time step by up to ~ 1 TECu before returning to its previous trend, an obviously unphysical variation. In order to identify these instances, the first difference of the dTEC series on each baseline was taken, and points in this series with a magnitude exceeding five standard deviations were identified as potential spikes, analogous to the approach used to identify RFI in LOFAR scintillation data (Fallows et al., 2020). If two consecutive points in the first difference time series exceeded 5 standard deviations and had opposite signs, the corresponding point in the dTEC series was flagged as a bad data point. If a given baseline had more than 5% of its data points flagged, it was excluded from further analysis. Observations with more than 40% of baselines removed in this way were not considered in the analysis, leaving 2,629 one-hour observations. Almost all observations required some data points to be replaced in this way, but the observations which were entirely removed were predominantly at low elevation ($< \sim 40^\circ$).

For baselines that were not removed from the analysis, the flagged data points were replaced by linear interpolation from neighbouring values, or extrapolation with

a constant value if they occurred at either end of the time series. This replacement of bad data was necessary as the spectral analysis method described in the previous chapter requires uniformly sampled time series of equal length on all baselines. Once the bad data points were replaced, the values were rescaled from differential slant TEC to the corresponding differential vertical TEC assuming a thin ionospheric shell at an altitude of 350 km. This correction is not exact for dTEC measurements due to the differences in elevation between different stations, but given that baselines are no more than ~ 50 km which corresponds to around 0.1° , this is negligible for baselines within the Dutch LOFAR network.

7.1.2 Event Identification

In order to identify waves in the large data set considered here, it was necessary to develop a method for automatically characterising the waves in a given observation. This is described in detail in Appendix E but the key points are summarised here. The method described in the previous chapter was applied at a sampled range of periods and times within the observation with sufficiently high SNR to suggest wave activity. These fit results were then passed to a clustering algorithm which identifies fits which were made for similar times and periods and which had consistent propagation directions and phase velocities. This was preferable to using the results of individual fits as any given wave activity will span a range of periods and times within the wavelet spectrum and should not be counted multiple times.

An example of the clustering results is shown in Figure 7.2. In this observation, three distinct wave events were identified, as highlighted by the coloured borders. As the clustering provides sets of points which correspond to the same wave event, each cluster can be used to calculate representative wave parameters for the whole event. First, the fit with the highest amplitude and all other fits which have amplitudes higher

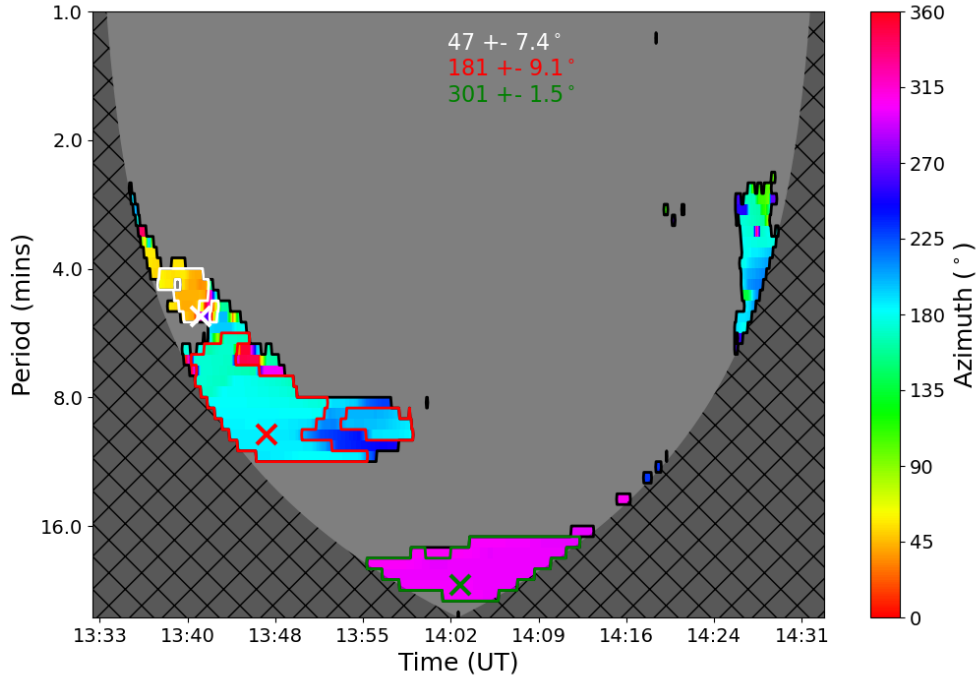


Figure 7.2. An example of the results of the clustering algorithm described above. Shown in the colours are the estimated azimuths as a function of time and period for an observation made on 25th July 2017 (observation ID: L601030). The coloured outlines indicate the three clusters identified within the fit results. The coloured crosses indicate the point within each cluster at which the highest amplitude was estimated. All coloured regions outside the clusters shown were classified as noise. The cross-hatched region shows the cone of influence.

than the maximum minus its estimated uncertainty were identified. The mean time and period corresponding to this subset of fits in the cluster defined the points shown as coloured crosses in Figure 7.2, the representative wave period and time of occurrence for the cluster. The representative wave amplitude and its uncertainty were calculated as the weighted mean and standard deviation of the amplitudes of this subset, with the weights being the inverse of the fit amplitude uncertainty multiplied by the period. The multiplication by the period was necessary due to the reduction in sampling density with increasing period (see Appendix E), to avoid biasing the results towards shorter periods. The representative phase velocity was then calculated in the same way, except in this

case using the whole cluster not merely the highest amplitude fits. The representative azimuth was calculated by finding the azimuth of a weighted sum of unit vectors along each fitted azimuth direction, again using the whole cluster. This was preferred to simply averaging the fitted azimuths as averaging azimuths runs into issues when the direction is close to the value at which the azimuth values wrap around (i.e. go from 360° back to 0°) and would provide unphysical values in those cases. The 5th and 95th percentiles of period and time within the cluster were also recorded as a measure of the extent of the wave activity.

7.2 Results

In total 7,195 wave events were identified within the 2,629 1-hour observations considered, with periods ranging from 1-27 minutes (the upper limit is defined by the length of the observing window). As shown in Figure 7.1, the observations covered low to moderate solar activity conditions and primarily low geomagnetic activity. Wave activity was observed for all solar and geomagnetic conditions as well as for all times and seasons.

7.2.1 Wave Occurrence Rates

The overall wave occurrence rate was determined by taking the total duration of the wave events in the data and dividing by the total duration of observations, binning by different conditions. The total duration of a given event was defined by the time between the 5th and 95th percentiles of times within the cluster, plus twice the cone of influence (CoI: the region in which the wavelet coefficients are unreliable due to edge effects from the ends of the time series) width at the representative wave period. For example, for the wave with a representative period around 10 min in Figure 7.2 (outlined

in red), the 5th and 95th percentiles are 13:42:18 and 13:57:40 UT respectively, and the CoI width is 11 min. Combining these gives a duration of around 37.5 min for this wave event. The addition of the CoI width factor was necessary as the period of time in which the wave activity is detected by the wavelet transform will not represent the whole duration of the wave activity. This is particularly clear at longer periods, such as the cluster outlined in blue in Figure 7.2, where the cluster is bounded on both sides by the CoI and it is unreasonable to assume that wave activity was not present throughout a much greater part of the observation than the cluster covers. Twice the CoI width is arguably a large correction in cases where the cluster is not actually limited by the CoI, but reducing this to adding only one times the CoI width was found to not significantly affect the variations in occurrence rate, only its magnitude. Hence, it does not change the conclusions that can be drawn about what conditions are more or less favourable for the waves observed by LOFAR.

This definition of the occurrence rate means that an occurrence rate over 100% is possible, which corresponds to multiple waves at different periods being observed simultaneously. The definition also does not include any weighting for the amplitude of the waves. This is because the observed amplitude is dependent not just on the physical properties of the wave in the ionosphere, but also the observing geometry as discussed in Section 6.1. In order to avoid biasing the results in this way, all wave events were considered equally for these calculations, regardless of amplitude.

Table 7.1. The number of wave events in each period range used for the analysis.

Period Range	No. of Waves	Mean Occurrence Rate
1 - 5 mins	2,131	12.9%
5 - 10 mins	1,732	27.5%
10 - 15 mins	1,386	35.2%
15 - 20 mins	844	30.2%
20 - 27 mins	1,102	42.6%
All Periods	7,195	148.5%




Figure 7.3. The occurrence rate of wave events as a function of UT hour on the horizontal and month on the vertical. The top left panel includes waves of all periods, other panels show only a subset of observed periods. All panels except the top left use a common colour scale capped at 100%. Only bins with at least 5 hours of observations are included, bins with fewer observations are coloured in grey.

The occurrence rate as a function of UT hour and month is shown in Figure 7.3, for all wave periods in the top left panel and for subsets of wave periods in each other panel. Focusing on the top left, the overall occurrence rate peaks between the winter solstice and spring equinox post midday, with a minimum occurrence rate in summer post midday. However, even at this minimum occurrence rate, the rate is still around 100%, indicating that wave activity of some sort is completely ubiquitous in LOFAR observations (note that this still does not mean that waves are guaranteed to be present at all times as many observations have multiple waves simultaneously observed at different

periods).

When the occurrence rate is split by observed period as in the other panels of Figure 7.3, however, differences in behaviour with period are clearly apparent. For example, the longest periods observed (20-27 minutes, bottom right panel) show peak occurrence rates around morning and evening in summer, and lower occurrence rates between winter and spring where the overall occurrence rate maximises. The peak between winter and spring only becomes apparent for periods below 15 minutes (top right and middle panels). This difference in the variation of occurrence rate for different wave periods indicates that either the sources of waves of different periods are different, or that the propagation conditions that favour AGWs of one period reaching ionospheric heights and being detected by LOFAR are not the same as the conditions favouring waves of another period. It is also clear that the average occurrence rate tends to increase for increasing periods, which is also shown in Table 7.1. This is despite there being more wave events at short periods, because the events at longer periods generally have longer durations and therefore contribute more to the calculated occurrence rate.

When splitting observations by period, there is a significant problem that must be considered, namely the Doppler shift between the observed period and the real physical period due to the motion of LOFAR's line of sight through the ionosphere. As long as this has a component along the wavevector, this will increase or decrease the apparent velocity and hence period of the wave as seen by LOFAR. This effect can be estimated for a given altitude, and is predominantly $< 10\%$ for an assumed altitude of 350 km (which is shown below to likely be an overestimate of the true altitude of the waves observed), meaning that for such a coarse binning by period the effect on the conclusions should be minimal. If the altitude of the structures could be estimated, this could be corrected, but this is not practical with the available data from the calibration solutions. The Brunt-Väisälä frequency ω_b is the natural frequency of vertical oscillations of air

parcels in the atmosphere, and typically corresponds to a period of ~ 10 min in the F-region ionosphere (e.g. Snively and Pasko, 2003). Waves with frequencies above ω_b will be unable to propagate effectively to higher altitudes, suggesting that the short period waves observed are likely to be at significantly different altitudes to the longer period waves (which are expected to predominantly exist at F-region altitudes (e.g. Fedorenko et al., 2011; Kirkland and Jacobson, 1998)). This prevents a single representative altitude being used to correct the Doppler effect. The shorter period waves may be in the E-region, where ω_b is typically higher than the F-region (e.g. Snively and Pasko, 2003), but waves at these periods have also been observed in the plasmasphere using similar observations from the VLA (Hoogeveen and Jacobson, 1997). To avoid these complications, the observed period was chosen as a means of distinguishing the waves for this analysis.

7.2.2 Wave Parameters

To place the observations made using LOFAR into context with other TID observations, it is also important to consider the range of wavelengths which are observed. Different observations and different analysis methods are more sensitive to different scales. This makes them complementary to one another, but also raises the question of whether any differences in the observed behaviour are due to real physical differences in the ionosphere-thermosphere system or systematic effects arising from the observing technique. The wavelengths of the waves detected with LOFAR are shown in Figure 7.4, with the same arrangement of panels as Figure 7.3. Firstly, this shows that the wavelength cutoff of 1000 km applied in processing the data had a negligible impact, as even for the longest periods observed very few waves are close to the cutoff. The longer periods tend to have larger wavelengths, with the modal wavelength varying between $\sim 80 - 160$ km between different period ranges, towards the short end of what

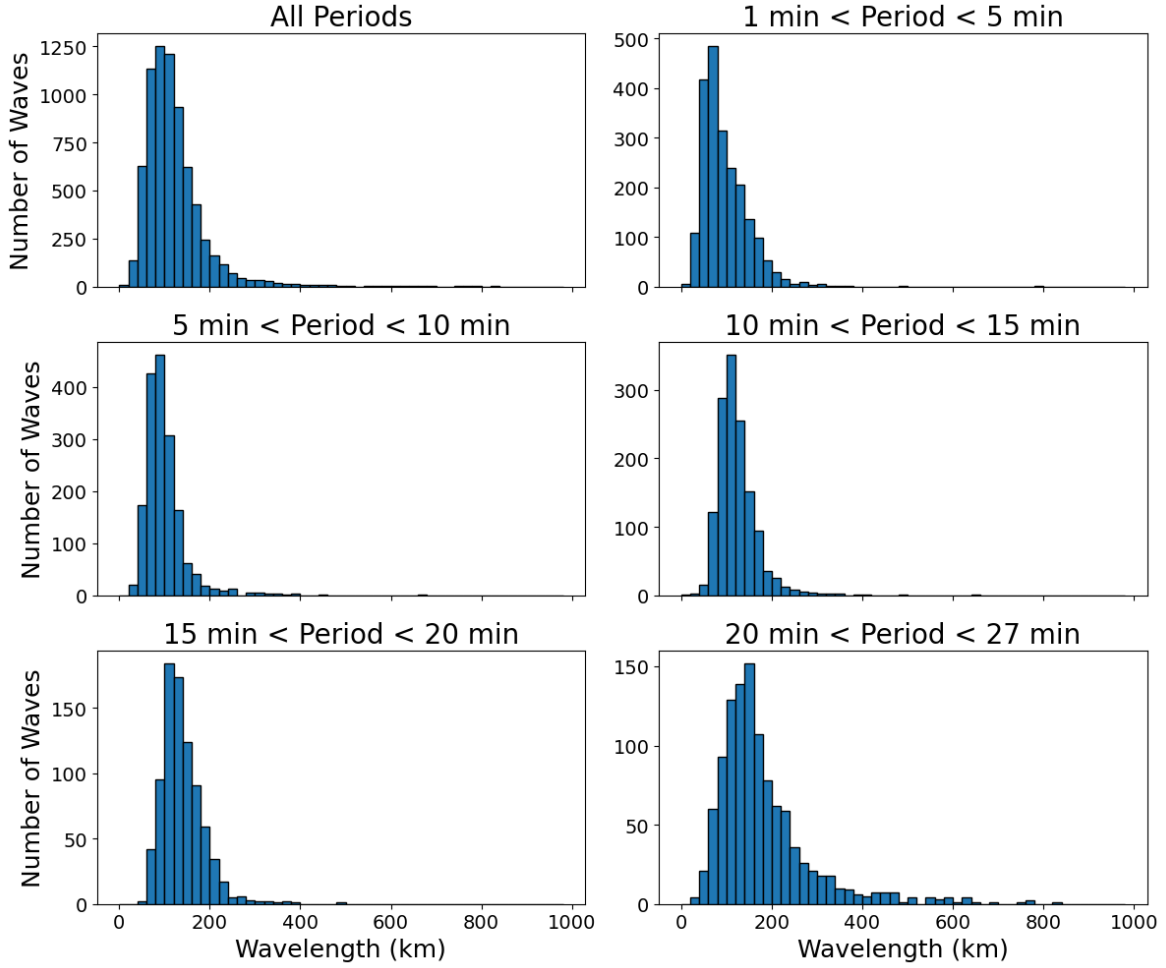


Figure 7.4. The distribution of observed wavelengths, with the panels corresponding to those in Figure 7.3. Note each panel uses a different scale for the vertical axis, but all are on the same horizontal axis scale.

is typically considered an MSTID.

The relative lack of waves observed at longer wavelengths is likely due to the reduced sensitivity of LOFAR to long wavelengths relative to the available baselines as discussed in the previous chapter. Also, as longer wavelengths will more often correspond to longer periods (as indicated by the larger tail of the distribution for periods of 20 – 27 min in the bottom right panel), many longer wavelength TIDs will have periods too large to be detected with 1-hour observing windows. Shorter wavelengths may also be suppressed to an extent, due to the assumption of an ideal plane wave spanning the whole LOFAR

array. For shorter wavelength waves this requires retaining the plane wave assumption over several wavelengths both along and perpendicular to the propagation direction, which is likely to not be the case for some of these waves.

A subtle difference observed in Figure 7.4 is the more significant tail of the distribution towards longer wavelengths for the lowest periods (1-5 minutes, top right) than at other periods. This may be an indication that at least some waves observed at these periods are from a different population to those observed at other periods, possibly being the waves in the plasmasphere reported by Jacobson and Erickson (1993) and Hoogeveen and Jacobson (1997).

The phase velocities of the detected waves are shown in Figure 7.5 in the same format as the wavelengths in Figure 7.4. They show that the majority of the waves have velocities between $\sim 50 - 250 \text{ m s}^{-1}$, as is typical for small- to medium-scale TIDs (e.g. Hunsucker, 1982; Hocke and Schlegel, 1996). The exception to this are the short period waves with periods below 5 min which show a much broader velocity distribution extending well above 1000 m s^{-1} . This is further evidence to support the idea that these waves are a distinct population to the longer period waves ($> \sim 5 - 10 \text{ min}$), and consistent with the short period waves being located in the plasmasphere where the higher effective scan velocity of the line of sight can cause such large apparent velocities (Hoogeveen and Jacobson, 1997; Jacobson and Erickson, 1993). The differences between the short period waves and the others are explored in more detail in Appendix F.

7.2.3 Propagation Directions

As mentioned in Section 2.1.1, the direction of AGWs, and hence of the TIDs they cause, is strongly influenced by the neutral wind direction, with waves propagating against the background wind able to propagate to higher altitudes and others filtered out (Cowling et al., 1971). No direct observations of the neutral winds in the thermosphere during

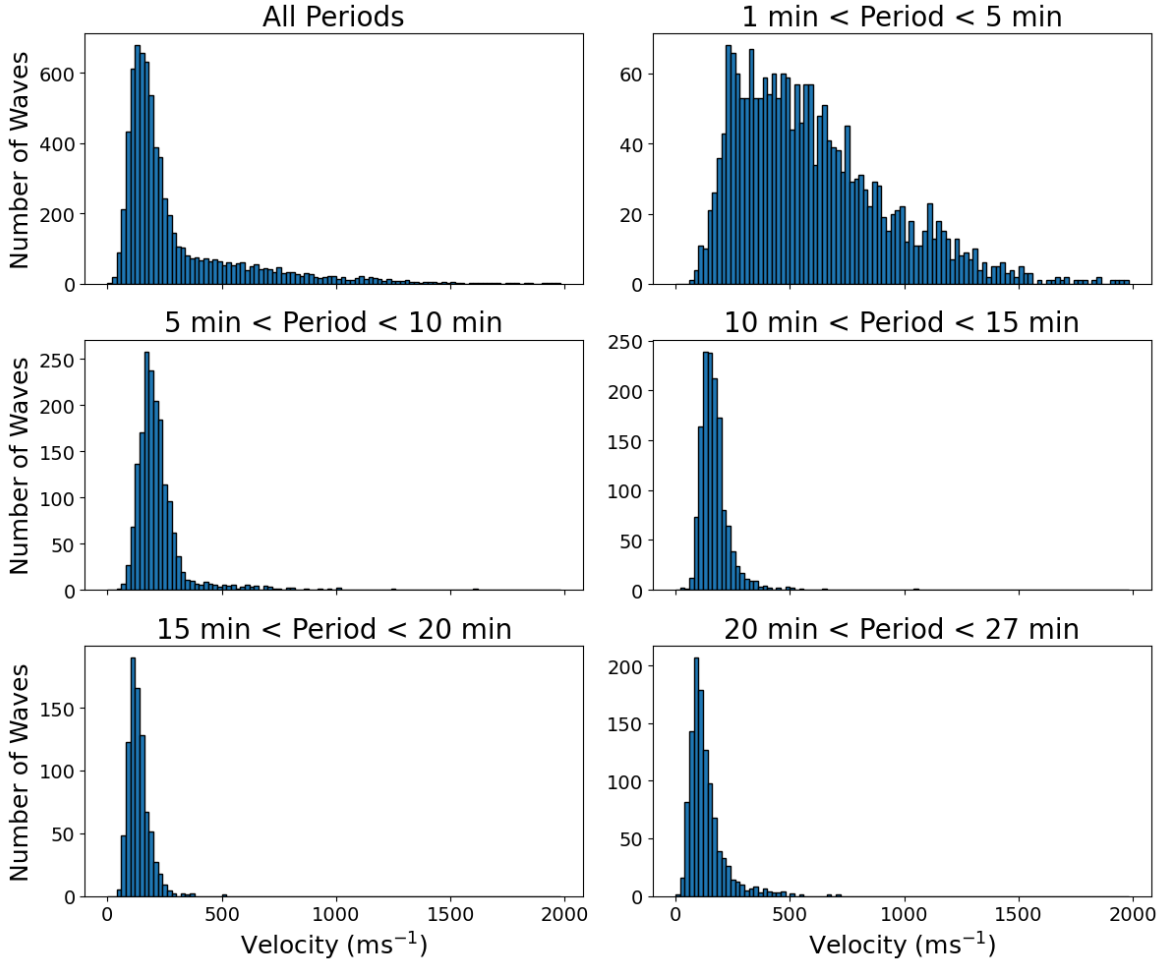


Figure 7.5. The distribution of observed phase velocities, with the panels corresponding to those in Figure 7.3. Note each panel uses a different scale for the vertical axis, but all are on the same horizontal axis scale.

these observations were available, but the Horizontal Wind Model 2014 (HWM14: Drob et al., 2015) can provide a climatological estimate of the winds. This model has substantial uncertainties in predicting the wind at a single time (root mean square error of $\sim 50 \text{ ms}^{-1}$ in each component, see Table 2 of Drob et al., 2015), but provides a good description of the systematic seasonal and diurnal variation in the wind. HWM14 has dependency on geomagnetic activity through the ap index (Matzka et al., 2021), although the impacts of this on wind direction in the mid-latitudes are typically very small, especially considering the generally low geomagnetic activity during the LOFAR

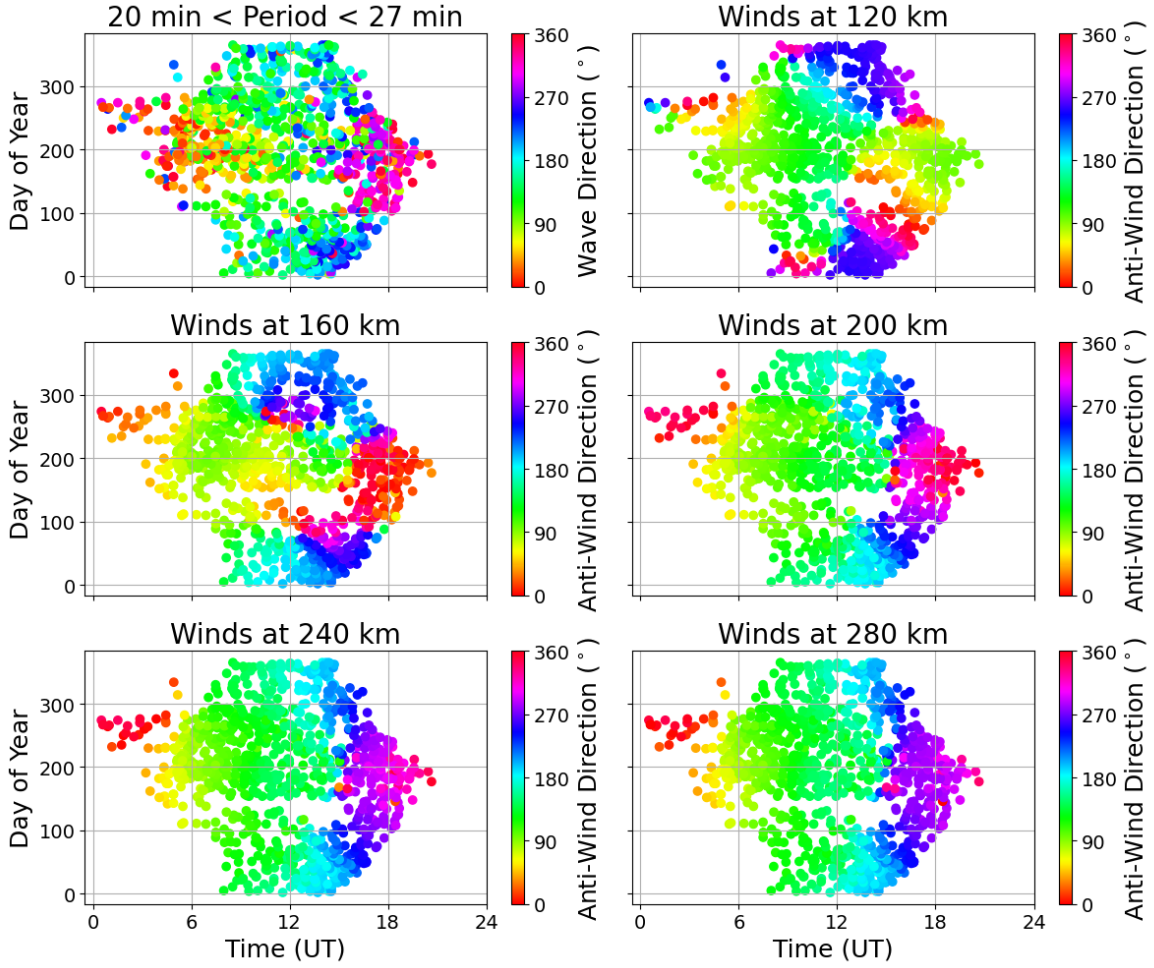


Figure 7.6. The variation of the observed wave direction for periods 20 – 27 min with time and day of year (top left), and the modelled anti-wind direction from HWM14 at a range of altitudes (other panels).

observations shown in Figure 7.1.

The altitude range at which the filtering occurs is uncertain, so the observations were compared to the modelled wind directions at various altitudes to determine if there is a relationship. The winds were calculated for the latitude and longitude of the IPP at 350 km from station CS002 (52.915°N, 6.870°E) at the time of the individual wave event. Choosing different altitudes for the assumed IPP was found to have negligible impact on the results (differences of a fraction of a degree in latitude and longitude) as observations are predominantly taken at high elevation. Overall the IPPs of wave

events ranged between approximately $0 - 14^\circ\text{E}$ and $52.5 - 60^\circ\text{N}$, with the majority between $3 - 9^\circ\text{E}$ and $52.5 - 54^\circ\text{N}$.

Figure 7.6 shows the observed wave directions for periods between 20 and 27 min in the top left panel, along with the corresponding HWM14 anti-wind directions at a range of altitudes from 120 – 280 km in the other panels, as a function of time and day of year. Simply by inspection of these plots, it is clear that the agreement between observed wave direction and modelled anti-wind direction is reasonable for altitudes above ~ 200 km and that below that it begins to break down. This strongly indicates that the TIDs observed by LOFAR are in the F-region.




Figure 7.7. The left panel shows the weighted standard deviation of angular offsets between HWM14 predicted anti-wind direction and observed wave direction for periods between 20 – 27 min as a function of altitude, while the right panel shows the angular offsets as a function of time and day of year for an altitude of 220 km, where the standard deviation is minimised. Negative offsets correspond to the wave being rotated anti-clockwise with respect to the anti-wind direction, and positive offsets correspond to clockwise rotation. The calculation of the standard deviation is weighted by the wind speed.

While the simple visual inspection of Figure 7.6 shows that the observed wave directions do correspond well to expected anti-wind directions in the thermosphere, and give an approximate sense of the relevant altitude, a more quantitative picture is provided by Figure 7.7. The left panel shows the weighted standard deviation (weighted by the modelled wind speed) of the difference between observed wave direction and modelled anti-wind direction for altitudes between 80 and 300 km with an altitude step of 10 km. The weighting reflects both the fact that the wind filtering effect is stronger for higher wind speeds and that when the expected wind speed is low the uncertainty in the wind direction is likely to be higher.

Figure 7.7 confirms what was apparent from Figure 7.6, with the lowest values of the standard deviation being at altitudes of 200 km and above. The standard deviation minimises at 220 km, and the distribution of offsets in time and day of year relative to the winds at this altitude is shown in the right panel of Figure 7.7. This shows that although the winds at this altitude provide a reasonable explanation for the observed variation of wave propagation directions, there are still unresolved aspects such as the overwhelmingly negative offsets for day of year 150-250 (roughly June-August) and UT <10 h. The systematic variation of the offset with time and day of year, rather than purely random noise, suggests that there is a systematic effect not accounted for in the analysis here such as variation in the optimum altitude with time and season. Such a variation could be explained by the AGW behaviour being primarily determined by pressure level rather than altitude, for example, but this is not explored further here.

It is important to emphasise that the filtering of AGWs by the wind does not happen at a single altitude, but is a process that is active throughout their vertical propagation from the lower atmosphere up to the thermosphere. No attempt is made here to estimate the integrated effect of this process on the observed waves, as the off-vertical propagation of AGWs and their finite vertical velocities cannot be accurately

accounted for with the data available in this study. The optimum altitude estimated here is also not a reliable measure of the altitude at which the TID amplitude maximises or of the ‘effective altitude’ of the integrated perturbations detected by LOFAR (e.g. Kirkland and Jacobson, 1998), but simply reflects the altitude at which the modelled winds best describe the observed seasonal and diurnal variations in wave directions.

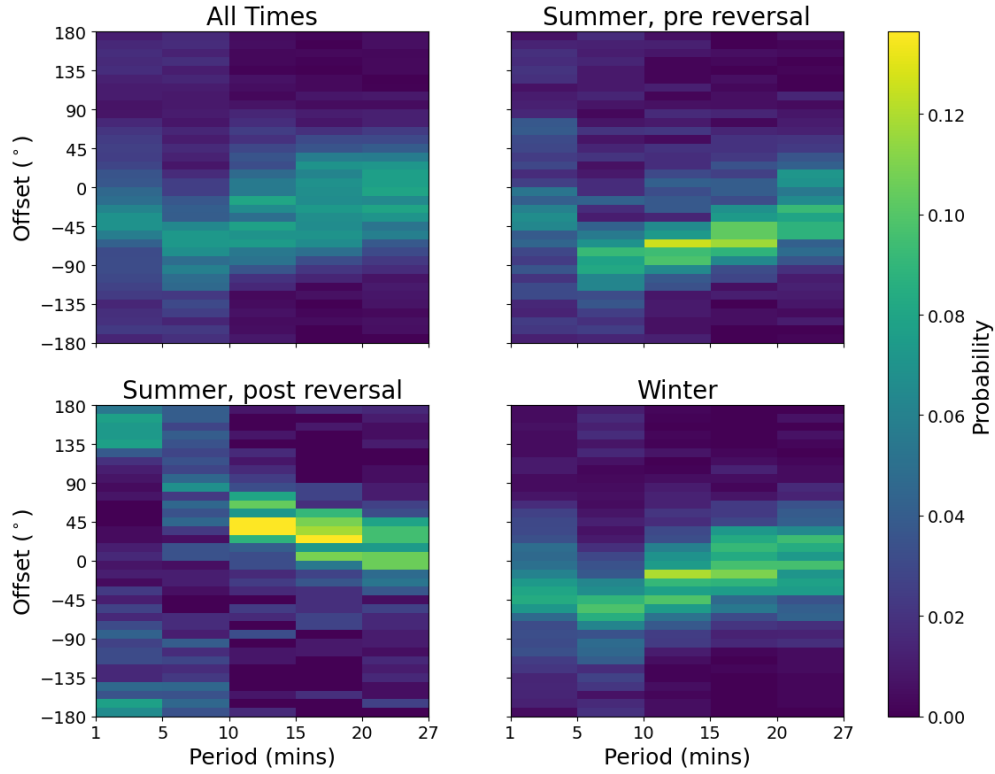


Figure 7.8. The distribution of wave-antiwind offsets for the different period ranges, each panel showing a different range of time and day of year. Summer is defined as day of year between 90 and 270 inclusive, whereas winter includes all days after day of year 270 and before day of year 90. Pre and post reversal refer to before and after 15:30 UT respectively. Winds are taken at an altitude of 220 km. The colours indicate the proportion of waves within the given period bin which fall into each 10° wide azimuth bin.

Based on Figure 7.7, an altitude of 220 km can be taken as reasonably representative for the longest period waves in the data. The anti-wind direction at this altitude was

then compared to the observed wave direction across the other period bins to investigate any systematic differences in propagation direction with period, which is shown in Figure 7.8. It is important to note at this point that given the typical Brunt-Väisälä period at 220 km altitude is ~ 10 min, it is not reasonable to expect waves with periods below 10 min to show strong agreement with the winds at this altitude. However, the comparison is carried out at all periods for completeness as it nevertheless highlights differences between the propagation directions of waves of different periods with respect to one another.

The vertical axis on each panel of Figure 7.8 shows the offset between the wave direction and anti-wind direction, with negative values corresponding to anticlockwise rotation of the wave with respect to the anti-wind direction as in Figure 7.7, while the horizontal axis distinguishes the period bins. The top left panel of Figure 7.8 shows the results including all wave observations, while the other three panels show subsets of time and season. The panels labelled Summer (top right and bottom left) include data with day of year between 90 and 270 inclusive, with pre reversal including all waves before 15:30 UT (2,541 waves) and post reversal including all after 15:30 UT (857 waves). This separation was made due to the clear reversal of both observed wave direction and modelled anti-wind direction around this time shown in Figure 7.6. The panel labelled Winter (bottom right) includes all observations with day of year less than 90 or greater than 270 (3,797 waves).

Figure 7.8 shows clear period dependence in the wave propagation directions and how they relate to the neutral winds. In all panels, there is a clear tendency for the waves to rotate with respect to the 220 km winds as a function of period (except for the 1 – 5 min bin which breaks this trend), but the sense and magnitude of this rotation is not constant. The clearest example of this is that for both ‘Summer, pre reversal’ and ‘Winter’, the rotation of the shorter periods with respect to the wind is anti-clockwise

(negative offset), whereas for ‘Summer, post reversal’ it is clockwise. Comparing the ‘Summer, pre reversal’ and ‘Winter’ panels more closely, it is apparent that despite the rotation having the same sense, the magnitude is different, with the 10 – 15 min period bin peaking around $10 - 50^\circ$ for ‘Winter’ and around $50 - 90^\circ$ for ‘Summer, pre reversal’.




Figure 7.9. The weighted standard deviation of angular offsets between HWM14 predicted anti-wind direction and observed wave direction as a function of altitude for the various period ranges (left panel) and the altitude at which the standard deviation minimises for each period range (‘optimum’ altitude, right panel). Standard deviation is weighted by wind speed as in Figure 7.7.

Given these differences in behaviour for different periods shown in Figure 7.8, the analysis shown in Figure 7.7 was repeated for the other period ranges to determine if different altitudes provide better agreement for the different bins. The results of this analysis are summarised in Figure 7.9, showing that the general trend is for the shorter periods to show better agreement with the winds at slightly lower altitudes

(again periods of 1 – 5 min provide an exception to this). Starting from the longest periods in the data, the optimum altitude, at which there is the best match between modelled winds and observed waves, decreases with decreasing period as shown in the right panel. This starts to break down for the 5 – 10 min period range, where the magnitude of the standard deviation also increases significantly as shown in the left panel. For the shortest periods of 1 – 5 min, the waves do not show any sign of tracking the diurnal rotation of wind direction and so there is no clear optimum altitude (see Figure F.1 in Appendix F).

The evolution of the observed propagation directions with time and day of year, along with the corresponding anti-wind direction at the optimum altitude, are shown in Figure 7.10 for the period ranges above 10 min (the shorter periods are shown and discussed in Appendix F). The general agreement with the solar heating driven wind variations for all periods, especially in terms of the diurnal variation, is clear. The differences between the winds for the three different altitudes shown in Figure 7.10 are relatively small, which explains the relatively small variations with altitude of the standard deviations shown in Figure 7.9. The greatest differences between modelled winds at different altitudes, such as the roughly northward propagation around 12 UT and day of year 275 in the bottom-right panel of Figure 7.10, correspond to small wind velocities ($< \sim 20 \text{ m s}^{-1}$) and hence have limited impact on the propagation of waves with typical phase velocities of $\sim 50 - 250 \text{ m s}^{-1}$. In terms of differences between the observed wave directions and the modelled winds, there is a clear tendency in the summer mornings for the waves to be rotated anticlockwise with respect to the winds, becoming clearer at the shorter periods. Other authors have suggested a time lag between the wind direction and observed wave direction from both observations (e.g. Waldock and Jones, 1986; Xu et al., 2024) and theoretical calculations (e.g. Cowling et al., 1971) due to the time taken for AGWs to propagate vertically from the altitude

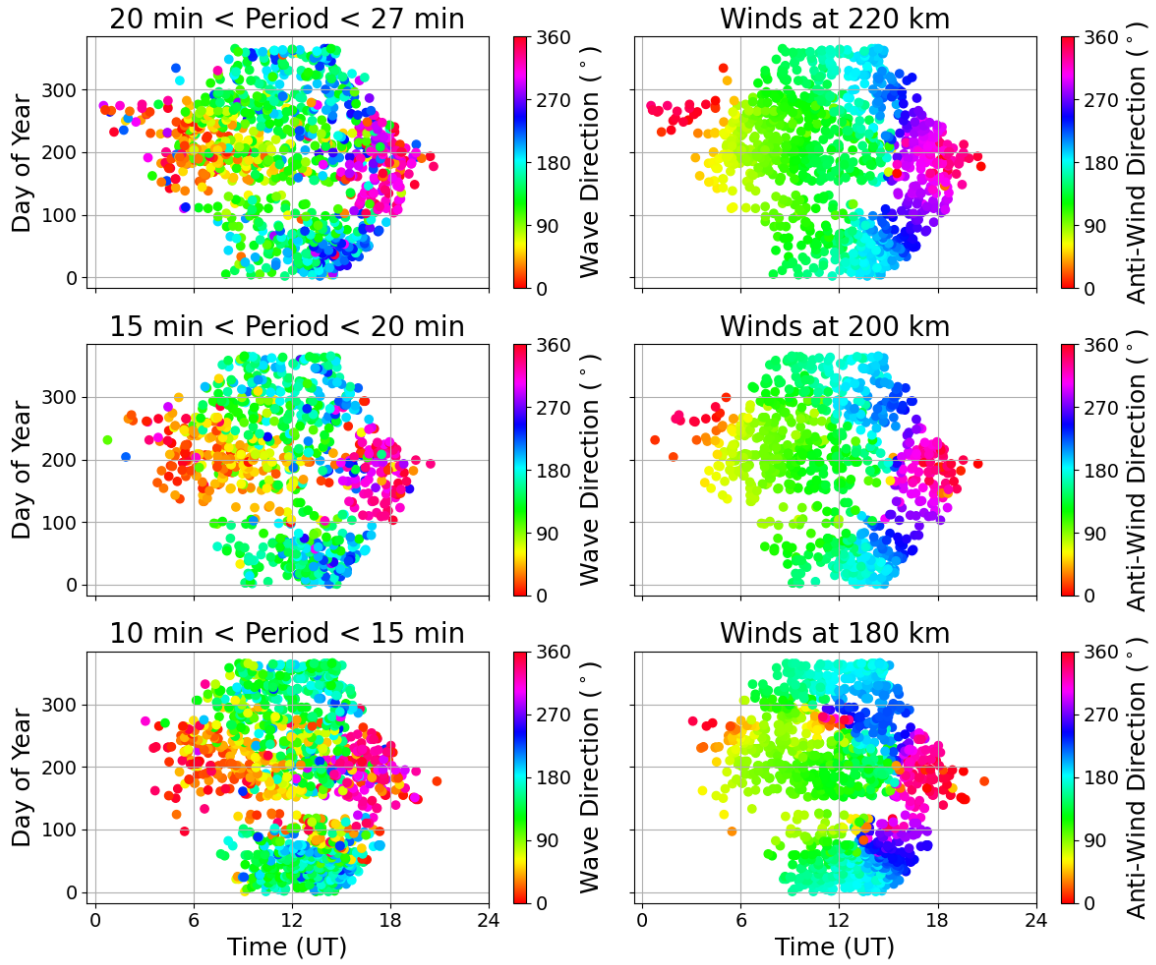


Figure 7.10. The left column shows observed wave propagation directions and the right column shows HWM14 modelled anti-wind directions at the corresponding times for the optimum altitude calculated in Figure 7.9. Each row corresponds to a different period range as indicated on each of the left hand panels.

at which wind filtering is most effective to the altitude at which they are observed as TIDs. However, the near immediate response to the change in direction in the afternoon (around 15-16 UT) in summer shows that the observed waves are able to track changes in the wind very rapidly.

As mentioned previously, the wind filtering explanation for the observed wave behaviour suggests that the filtering effect should be stronger for higher wind speeds relative to the wave phase speed. In order to test whether this is supported by the data




Figure 7.11. The left panel shows the distribution of wind-wave offsets as a function of the ratio of HWM14 modelled wind speed divided by observed wave phase speed (‘velocity ratio’), with the wind taken at an altitude of 200 km and wave periods from 10 – 27 min considered. The right panel shows the standard deviation of offsets in each velocity ratio bin of the left panel.

obtained using LOFAR, Figure 7.11 shows the distribution of wave-antiwind offsets as a function of the ratio of modelled wind speed to observed wave speed (‘velocity ratio’) in the left panel and the standard deviation of the offsets in each velocity ratio bin in the right panel. In order to carry out this part of the analysis, waves of all periods from 10 – 27 min were considered together to ensure there were enough waves in each velocity ratio bin. This is particularly necessary given the fact noted above that HWM14 is only reliable as a climatology and will have substantial uncertainties for any individual wave event. The winds were considered at an altitude of 200 km as this was the ‘optimum’ altitude for all these waves considered together derived in the same way as shown in Figure 7.9 (not shown). Figure 7.11 shows that for lower values of the velocity ratio

(i.e. when wind speed is significantly lower than the wave phase speed) the agreement between wave and anti-wind direction is significantly weaker than for higher velocity ratios. For velocity ratios greater than ~ 0.5 there is no discernible variation in how well the modelled wind directions predict the wave directions, which may reflect a physical limit but could also simply reflect the limitation of the accuracy of HWM14.

7.2.4 Wave Amplitude




Figure 7.12. The mean wave amplitude as a function of UT on the horizontal axis and F10.7 solar flux index on the vertical axis. The panels correspond to the same periods as the panels in Figure 7.3. All panels use the same colour scale. Bins are 1 hour by 20 sfu. Only bins with at least 5 wave observations are included, all others are coloured in grey.

Although the individual wave amplitudes should be treated with caution due to the impact of observing geometry, as discussed previously, the average wave amplitude under different conditions can provide useful information. Figure 7.12 shows the mean wave amplitude as a function of time and solar activity for all waves (top left) and subdivided into the different period bands in the other panels. As expected, it shows that higher amplitudes were observed under higher solar activity and post-midday when background ionospheric densities are highest and therefore the same AGW will induce a greater change in plasma density (see equation (2.29)). Comparing the different period ranges, there is little difference in the observed amplitudes for the highest bands (10 – 15 min, 15 – 20 min and 20 – 27 min) whereas the shorter periods show significantly lower amplitudes, especially in the 1 – 5 min period band. This supports the interpretation that the short period ($< \sim 10$ min) waves are dominated by a distinct population to the longer period waves as suggested by the difference in propagation directions relative to the neutral wind (see Appendix F for more details).

7.3 Discussion

7.3.1 Occurrence Rates and Possible Wave Sources

The occurrence rate of daytime MSTIDs in the mid-latitudes as reported in previous works were summarised in Section 2.5.1, but it is useful to recap the key results here to compare with the observations made using LOFAR. The occurrence rate is generally reported to be independent of solar and geomagnetic activity (e.g. Frisell et al., 2016; Otsuka et al., 2021), but show strong seasonal and diurnal variations. The seasonal dependence has a maximum in winter (e.g. Ding et al., 2011; Sivakandan et al., 2021) with some authors also reporting a second local maximum in summer (e.g. Fišer et al., 2017). The diurnal pattern is reported to vary with season, with a single peak around

midday in winter and two distinct peaks in summer, one in the morning and one in the afternoon (e.g. Fišer et al., 2017; Jacobson et al., 1995).

Comparing these to the results from LOFAR shown in Figure 7.3, there are broad similarities but some subtle differences. The overall occurrence rate (top left, with the caveat that this likely includes some non-ionospheric waves) does show a peak in winter and a minimum in summer as is generally reported. The diurnal behaviour is also broadly consistent with the pattern established by Jacobson et al. (1995) with two distinct peaks in summer and a single peak in other seasons (although the second peak identified by Jacobson et al. (1995) maximises around 20-23 LT, it becomes apparent around 16-17 LT similar to the observations made by LOFAR). Once this is broken down by period in the other panels of Figure 7.3, a slightly more complicated picture emerges. The winter peak is prominent for periods of 5 – 15 min, whereas for periods of 15 – 20 min there is no discernible seasonal variation. For periods of 20 – 27 min there are two peaks in summer, one in the morning and another in the evening close to sunset, although it is possible that this double peak is also present in winter but cut off by the limited time coverage of the observations.

These differences with period may suggest an explanation for some of the inconsistencies between the TID occurrence rates reported by other authors. Different observing techniques will preferentially detect waves with different parameters (e.g. period, wavelength) which can lead to different conclusions being drawn. For example, Otsuka et al. (2013) used a relatively high amplitude threshold of 0.2 TECu to identify TIDs in GNSS TEC data for their analysis (contrast with the ~ 0.001 TECu threshold used here) which suggests that their conclusion of near zero TID occurrence in Summer may instead reflect lower amplitudes. Due to the broader geographical coverage of GNSS receivers and continuous observations, however, Otsuka et al. (2013) were able to extend their analysis to much larger wavelengths (up to 1500 km) and longer periods

(10 – 60 min) than is possible with the LOFAR dTEC solutions. This emphasises the value of complementary instrumentation and analysis techniques to provide a complete picture of TIDs, as well as the need to account for the different biases when analysing and comparing the results.

The suggestion that reported differences in occurrence rates from different authors are at least somewhat due to variations in amplitude is supported by the results of Ding et al. (2011) and Jacobson et al. (1995), who report that TID amplitudes maximised in winter at the same time as the reported maximum occurrence rate. The variation in amplitude also has a more direct impact on the comparisons due to the fact that some authors estimate occurrence rates while others estimate ‘activity levels’ as discussed in Section 2.5.1.

Overall, the generally high occurrence rates reported here suggest that the dominant drivers of the observed waves are not sporadic impulsive events such as geomagnetic disturbances, earthquakes and major tropospheric storms. Rather, it points to drivers such as jet streams and the polar vortex (Buss et al., 2004; Sato and Yoshiki, 2008) as more likely candidates for the consistent generation of AGWs which then propagate to ionospheric altitudes.

7.3.2 Period Dependence of Propagation Direction and Inferred Altitudes

Some authors have previously reported variations in TID propagation with height or period. For example, Negrea et al. (2016) found that the propagation direction of observed TIDs in ionosonde data from October 2013 in the North American mid-latitudes rotated clockwise with increasing altitude during the daytime. However, this was only clearly observed for the longest periods in their data (frequencies < 1 mHz, i.e. periods $> \sim 17$ min) with shorter periods showing much weaker directional preference,

and is limited by averaging all daytime observations together thereby removing any information on variations with local time. This absence of local time information prevents comparison with neutral wind directions as these vary substantially with local time (see e.g. Figure 7.6). Subsequently, using roughly 3 years of data from the same location Negrea et al. (2018) showed that the activity of shorter period waves was higher at slightly lower altitudes, which was apparent for frequencies up to ~ 1.5 mHz (periods above ~ 11 min). This is consistent with the results presented here showing better agreement between wave directions and modelled neutral winds at lower altitudes for lower periods, and the apparent break down of this relationship for periods below ~ 10 min.

On a case study basis, a further example is provided by Crowley and Rodrigues (2012) who showed TIDs with $\sim 15 - 30$ min periods displaying anticlockwise rotation with increasing period in data from HF Doppler sounding (see their Figure 7, red box labelled ‘Medium Scale’). These observations were made between 22-01 LT in the North American sector, meaning they fall outside the geographical and time of day coverage of the present study, but they provide further evidence that this rotation of TIDs with period is detected by a variety of techniques. In this case, it is also interesting to note that as Crowley and Rodrigues (2012) used HF Doppler sounding, all measurements were obtained at the same altitude. This suggests that the observed rotation with period in the data obtained using LOFAR may not be entirely due to the waves peaking in amplitude at different altitudes, but rather the neutral winds at different altitudes having stronger effects on waves of different periods, with these effects persisting as they propagate to higher altitudes.

Extensive theoretical calculations on the vertical propagation of AGWs have been carried out by Vadas (2007), showing the dissipation altitudes (altitude at which the AGW momentum flux maximises) for various AGW parameters. For the range of wave

parameters observed here (periods up to 27 min, velocities of $50 - 250 \text{ m s}^{-1}$, wavelengths of $50 - 250 \text{ km}$) the dissipation altitude would be expected to increase with decreasing period (see their Figures 6 and 7, increasing from $\sim 200 \text{ km}$ for 20 – 27 min waves to $\sim 225 - 250 \text{ km}$ for 10 – 15 min for moderate solar activity). This seems to run counter to the observations from LOFAR presented here which show wave propagation directions at shorter periods better explained by wind directions at lower altitudes. One possible factor that could explain this discrepancy is that the vertical wavelength of the shorter period waves reported here would be much greater than the longer period waves (Vadas, 2007), potentially making the shorter period waves more sensitive to the winds at lower altitudes. With this explanation, the observations at all periods could be dominated by contributions from the same altitude (e.g. the altitude of the F2 ionospheric density peak) rather than arising from different altitudes determined by individual wave dissipation altitudes. However, this explanation would not explain the direct observations of shorter period waves being more significant in the ionosphere at lower altitudes made by Negrea et al. (2018).

Previous authors have suggested similar values for the altitude of waves observed in TEC to those at which the direction of the waves considered here best match the neutral winds. For example, Kirkland and Jacobson (1998) used a radio interferometer in New Mexico which observed signals from two closely spaced geostationary satellites to deduce the effective altitude of the observed TEC perturbations by parallax, and found typical altitudes around $180 \pm 40 \text{ km}$, comparable to the range of $180 - 220 \text{ km}$ found by matching the wind directions on a statistical basis here. They also reported that events with higher amplitude were biased towards higher altitudes, although still mostly within the standard deviation of the overall altitude distribution. The data reported by Kirkland and Jacobson (1998) covered only October to March, and they do not provide information on the typical wave parameters (e.g. period, wavelength)

in their observations which limits the ability to make more detailed comparisons.

The relationship between the different period waves observed by LOFAR and the neutral winds at different altitudes, and the uncertainties in the underlying physical mechanism, suggest a useful direction for future work. Modelling work to investigate the AGW propagation under the conditions relevant to the LOFAR observations, case studies incorporating data from other instruments such as ionosondes which can provide altitude resolved measurements, or a combination of these approaches, may enable the causes of this relationship to be determined.

7.4 Conclusions

The calibration solutions derived as part of astronomical observations with LOFAR provide a new source of TID observations. Due to the extremely high sensitivity of the measurements, they reveal parts of the TID population which may be inaccessible with other more traditional techniques. However, this does not mean they are without limitations, such as the lack of altitude information and limited spatial and temporal coverage. This highlights the need for a range of measurements to fully characterise the TID population, and for data from any given technique to be interpreted carefully to avoid confusing a bias or limitation of the technique for a characteristic of the TIDs themselves.

Using the LOFAR calibration solutions and the analysis method developed in Chapter 6, it has been shown that TIDs and other wave-like disturbances are ubiquitous above LOFAR during the daytime. No relationship with geomagnetic or solar activity could be identified, but this must be qualified by acknowledging that the observations span a relatively limited range of geomagnetic activity levels and only part of a solar cycle. Overall, the diurnal and seasonal occurrence rates show a complicated pattern, with

different period ranges showing activity peaks at different times. However, generally the high sensitivity of the measurements made with LOFAR indicates that variations in TID occurrence with time and season are less dramatic than has been suggested previously.

LOFAR reveals wave activity across a wide range of wavelengths, phase velocities and propagation directions. As far as the wavelengths are concerned, it is difficult to make definite statements about relative occurrence rates for waves of different wavelength due to the limitations on LOFAR's sensitivity imposed by the finite baseline lengths and relatively short observing windows as discussed in Section 7.2.2. The wave amplitudes on average behave exactly as would be expected, with peak amplitudes around midday or just afterwards, and increasing amplitude with increasing solar activity, due to the corresponding increases in ionisation. As a given AGW induces a given relative perturbation in electron density (Hooke, 1968), the increase in the background ionisation causes an increase in the observed absolute amplitudes.

The wave propagation directions show clear evidence that the dominant mechanism controlling them is filtering by thermospheric winds. Comparison across period ranges has indicated that the altitude at which the observations best match the climatological winds varies with period, with shorter periods corresponding to lower altitudes. As LOFAR does not provide any direct altitude information, it is unclear from this data whether the shorter period waves are actually being observed at lower altitudes or if they are simply more controlled by winds at lower altitudes due to having longer vertical wavelengths, as discussed in Section 7.3.2. For the periods below 10 min, this relationship with the neutral winds breaks down, as might be expected given that these waves would be below the Brunt-Väisälä period at the altitude of the F-region ionosphere and hence less likely to be explained by thermospheric AGWs. The possible contribution of plasmaspheric waves in this period range, especially below 5 min, as

discussed in Appendix F may provide a partial explanation for the significant differences to behaviour at longer periods.

In general, the observations with LOFAR show that the dominant wave population detected are TIDs associated with upward propagating AGWs. There is no apparent relationship with geomagnetic activity, suggesting that the waves are primarily driven from within the Earth system. These observations therefore provide a means of investigating vertical coupling between the lower atmosphere and thermosphere.

Chapter 8

U-Shaped Scintillation Features

Across the ionospheric scintillation observations made using LOFAR, one recurring type of event is what will be referred to here as ‘U-shape’ features. These events are characterised by moderate to intense scintillation occurring above a frequency cutoff with essentially no scintillation below this frequency. The cutoff frequency varies slowly with time, initially increasing from the low end of the observing band (25 MHz) and later doing the reverse, forming a U-shaped curve in the dynamic spectrum. The whole feature typically lasts for 5-7 hours.

Examples of these features seen from station CS002 are shown in Figure 8.1 (note that these are quicklook plots, which is the reason for the difference in colour scale and horizontal axis compared to other dynamic spectra shown in Chapter 5 and later in this chapter). All of these examples are seen in Cygnus A, as is the case for the vast majority of U-shape features observed. Only a small number of these features have been observed in Cassiopeia A, and they have always been ‘shallow’, that is to say the frequency cutoff never extends to frequencies above ~ 40 MHz. This suggests that the occurrence of these features is likely to be highly dependent on observing geometry, as the main difference between the two sources in terms of the ionospheric scintillation they




Figure 8.1. Four examples of U-shape features observed in Cygnus A. All are taken from the ‘quicklook’ plots routinely produced by ASTRON. One sample is approximately 1 second. The start times of the observations are: L843074 - 14/01/2022 16:00 UT, L878472 - 08/12/2022 22:01 UT, L880388 - 03/01/2023 17:15 UT, L881256 - 31/01/2023 19:00 UT.

experience is their positions in the sky. The examples shown here are all ‘complete’ U-shapes, where the observing window started before the feature was apparent and continued through until it had receded. This is not the case in most observations due to the long duration of the U-shape features, which means many observations only capture part of the feature, such as only one half of the overall U.

8.1 Occurrence Statistics

For the purposes of analysing the events statistically, it is useful to distinguish between ‘descents’ as the first phase of the event where the cutoff frequency is increasing with

time, and ‘ascents’ as the second phase where the cutoff frequency is decreasing with time (these are named based on their visual characteristics in dynamic spectra, where low frequencies are towards the top of the plot). Most observations only see one phase of the event or the other due to the extended duration of the events. Each phase is then characterised by a start and end time, covering the region in which the cutoff frequency is appreciably evolving with time and within the observing frequency band. This means that these times will be influenced by how ‘deep’ the event is, that is to say how high the cutoff frequency becomes in the middle of the event. Due to the volumes of data produced by LOFAR, it was impractical at the time of writing to download and analyse every event individually, as LOFAR ionospheric scintillation observations comprise $\sim 30 \text{ GB h}^{-1}$ of data. Considering only observations with suitable geometry for U-shape occurrence (see below for how this was defined) the archive contains 464 such observations with a combined duration of approximately 110 days. This gives a total data volume of $\sim 80 \text{ TB}$, which was not practical with the resources and time available.

Instead of downloading the data and working with it directly, the timings of the U-shape features were estimated visually by looking at the ‘quicklook’ plots such as those shown in Figure 8.1, meaning that there is some uncertainty in the exact timings. These quicklook plots were downloaded and stored locally for all observations starting from October 2021. By repeating the timing estimation for all events when new observations were added to the list, it was possible to estimate the consistency of these estimated timings. The values were found to be consistent to within $\pm 5 \text{ min}$, but this does not allow an estimate of any systematic bias in the manual estimation process. For this analysis, only U-shape features observed in Cygnus A from the LOFAR core stations were considered. The small number of U-shape features observed in Cygnus A only from international stations or in Cassiopeia A are discussed at the end of this section.




Figure 8.2. The azimuth and elevation values corresponding to the observed start and end times of each half of the event as seen from CS002. The apparent motion of Cygnus A with time is clockwise in this plot.

The observing geometry under which the U-shaped features are observed, shown in Figure 8.2, is clearly a controlling factor. All descents occur as Cygnus A is decreasing in elevation and approaching due North, while all ascents occur after it passes due North and is increasing in elevation. From event to event, there is significant variability in exactly how long before/after passing North the feature is observed, but none extend much above 20° in elevation.

Given that U-shape features are only observed at such low elevations, it is reasonable to question whether this is simply an instrumental effect. LOFAR was not originally designed to observe at such low elevations (astronomical observations are not typically

made below $\sim 20 - 30^\circ$, e.g. de Gasperin et al., 2019), and the sensitivity of any phased array decreases with distance away from zenith (see discussion in Section 5.2.1). The ionospheric scintillation observations are only able to reach such low elevations due to the brightness of Cygnus A at these frequencies. If the U-shape scintillation were purely an instrumental effect, minimal variation between events would be expected as the instrument and observing geometry are consistent. However, as shown in Figure 8.2, there is substantial variation in the timings of the events, which significantly exceeds the $\sim \pm 5$ min uncertainty in the timing estimates (this corresponds to $\sim 1^\circ$ in azimuth and substantially less in elevation). This, along with the other results detailed in the remainder of this chapter, provides strong evidence that the U-shape features are genuine ionospheric scintillation, although the details of their appearance may certainly be influenced by instrumental factors. The variation of LOFAR’s sensitivity with elevation and observing frequency for Cygnus A specifically is addressed in more detail in Section 8.3.3.

The events including both an ascent and a descent can be examined to determine whether they are symmetric around due North, which may indicate the approximate timescale of whatever process is responsible for the variability shown in Figure 8.2. If they are strongly symmetric, this would suggest that the controlling process varies on a timescale of significantly more than 6 hours, whereas weak correlation could indicate that it varies appreciably over this timescale. The correlation between elevation at the ‘outer edges’ (descent start and ascent end) is ~ 0.4 and for the ‘inner edges’ (descent end and ascent start) it is ~ 0.2 . This therefore may suggest some variability in the controlling process, but may also be negatively influenced by any potential bias in estimating event timings from the quicklook plots. Given that the correlation is weaker for the inner edges than for the outer edges despite the inner edges being closer together in both space (i.e. observing geometry) and time, the uncertainty in timing estimation

is likely to be a dominant factor in this.




Figure 8.3. The times in UT of all events observed from CS002 as a function of day of year. Also shown in black is the time at which Cygnus A is due North of CS002.

The dependence on observing geometry also dominates the relationship between time of day and day of year shown in Figure 8.3. This is because the observing geometry for a given source is a function of sidereal time, meaning that it drifts relative to local solar time over the course of the year. This is an inherent limitation of using a single astronomical source to probe the ionosphere (e.g. Wild and Roberts, 1956b), and in this case no other sources of a suitable intensity were available to resolve this ambiguity. However, despite this correlation between time of day and season in the observations it is clear that the majority of the U-shaped features are observed during winter and at night. It is also important to note that the observing time given to ionospheric scintillation observations on LOFAR is generally higher at night than during the day,

which may also contribute to the bias seen in Figure 8.3. However, even accounting for this, it will be shown subsequently (see Figure 8.4) that the occurrence rate is indeed higher in winter nighttime than summer daytime.

Using the clear dependence on observing geometry shown in Figure 8.2, it is possible to define a set of all LOFAR observations in which a U-shape feature is possible (candidate observations). From this, the occurrence rate of U-shapes can be estimated under different conditions without any contamination from the distribution of observations. The criteria for identifying candidate observations were that they must: first, be at least 2 hours in duration; and second, include the period 1.5-2.5 hours before and/or after Cygnus A passes North as viewed from CS002. The offset either side of North was based on the distributions of inner and outer edge timings, which had means of ~ 1 hour and ~ 3 hours respectively and standard deviations of ~ 1 hour. This means that if a U-shape feature occurred on a given day, it should be reliably present in the data around 2 hours either side of due North. Only observations made after October 2021 were considered as these were the only ones for which quicklook plots were available.

The candidate observations defined in this way were then all manually reviewed to ensure that all U-shapes were identified. Observations with strong contamination from RFI were removed as it is impossible to confirm whether a U-shape was present or not in these cases. In total, there were 255 candidate observations identified, of which 166 contained at least one side of a U-shape feature (2 of the 32 observations which were expected to contain both a descent and an ascent contained only one or the other).

The occurrence rates as a function of month and time of day estimated based on these candidate observations are shown in Figure 8.4. The rate as a function of month clearly maximises around winter/early spring and the rate as a function of UT maximises somewhere between 16-04 UT. This demonstrates that the bias towards winter/nighttime in Figure 8.3 is not merely a reflection of the greater observing time




Figure 8.4. The estimated rates of U-shape features as a function of month (left) and UT (right). Rates are indicated for descents (blue crosses, dashed line) and ascents (red triangles, dash-dotted line). The bar charts indicate the number of candidate observations in each bin for each type, normalised to a maximum of 1 (29 observations in January descent bin, 29 observations in 20-22 UT descent bin).

during nighttime (shown by the coloured bars). Comparing the ascents and descents, the variations in rate appear to be more aligned as a function of UT than month, which is a possible indication that the seasonal bias is primarily a function of observing geometry and time of day and not a preference for winter per se, although no strong conclusions should be drawn from this qualitative comparison.

To compare the occurrence rate to the occurrence of geomagnetic disturbances, two indices of geomagnetic activity were considered. The Kp index (Matzka et al., 2021) is derived from magnetic field measurements at 13 sub-auroral stations and has a 3 hour cadence. It has been shown to correlate with many variations in the magnetosphere, ionosphere and thermosphere (Matzka et al., 2021, and references therein). SME is a

dedicated auroral index calculated at 1 minute cadence, using data from roughly 130 stations between $40^\circ - 80^\circ$ geomagnetic latitude (the exact number of stations used varies depending on data availability, Newell and Gjerloev, 2011a; Gjerloev, 2012). It is a measure of the intensity of the auroral electrojet, and has been shown to strongly correlate with total auroral power (Newell and Gjerloev, 2011a) and substorm activity (Newell and Gjerloev, 2011b).




Figure 8.5. The estimated rates of U-shape features as a function of Kp (left) and SME (right). Rates are indicated for descents (blue crosses, dashed line) and ascents (red triangles, dash-dotted line). The bar charts indicate the number of candidate observations in each bin for each type, normalised to a maximum of 1 (73 observations in $Kp < 2$ descent bin, 41 observations in $80 \text{ nT} < SME \leq 150 \text{ nT}$ descent bin).

The variation of U-shape occurrence rate with geomagnetic indices is shown in Figure 8.5. The Kp values used here are taken based on the time 2 hours before the expected midpoint of the descent or ascent (i.e. taken 4 hours before Cygnus A is due North for descents and taken when Cygnus A is due North for ascents). The

SME values used are the mean value in the 2 hour window before the midpoint of the descent or ascent (i.e. from 4 to 2 hours before Cygnus A is due North for descents and from when Cygnus A is due North to 2 hours after for ascents). The lead times were selected as those which maximise the observed correlation between the index and U-shape occurrence rate (although given the 3 hour cadence of Kp and that it is based on a maximum value within that 3 hour interval, it is still possible the Kp value used will be based on a disturbance occurring after the relevant phase of the U-shape). The bins for Kp were defined to represent ‘quiet’, ‘moderate’ and ‘disturbed’ conditions respectively, while the bins for SME were defined to provide roughly even sample size in each. Both indices show a clear tendency for higher occurrence rates during elevated geomagnetic activity, the relationship being stronger for Kp. It is also clear that for both indices this relationship is stronger for the ascents than descents. This was consistently the case for all lead times considered in calculating the indices, suggesting it is a genuine difference between the two portions of the event and not an artefact of the analysis. However, as yet there is no proposed explanation for this difference in behaviour as both of the physical mechanisms proposed in the following section suggest symmetry between ascents and descents.

It is useful here to briefly discuss the characteristics of the U-shape features not considered in the previous plots: those observed in Cygnus A only from international stations or in Cassiopeia A. Those observed only from international stations were typically restricted to the northernmost LOFAR stations (those in Sweden and Latvia) which are in fact those from which the minimum elevation of Cygnus A (or any other source for that matter) is greatest. This may indicate that rather than requiring a low elevation per se, the U-shape features require a line of sight passing through a specific region of the ionosphere to the north of LOFAR. Those observed in Cassiopeia A were far less common than in Cygnus A (5 events identified compared to 166 in Cygnus A)

likely due to the differences in observing geometry. The minimum elevation reached by Cassiopeia A as viewed from the LOFAR core is approximately 20° , while the U-shape features observed in Cygnus A are largely restricted to elevations below 20° as shown in Figure 8.2. The features which are observed in Cassiopeia A all occurred under substantially disturbed geomagnetic conditions, consistent with the increase in occurrence rate with increasing geomagnetic activity for features seen in Cygnus A that is shown in Figure 8.5.

8.2 Possible Physical Explanations

While these statistics provide a lot of useful information about the conditions under which these U-shape features occur, they do not give any direct indication of why they occur. The statistics do strongly indicate that the scintillation is of real ionospheric origin rather than simply an instrumental effect, given the significant inter-event variability shown in Figure 8.2 and the relationships with UT/season and geomagnetic activity shown in Figures 8.4 & 8.5. The restriction of scintillation to only exist above a certain frequency suggests some corresponding physical quantity controlling the cut-off. This may be a property of the irregularities that produce the scintillation, such as a characteristic spatial scale, or a property of the background ionosphere that they exist in, such as refraction causing some frequencies not to pass through the region in which the irregularities exist.

It is useful here to briefly re-summarise the key characteristics of the U-shape features identified in the previous section. They occur at low elevations on northward lines of sight, preferentially at night and/or during winter and preferentially during elevated geomagnetic activity. When they are observed, they typically span the whole international LOFAR network, indicating that the structures responsible are not highly

localised within the ionosphere, and the long duration of the events indicates that the structures are present over long timescales rather than being sporadically generated.

8.2.1 Irregularity Scale

As discussed in Section 3.1, the Fresnel scale defines a characteristic scale size for irregularities which cause scintillation. As this is frequency dependent, this means that irregularities of a given scale will only create scintillation in a certain range of radio frequencies. In general, irregularities do not occur at a single scale, but can be characterised as existing across a range of scales. Irregularity mechanisms typically have preferred scales at which their growth rate maximises (e.g. Eltrass et al., 2016), providing a maximum scale for the irregularities generated. Once these reach sufficient amplitude, the linear behaviour breaks down and smaller scale irregularities are generated by non-linear processes, creating an energy cascade to smaller scales. The lower scale limit then occurs due to the smallest scales experiencing the greatest dissipation due to viscous effects.

For irregularities at a given distance along the raypath, these inner and outer scales can then be related to an upper and lower frequency respectively based on the Fresnel scale using equation (3.5). The cutoff frequency observed in the U-shape features therefore could correspond to the outer scale of the irregularities responsible for the scintillation.

Altitude (km)	Fresnel Scale (km)
100	2.52
200	3.36
300	3.93
400	4.38

Table 8.1. The Fresnel scales corresponding to a cutoff frequency of 45 MHz at a source elevation of 10° . Path length is calculated assuming a spherical Earth and no refraction according to equation (5.3).

Given the geometric constraints on the U-shape occurrence, the Fresnel scale for various ionospheric altitudes for a typical U-shape can be estimated. As an approximation, consider a source elevation of 10° and cutoff frequency of 45 MHz, typical of the midpoint of an ascent or descent. Fresnel scale values for assumed altitudes between 100 – 400 km are shown in Table 8.1. This suggests that an outer scale of a few kilometres could potentially cause a cutoff in scintillation at a frequency consistent with the U-shape features. The evolution of the cutoff frequency could then arise purely from the variation in elevation, with the irregularity scale and altitude remaining essentially constant. At lower elevations, the path length will be greater, meaning that the frequency will have to increase proportionally to correspond to a constant Fresnel scale.

To test the viability of this idea, the phase screen model developed in Chapter 5 was employed to represent the effect of such a spectrum of irregularities. A simple irregularity spectrum was defined and used to generate a phase screen realisation. This phase screen could then be used to generate the synthetic dynamic spectrum. This would not be able to replicate the time variation of the cutoff frequency due to the model assuming a constant path length, but it could produce a cutoff at some fixed frequency, which may vary depending on the path length and irregularity parameters. For more details of the modelling approach, see Appendix G.

Examples of dynamic spectra generated in this way are shown in Figure 8.6. The outer scale of 4 km used here is set to provide a corresponding frequency of approximately 45 MHz at an elevation of 10° (see Table 8.1). However, the dynamic spectra do not show any sign of a frequency cutoff, with scintillation present throughout the frequency band. There is an apparent transition between weak scattering behaviour at the higher frequencies and strong scatter at the lowest, indicated by the lack of coherence across frequency of the scintillation at the lowest frequencies (see Section 3.3.1). The

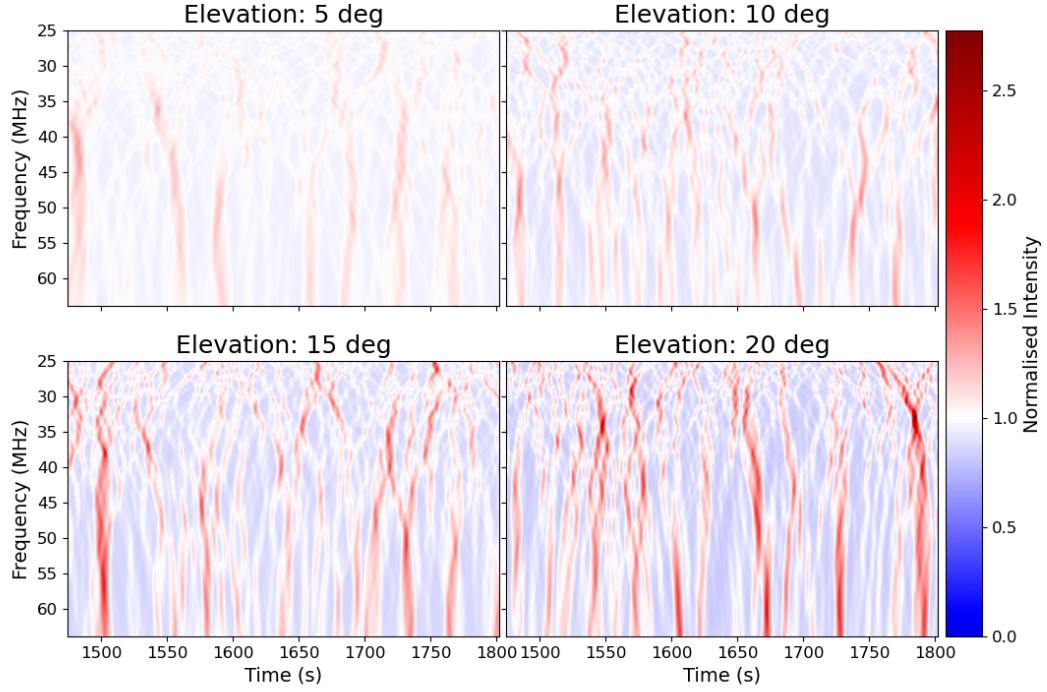


Figure 8.6. Example synthetic dynamic spectra generated with parameters: altitude 300 km, outer scale 4 km, inner scale 20 m, spectral index $\frac{-11}{3}$, velocity 100 m s^{-1} and root mean square TEC fluctuations of 10 mTECu. Each panel corresponds to a different elevation as indicated. The absolute value of the time stamps have no specific meaning, with 0s simply corresponding to the start of the modelled time period. The synthetic dynamic spectra are taken from the midpoint of the time window.

increase in magnitude of the intensity variations with increasing ~~intensity~~ is due to the increasing sensitivity of LOFAR as discussed in Section 5.2.1. Varying the irregularity parameters did not provide any examples of a cutoff in scintillation (see Appendix G).

The variation of the simulated scintillation characteristics with elevation is shown in Figure 8.7. The S_4 index is calculated for each frequency at elevations ranging from 5° to 20° over the time window shown in Figure 8.6. The general increase in S_4 with increasing elevation is due to the increasing sensitivity of LOFAR to the compact source as discussed in Section 5.2.1. The variation of S_4 with frequency is generally minor, and the slight decrease in S_4 towards the lowest frequencies is neither significant




Figure 8.7. The S_4 index as a function of frequency for various elevations calculated from the synthetic dynamic spectra. The x indicates the frequency with a Fresnel scale equal to the outer scale of the irregularities for each elevation. The parameters used are the same as in Figure 8.6.

enough nor sharp enough to replicate the behaviour of the observations. There is no clear relationship to the outer scale of the irregularities, which has a corresponding frequency at each elevation denoted by a cross.

Given that the phase screen approach has been unable to replicate the observed scintillation behaviour, it is useful to enumerate the assumptions made in applying the model, to identify possible reasons that it is unable to represent these observations. The model assumes that all frequencies ‘see’ the same irregularities, meaning they take the same path through the ionosphere. Especially at low elevations, refraction will be signif-

icant and different frequencies may take noticeably different paths, which is addressed further in the following section. Also, the model assumes that the plasma frequency is always much less than the radio frequencies considered, in order to approximate the phase accumulated using equation (3.6). While this is likely a reasonable assumption for the higher frequencies, the lower frequencies observed by LOFAR will likely be close enough to the plasma frequency (up to $\sim 5 - 10$ MHz at the F2 peak) that this approximation will start to break down. Finally, the model collapses an ionosphere of finite thickness into an infinitesimally thin phase screen. Especially at low elevations, the effective thickness of the ionosphere will be significant, meaning that the irregularity population can in principle be spread over a range of distances from the observer (although the dominant irregularities may be localised within the ionosphere and not distributed throughout).

8.2.2 Large Scale Refraction




Figure 8.8. A simplified representation of the possible explanation for the U-shape features. Dashed arrows represent the path taken by radio waves at a higher frequency (above the cutoff) while solid arrows represent the path taken by waves at a lower frequency (below the cutoff).

At low elevations, the vertical gradients in electron density cause significant refrac-

tion of the radio signals from the direct path. This effect will naturally be strongest for the lowest frequencies (see Section 3.3.3), meaning that the different frequencies will experience different ionospheric conditions as they propagate. For the observing geometry of the U-shapes, the result will be that the lower frequencies will generally pass through the ionosphere further north than the higher frequencies. If the irregularities responsible for the scintillation exist in a restricted range of latitudes and altitudes, this may lead to the lower frequencies simply not encountering them and hence experiencing no scintillation, a scenario illustrated in simplified form in Figure 8.8.

As the source rises in elevation, the frequency cutoff would therefore be expected to reduce. This is due to both the movement of all raypaths further South within the ionosphere and the reduction in separation between frequencies caused by the refractive effects. This time/elevation variation is therefore qualitatively consistent with the observed variation of cutoff frequency in the U-shape features. The variation in cutoff frequency between different events would in this case be explained by either a variation in the background ionosphere changing the refraction experienced by the different frequencies, a change in the latitudinal and/or altitudinal extent of the irregularities, or a combination of both of these effects.

The viability of this as an explanation for the observed features is dependent on two unknown quantities: the magnitude of the difference in paths taken by the different frequencies, and the spatial extent of the irregularities. Addressing either of these in any detail was beyond the scope of this work, but a brief summary of possible approaches is warranted here. For addressing the question of the path taken through the ionosphere by the various frequencies, a raytracing simulation (see Section 3.2.1) would be the natural solution. This would require a model of the ionosphere in at least two dimensions along the great circle path from LOFAR, and in practice to capture the variations with elevation (and associated variations in azimuth) this would need to

be three-dimensional. The extent of the irregularities responsible could be investigated directly using techniques such as GNSS TEC and/or scintillation measurements and in-situ measurements from satellites. Alternatively, it could be addressed indirectly by considering variations in the large-scale structure of the high-latitude ionosphere such as the position of the auroral oval, which can be derived from magnetometer measurements (e.g. Johnsen, 2013), satellite in-situ measurements (e.g. Xiong et al., 2014) and satellite imaging (e.g. Baker et al., 2000). In the future, the EISCAT_3D radar system in Northern Fennoscandia (McCrea et al., 2015, not yet operational at time of writing) would also be a promising candidate to investigate the structures responsible given its high spatial resolution and coverage, although this would depend on the structures being within its field of view.

8.3 Case Studies: 1st April 2022 and 21st January 2023

As well as the statistical conclusions that can be drawn from the event timings across many observations, it is useful to examine the scintillation of individual events in detail. This makes it possible to estimate properties such as characteristic timescales of the scintillation, velocity and any anisotropy of the structures responsible. It also permits the testing of possible explanations for the scintillation cutoff by comparing its variation with time to that expected for a given hypothesis.

The observations examined here were taken on 1st April 2022 from 19:31-23:59 UT (observation ID L856292) and 21st January 2023 from 00:46-03:35 UT (observation ID L880978). Both contain only the ascending phase of a U-shape, as observations containing full events were too large to practically manage. Both events occurred under a Kp of 2+, meaning conditions were relatively favourable for U-shape occurrence.




Figure 8.9. The dynamic spectrum of Cygnus A (top panel) viewed from station CS002 (52.915°N , 6.870°E , 49.35 m above sea level) for the first event (1st April 2022). Intensity range is limited to $\pm 20\%$ to highlight scintillation, maximum variations are $\sim \pm 50\%$. Elevation (second panel) and azimuth (third panel) of Cygnus A viewed from CS002, and the SME index (fourth panel) are also shown.

The dynamic spectrum of Cygnus A from station CS002 (52.915°N , 6.870°E , 49.35 m above sea level) is shown in Figure 8.9 for the first event, and Figure 8.10 for the second event. In the first event, the ascent occurs in the earlier part of the observation, and scintillation continues across the whole observing band for the remainder of the observation whereas in the second event the scintillation intensity decreases noticeably after the ascent is complete. Based on the quicklook plots, the ascent was estimated to occur from 20:05-21:50 UT in the first event and 01:10-02:25 UT in the second event. In both cases, the estimated start time occurs well before the scintillation is clearly developed, with only very minor scintillation present before approximately 21:00 UT in the first event and 01:40 UT in the second. This indicates the limitations of manually




Figure 8.10. The dynamic spectrum of Cygnus A (top panel) viewed from station CS002 (52.915°N , 6.870°E , 49.35 m above sea level) for the second event (21st January 2023). Intensity range is limited to $\pm 20\%$ to highlight scintillation, maximum variations are $\sim \pm 30\%$. Elevation (second panel) and azimuth (third panel) of Cygnus A viewed from CS002, and the SME index (fourth panel) are also shown.

estimating the timings from the quicklook plots as there is a clear bias when comparing to the full resolution data, especially for the inner edges for which the elevation is especially low, leading to low sensitivity and low scintillation amplitudes.

The elevation and azimuth of Cygnus A are also shown in Figures 8.9 and 8.10, along with the SME index. This indicates an enhancement of the SME index around 20:30 UT which then remained elevated for the remainder of the observation for the first event, and generally lower values of SME with an enhancement occurring around 02:30 UT in the second event. The scintillation intensity also increases over the course of the first observation. However, this may be a result of the increasing elevation of Cygnus A rather than any change in the ionospheric conditions. As discussed in Section 5.2.1,

the sensitivity of LOFAR to the compact source decreases with decreasing elevation, meaning that scintillations will appear to be lower in amplitude at lower elevations.

8.3.1 Timescales of the Scintillation




Figure 8.11. The 2D periodograms for four consecutive 30 minute windows in the first event (1st April 2022) spanning the ascent. Power units are arbitrary but consistent across all panels (note different colour scales in each panel, each is capped at the 95th percentile of power in that panel).

To quantify the timescale(s) of the scintillation, periodograms covering four consecutive 30 minute windows are shown in Figure 8.11 for the first event and Figure 8.12 for the second event. In each case these cover roughly the full duration of the ascent. Initially, scintillation power is clearly confined to the higher radio frequencies, whereas by the final window after the ascent is finished it spans essentially the full bandwidth, as would be expected from the dynamic spectrum. The scintillation power

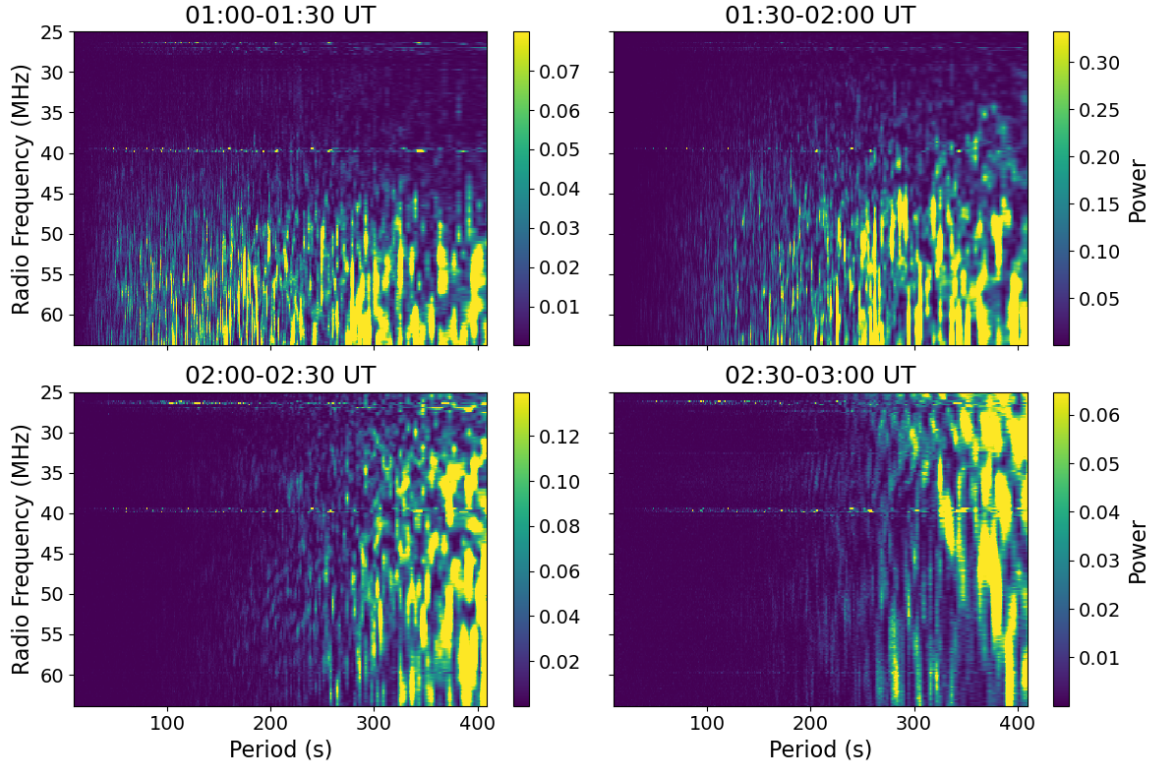


Figure 8.12. The 2D periodograms for four consecutive 30 minute windows in the second event (21st January 2023) spanning the ascent. Power units are arbitrary but consistent across all panels (note the different colour scales in each panel, each is capped at the 95th percentile of power in that panel).

also increases significantly with time for the first event, as shown by the difference in the colour scales, but this is likely largely influenced by the change in elevation as discussed above. Generally, the dominant timescales are ~ 100 s or more in the first event, while in the later windows the dominant scintillation power shifts to above ~ 200 s. For the second event, there is initially some power at shorter periods < 100 s (see top left panel of Figure 8.12) but the power at longer periods becomes increasingly dominant in each successive time window. It is also interesting to note that in the second event the scintillation power does not increase consistently with time but rather maximises between 01:30-02:00 UT and then begins to decrease as shown by the differences in the colour scales of Figure 8.12.

There is some evidence that scintillation at longer periods extends further to the lower radio frequencies in both events, clearest in the top right panel of each figure. These longer timescales are likely indicative of scintillation from larger structures, and may be present at lower radio frequencies due to the larger Fresnel scale at these frequencies, although a general increase in scintillation timescales across all radio frequencies could also reflect a reduction in the velocity at which the irregularities are moving across the line of sight. Also, a counterexample to this phenomenon is given by the bottom right panel of Figure 8.11, in which the power at the longest periods > 300 s is more pronounced at the higher radio frequencies.

8.3.2 Power Spectrum Analysis

To obtain more detailed information on the scintillation observed, the power spectrum of the scintillation (i.e. the Fourier transform of the intensity variations at a given frequency) is useful. The form of this spectrum can be theoretically related to the spectrum of irregularities responsible for the scintillation (e.g. Carrano and Rino, 2016), as well as properties such as the Fresnel frequency (see Section 3.3.1).

A generalised approach to characterising the power spectrum is provided by Braden (2022), using a three component piecewise fit with adaptive breakpoints. The three components represent the region below the Fresnel frequency, the region above the Fresnel frequency, and the noise dominated region respectively and are each linear in log-frequency log-power space. While Braden (2022) suggests allowing an arbitrary spectral index for the region below the Fresnel frequency and for the noise dominated region, it was found for the LOFAR data considered here that a spectral index of zero in these regions described all cases well and avoided certain instances of spurious fitting where the break points of the fit were matched to the incorrect features of the power spectrum (such as when the Fresnel frequency is particularly low and the low frequency

regime is therefore not well defined). This fit therefore has four free parameters: the Fresnel frequency (ν_F), the noise power level, the power at the Fresnel frequency, and the spectral slope (p). This is slightly different to the way the fit is defined by Braden (2022), where the equivalent four parameters are the power and frequency at the two breakpoints respectively. This change is made to enable the imposition of a limitation that the spectral slope must be negative, which again was found to limit the possibility of spurious inaccurate fits.




Figure 8.13. Example intensity power spectra for a range of radio frequencies over a single time window (21:10:59 - 21:15:59 UT) for the first event (1st April 2022). Solid grey lines show the calculated power spectra and red dashed lines show the fit. Note all panels use the same scale for the vertical axis (showing power in arbitrary units) for ease of comparison. Each panel is annotated with the estimated values of Fresnel frequency ν_F and spectral slope p .

An example of the results of this fitting procedure for a single five minute window from the first event at various radio frequencies is shown in Figure 8.13. The frequencies shown were chosen to evenly sample the observing band, and the time window was

chosen as it lies within the middle of the U-shape ascent phase (see Figure 8.9). In general, it is clear that the simple piecewise fit proposed by Braden (2022) is able to replicate the form of the power spectra calculated from the LOFAR data. The spectra become noise dominated for frequencies greater than ~ 1 Hz, which was typical for all times and frequencies within the U-shape scintillation events studied. The decrease in scintillation power at the lower radio frequencies is clear from the reduction in spectral power at the lower radio frequencies, while the noise level remains essentially constant across all radio frequencies. At the intermediate radio frequencies (top left, bottom right and bottom middle panels) there is an enhancement in the Fresnel frequency ν_F and an associated steepening of the spectral slope p , while at the lowest radio frequency (bottom right panel) the Fresnel frequency is comparable to that at the highest radio frequencies and the spectral slope is markedly shallower.

The main benefit of this fitting procedure is obtaining the Fresnel frequency and spectral slope, both of which provide information on the irregularities responsible for the scintillation. Another benefit is quantifying the noise power level and the frequency range in which it dominates. This can be used to define a cutoff frequency for a low-pass filter to remove the noise before calculating S_4 , providing a more accurate measurement of the scintillation. This is particularly useful across extended LOFAR observations, because, once the noise contribution is removed, the known variation of sensitivity with elevation can be used to remove this instrumental effect and reveal any residual variations in scintillation intensity due to changes in the ionospheric conditions and/or observing geometry.

It is important to note at this point that the assumption that the spectral break point and spectral slope directly correspond to the Fresnel frequency and slope of the irregularity spectrum is reliant on dealing with weak scintillation (see Section 3.3.1). In the case of strong scintillation, the intensity power spectrum can deviate significantly

from the form described by this simple piecewise fit (e.g. see examples in Carrano and Rino, 2016). If the whole of one of the spectra shown in Carrano and Rino (2016) were observed, it would be clear that the weak scattering approximation was inapplicable and the piecewise fit would fail. However, in reality the range of frequencies in the intensity power spectrum is limited. At the high frequencies the noise floor of the receiver system dominates the scintillation power, while at the lower frequencies the statistical non-stationarity of the irregularity population becomes an issue (i.e. over a sufficiently long time window it is no longer possible to assume that the irregularities being observed are all samples from the same distribution, see discussion in: Carrano and Rino, 2016; Forte et al., 2022). In reality therefore, only a small part of the theoretical spectra presented by Carrano and Rino (2016) can be sampled, and therefore a spectrum arising from strong scintillation may appear consistent with weak scattering. This possibility and its implications will be discussed later in this section, although for simplicity the term ‘Fresnel frequency’ will continue to be used to refer to the fitted spectral break frequency.

In a simple idealised case, in which scintillation at all frequencies arises from the same irregularity population, the dependence of the Fresnel frequency on the radio frequency is simple. The Fresnel frequency ν_F is defined as

$$\nu_F \propto \frac{v}{D_F} = v \sqrt{\frac{f}{2cL}}, \quad (8.1)$$

where v is the velocity of the irregularities relative to the line of sight, f is the radio frequency, c is the speed of light and L is the line of sight distance from the irregularities to the observer. Therefore, if all frequencies are affected by the same irregularities, both v and L will be constant and the only variable will be f , giving a Fresnel frequency proportional to the square root of the radio frequency. Deviations from this may indicate that the scintillation at different frequencies is caused by different irregularity

populations, consistent with the explanation suggested in Section 8.2.2.




Figure 8.14. The estimated Fresnel frequency vs. time and radio frequency for the first event (1st April 2022). The red curve shows the frequency corresponding to a Fresnel scale of 4.2 km for an altitude of 300 km.

For this fitting, each observation was divided into 5 minute segments, and for each of these at each radio frequency the power spectrum of intensity variations was calculated and parameters estimated. The Fresnel frequencies derived from this are shown in Figures 8.14 and 8.15 for the first and second event respectively. In both cases, it shows that the simple prediction of Fresnel frequency being proportional to the square root of radio frequency based on equation (8.1) does not correspond to the observations. Instead, there is an enhancement of the Fresnel frequency around the cutoff radio frequency, clearest in the first event but also apparent in Figure 8.15 between approximately 01:30-02:00 UT (the red curve gives a rough approximation of the cutoff frequency for reference, note that the approximation is more accurate for the first




Figure 8.15. The estimated Fresnel frequency vs. time and radio frequency for the second event (21st January 2023). The red curve shows the frequency corresponding to a Fresnel scale of 4.2 km for an altitude of 300 km.

event than the second as shown in Figures 8.18 and 8.19). Referring back to equation (8.1), this enhancement could be explained by either an increase in the velocity v or a decrease in the line of sight distance L , or a combination of the two. The particularly high estimated Fresnel frequencies before 01:30 UT in Figure 8.15 are dubious, given the almost negligible scintillation shown in Figure 8.10 for this time period, and so should be treated with caution.

This behaviour of the estimated Fresnel frequencies provides possible evidence that the observed frequency cutoff of the scintillation is due to different frequencies being affected by different irregularity populations, consistent with the proposed explanation in Section 8.2.2. However, there are some caveats to this that must be considered. This conclusion relies on the assumption that the spectral break detected by the fitting

routine corresponds to the Fresnel frequency, which is dependent on the scattering being weak. As moving to lower frequencies can lead to a transition from weak to strong scattering (see Section 3.3.1) it is possible that the apparently anomalous behaviour around the cutoff frequency in Figures 8.14 and 8.15 reflects this transition. However, it is difficult to explain why a transition from weak to strong scattering would lead to a decrease in the scintillation below the transition frequency as is observed in this case. Another possible confounding factor is that if the spectrum of irregularities is not a single power law but contains a spectral break of its own, for example corresponding to a preferred scale of the instability mechanism generating the irregularities, this spectral break can also be present in the intensity power spectrum (Carrano and Rino, 2016) and could be confused for the Fresnel frequency. If this were the case, however, the estimated Fresnel frequency would be approximately constant with radio frequency as the spectral break frequency would be determined purely by the break scale of the irregularity spectrum and the irregularity drift velocity.

The distribution of estimated spectral slope values for the first and second event are shown in Figures 8.16 and 8.17 respectively. In both cases the spectral slopes cluster around $-8/3$, albeit with quite substantial variation, as shown in the left panels. This value of $-8/3$ is the intensity spectral slope corresponding to an irregularity spectrum with a slope of $-5/3$ (i.e. Kolmogorov turbulence, assuming weak scatter see Section 3.3.1) so it is somewhat unsurprising that it appears here. The distribution of estimated spectral slopes with time and radio frequency is shown in the right panel of each figure. It is clear that the shallowest (i.e. closest to zero) values of the spectral slope come from regions of time and radio frequencies for which there was little to no scintillation (see Figures 8.9 and 8.10, also note that the colour scales in Figures 8.16 and 8.17 are capped such that the colour scale saturates in these regions). This indicates that the spread of the histogram in the left panels towards shallower slopes, including the




Figure 8.16. The distribution of estimated spectral slopes for the first event (1st April 2022), shown as a histogram in the left panel and as a function of time and radio frequency in the right panel. The colour scale in the right panel has been capped to highlight variations within the regions which experience scintillation and therefore saturates in others (e.g. the upper left corner). The red curve shows the frequency corresponding to a Fresnel scale of 4.2 km for an altitude of 300 km.

secondary peak around -0.7 in Figure 8.17, are likely more a result of fitting to noise dominated spectra than genuine shallowing of the spectra. Outside of these regions, there is an apparent slight tendency for the spectra to steepen with time throughout the observations, although there is a great deal of variability around this trend.

8.3.3 Scintillation Power

Given that significant scintillation occurs on timescales of $\sim 100 - 200$ s or longer as shown in Section 8.3.1, calculating S_4 over a 60 s timescale as is common for GNSS (e.g. Pi et al., 2017; Spogli et al., 2013) would not be appropriate to fully capture the scintillation power. Instead, for this study a timescale of 300 s is used in calculating S_4 to account for this. This provides a more accurate measurement of total ionospheric

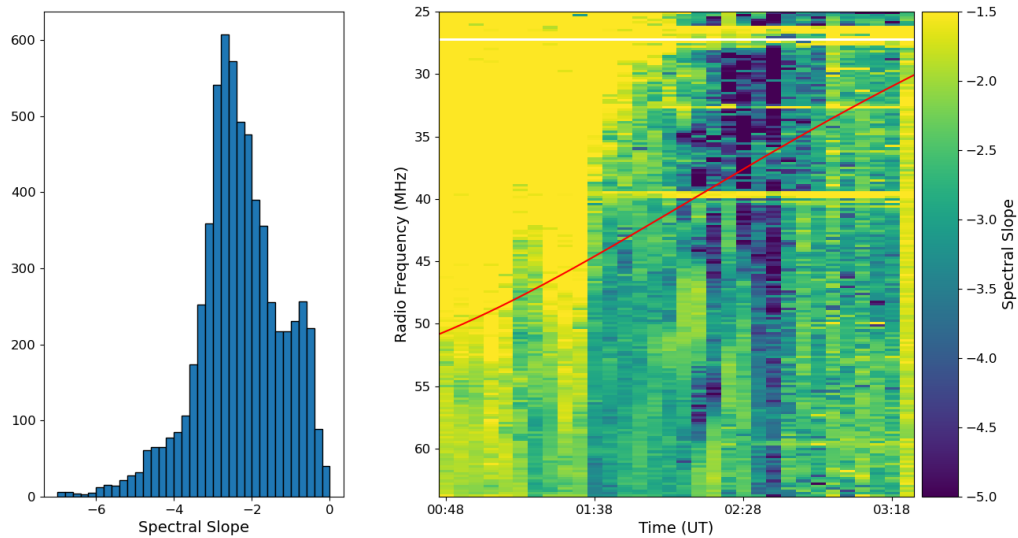


Figure 8.17. The distribution of estimated spectral slopes for the second event (21st January 2023), shown as a histogram in the left panel and as a function of time and radio frequency in the right panel. The colour scale in the right panel has been capped to highlight variations within the regions which experience scintillation and therefore saturates in others (e.g. the upper left corner). The red curve shows the frequency corresponding to a Fresnel scale of 4.2 km for an altitude of 300 km.

scintillation power, but naturally reduces the time resolution of the measurement. The difficulties of defining an appropriate timescale for calculating S_4 for LOFAR data are explored in more detail in Wood et al. (2024). Calculating S_4 in this case is useful because it allows more accurate identification of the cutoff frequency than is possible from the dynamic spectrum directly. This can then be used to compare to the expected variation corresponding to possible physical explanations for the frequency cutoff.

Before calculating the S_4 index, the data was low-pass filtered with the cutoff frequency set based on the fits described in Section 8.3.2 (predominantly ~ 1 Hz for the first event and ~ 0.2 Hz for the second event). This removes instrumental noise from the data meaning that S_4 will provide a more accurate measure of the true amount of ionospheric scintillation present. Removing the instrumental noise also enables the S_4 index to be normalised to remove the theoretical elevation dependent sensitivity

variations. This is done by assuming that the background sky temperature emission is unaffected by the ionosphere and that all the observed variability is variability in the compact source (Cygnus A) (note this is the inverse of the correction applied to the phase screen modelling in Section 5.2.1 to match model output to the data). Given the wide frequency range covered by the U-shape features, it is also necessary to account for the frequency dependence of the intensity of Cygnus A and of the SEFD. The spectrum of Cygnus A is taken from McKean et al. (2016) while the SEFD is manually estimated from Figure 22 of van Haarlem et al. (2013). This estimate is

$$SEFD(f) = \begin{cases} 27 \text{ kJy} + 8 \text{ kJy} \frac{(f-60 \text{ MHz})^2}{(30 \text{ MHz})^2} & \text{if } f \leq 60 \text{ MHz}, \\ 27 \text{ kJy} + 33 \text{ kJy} \frac{(f-60 \text{ MHz})^2}{(20 \text{ MHz})^2} & \text{otherwise.} \end{cases} \quad (8.2)$$

where f is the radio frequency. The elevation corrected S_4 (here denoted as S'_4) is then

$$S'_4 = \frac{SEFD(f) + I_{Cyg}(f) \sin(\theta)}{I_{Cyg}(f) \sin(\theta)} S_4, \quad (8.3)$$

where $I_{Cyg}(f)$ is the intensity spectrum of Cygnus A.

The S_4 spectra are shown in Figures 8.18 and 8.19 for the first and second events respectively. The broadband S_4 enhancement around 21:45 UT in Figure 8.18 is due to a data gap and not genuine scintillation. The frequency cutoff of the scintillation and its variation with time is much more distinct here than in the dynamic spectra. Compared to the previous estimated duration of 20:05-21:50 UT from the quicklook plot for the first event, a more detailed inspection of the S_4 spectrum in Figure 8.18 is broadly consistent with this, at least to within the estimated ~ 5 min precision of the estimated timings from the quicklook plots. For the second event, the estimated timings were 01:10-02:25 UT, whereas the ascending slope in Figure 8.19 is not apparent until closer to 01:30-01:40 UT. This highlights the point that the estimated timings based




Figure 8.18. The S_4 index for all radio frequencies calculated from the dynamic spectrum for the first event (1st April 2022) shown in Figure 8.9. Also shown (red line) is the variation in the frequency corresponding to a Fresnel scale of 4.2 km at an altitude of 300 km. Noise has been filtered based on the fits described in Section 8.3.2 and sensitivity variations with elevation and the constant background contribution to the intensity have been removed according to equation (8.3).

solely on the quicklook plots are inaccurate for individual cases and only useful for a coarse analysis.

The red line in Figures 8.18 and 8.19 shows the frequency corresponding to a Fresnel scale of 4.2 km at an altitude of 300 km (note this is also plotted in Figures 8.14, 8.15, 8.16 and 8.17). This is provided as a quantitative means of testing the explanation proposed in Section 8.2.1, as well as giving a direct comparison of the two events in terms of observing geometry. The parameters used were selected to approximately match the observed variations in these two cases, but a wide range of altitudes and corresponding Fresnel scales could be used to provide a similar match. The variation




Figure 8.19. The S_4 index for all radio frequencies calculated from the dynamic spectrum for the second event (21st January 2023) shown in Figure 8.10. Also shown (red line) is the variation in the frequency corresponding to a Fresnel scale of 4.2 km at an altitude of 300 km. Noise has been filtered based on the fits described in Section 8.3.2 and sensitivity variations with elevation and the constant background contribution to the intensity have been removed according to equation (8.3).

shown by the red line is fairly similar to the observed variation in cutoff frequency, particularly in Figure 8.18 whereas in Figure 8.19 the cutoff frequency varies more rapidly with time / observing geometry than the simple relationship predicts.

Unfortunately, there is no comparable simple predicted relationship between elevation and cutoff frequency for the large scale refraction explanation proposed in Section 8.2.2. Instead, as discussed previously, this would require information both on the large scale electron density distribution responsible for the refraction of the radio waves and the spatial extent of the irregularities. With the data available here, all that can be said is that the observed trend of cutoff frequency decreasing with increasing elevation

is qualitatively consistent with the explanation proposed in Section 8.2.2.

8.3.4 Cross-Correlation Analysis

In order to further characterise the scintillation, the full network of Dutch LOFAR stations is useful. The cross-correlation between the signals observed at each pair of stations provides a great deal of information on both the spatial structure of the irregularities and their propagation. The correlation function can be characterised by fitting a quadratic in two spatial dimensions and time to the main peak (Grzesiak et al., 2022). Using the actual values of the cross-correlation function rather than just the time at which the cross-correlation on each baseline is maximised (as has been done before with LOFAR data: e.g. Fallows et al., 2020, Chapter 5 of this work) is preferable because it allows the estimated velocity to account for both spatial anisotropy and temporal decorrelation. In addition to this improved estimate of the velocity, it also provides a direct estimate of the anisotropy of the irregularities responsible for the scintillation. For a detailed description of the approach used here to analyse the cross-correlations, see Appendix H.

The orientation and magnitude of the anisotropies in the plane of sky are shown in the left and right panels of Figures 8.20 and 8.21 for the first and second events respectively. The angle within the plane of sky is defined such that an orientation of 90° corresponds to the purely horizontal direction within the plane of sky. In both observations, the anisotropy of the irregularities clearly tracks the orientation of the geomagnetic field projected into the plane of sky (shown by the red line in the left hand panel). This field orientation was calculated assuming an altitude of 300 km for the irregularities and neglecting any effect of refraction. For the observing geometries considered here, the line of sight is quasi-perpendicular to the geomagnetic field, with angles between them ranging from $\sim 74 - 86^\circ$. This means that there is little distortion

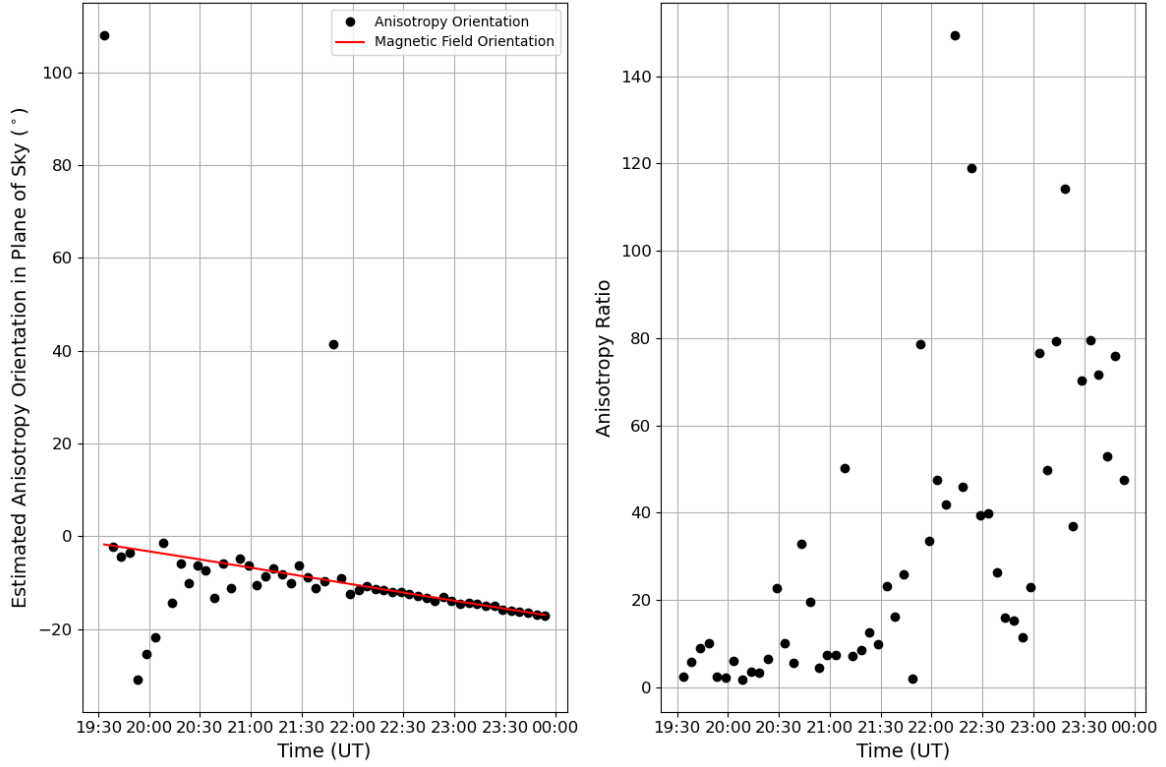


Figure 8.20. The estimated anisotropy parameters in the plane of sky for the first event (1st April 2022), using cross correlations at 61.1 MHz. The left panel shows the estimated anisotropy orientation in the plane of sky (black dots) and the projected orientation of the geomagnetic field at 300 km (red line). The right panel shows the anisotropy ratio (major axis / minor axis) as a function of time. The geomagnetic field is calculated using the International Geomagnetic Reference Field (IGRF: Alken et al., 2021).

in projecting the magnetic field into the plane of sky and field-aligned irregularities can be very easily identified. The anisotropy ratio (anisotropy major axis divided by minor axis) shown in the right hand panel shows a wide range of values and a great deal of variability from one time window to the next. The highest values also generally correspond to the most intense scintillation as indicated by Figures 8.18 and 8.19. This is indicative of the fact that the anisotropy is not well constrained by these observations, with the baseline foreshortening in the plane of sky leaving a maximum baseline along the magnetic field direction of ~ 10 km or less and the correlation peak clearly extending beyond this distance (see Appendix H for examples of this). As a result, the precise




Figure 8.21. The estimated anisotropy parameters in the plane of sky for the second event (21st January 2023), using cross correlations at 61.1 MHz. The left panel shows the estimated anisotropy orientation in the plane of sky (black dots) and the projected orientation of the geomagnetic field at 300 km (red line). The right panel shows the anisotropy ratio (major axis / minor axis) as a function of time. The geomagnetic field is calculated using the International Geomagnetic Reference Field (IGRF: Alken et al., 2021).

values of the anisotropy ratio are not reliable beyond as an indication that the anisotropy is significant, and indeed low values may at least in some cases reflect low scintillation intensity rather than quasi-isotropic irregularities.

The projected orientation of the geomagnetic field is relatively insensitive to changing the assumed altitude, with values between 100 – 600 km only changing the angle in Figure 8.20 by a few degrees. This means there is no way to draw any conclusions as to the likely altitude of the irregularities based on their alignment with the geomagnetic field, as has been done previously for other observations (e.g. Appendix F of this thesis; Hoogeveen and Jacobson, 1997). However, the fact that they are field-aligned seems to

be beyond doubt.

The strong anisotropy and insufficient baseline coverage to fully characterise the major axis mean that only the component of velocity perpendicular to the anisotropies can be reasonably estimated from the data available. The fitting process also outputs estimate of the parallel velocity, but these are generally unphysical ($\sim 1 - 100 \text{ km s}^{-1}$) and have large estimated uncertainties, and so are ignored for the purposes of analysis. Given the observing geometry and orientation of the anisotropy, this means LOFAR is effectively sensitive to quasi-zonal motion of the irregularities and insensitive to quasi-meridional motion.




Figure 8.22. The azimuth (left panel) and magnitude (right panel) of the velocity normal to the anisotropy converted into the horizontal plane for the first event (1st April 2022), using cross correlations at 61.1 MHz. The irregularities are assumed to be at an altitude of 300 km. The azimuth is defined clockwise from geographic North. Note that this does not include any correction for the motion of the ionospheric pierce point (~ 10 s of metres per second, see text for details).




Figure 8.23. The azimuth (left panel) and magnitude (right panel) of the velocity normal to the anisotropy converted into the horizontal plane for the second event (21st January 2023), using cross correlations at 61.1 MHz. The irregularities are assumed to be at an altitude of 300 km. The azimuth is defined clockwise from geographic North. Note that this does not include any correction for the motion of the ionospheric pierce point (~ 10 s of metres per second, see text for details).

Figures 8.22 and 8.23 show this component of the velocity expressed in terms of its azimuth and magnitude, after conversion from the plane of sky to the horizontal. In both cases, there is noticeable variability in the estimated direction early in the observation, when scintillation amplitudes are low and the ascending phase of the U-shape has not yet started. Later in both observations, the velocity settles to westward flow. The rotation of the anisotropy in Figures 8.20 and 8.21 is not present in the velocities as this is largely due to the rotation of the plane of sky with respect to the geomagnetic field and the geographic horizontal co-ordinates. There is an apparent rotation of the velocity in the opposite sense shown in Figure 8.22, but given that

LOFAR is only sensitive to one component of the velocity, it is possible that this does not reflect any real rotation of the full physical velocity vector. In combination with the decreasing velocity towards the end of the observation, it may instead indicate a constant quasi-westward velocity to which LOFAR's observing geometry becomes increasingly less sensitive with time as the apparent anisotropy rotates.

It is important to note here that these are the velocities of the irregularities relative to the line of sight, which is itself moving relative to the ionosphere. This can in principle be corrected for, but only if the location of the irregularities along the line of sight is known, which it is not in this case. However, it is still possible to obtain an estimate of the impact by assuming a location for the irregularities, which here is taken to be at an altitude of 300 km. As this velocity is purely a function of the observing geometry, it is consistent between both cases considered here. The velocity of the line of sight relative to the ionosphere in the plane of sky decreases from $\sim 90 \text{ m s}^{-1}$ around the start of each observation to $\sim 40 \text{ m s}^{-1}$ around the end, always approximately in the positive x -direction (i.e. clockwise). If the velocities shown in Figures 8.22 and 8.23 are corrected to account for this, the effect is to reduce the magnitude of the observed westward velocities. In the second event, this reduction for an altitude of 300 km is enough that the velocity reduces to $< 50 \text{ m s}^{-1}$ after 01:45 UT and occasionally switches direction from westward to eastward. This suggests that the apparent westward flow in Figure 8.23 may be a result of the motion of the line of sight, with the motion of the irregularities themselves being negligible in comparison. However, without an independent measurement indicating where the irregularities occur along the line of sight, it is impossible to say for certain exactly how significant the line of sight motion is.

In both cases, the magnitudes of the velocity show significant variability on timescales of tens of minutes, only stabilising late in the first event after the U-shape ascent has fin-

ished (Figure 8.22). The velocities in the first event range from $\sim 50 - 300 \text{ m s}^{-1}$ whereas for the second event they are generally slightly lower, ranging from $\sim 25 - 150 \text{ m s}^{-1}$.




Figure 8.24. The estimated plane of sky velocity as a function of radio frequency and time for the first event (1st April 2022). Also shown (red line) is the variation in the frequency corresponding to a Fresnel scale of 4.2 km at an altitude of 300 km. The channel at 39.6 MHz was removed as the results were dominated by noise and/or RFI.

The orientation of the anisotropies estimated at different frequencies is essentially constant for a given time, but there is some variation in estimated velocities. The estimated velocities as a function of time and frequency are shown in Figures 8.24 and 8.25 for the first and second event respectively. This shows that the velocities estimated for different radio frequencies at the same time are generally consistent, except in regions where there is negligible scintillation.

Comparing Figures 8.24 and 8.25 to Figures 8.14 and 8.15 there is no corresponding enhancement in the velocities near to the scintillation cutoff as there is in the Fresnel




Figure 8.25. The estimated plane of sky velocity as a function of radio frequency and time for the second event (21st January 2023). Also shown (red line) is the variation in the frequency corresponding to a Fresnel scale of 4.2 km at an altitude of 300 km. The channel at 39.6 MHz was removed as the results were dominated by noise and/or RFI.

frequency. Given equation (8.1) this would suggest that the enhancement in Fresnel frequency must correspond to a reduction in line of sight distance L to the irregularities. This seems unlikely to be the case as the difference would need to be extremely large (approximately a factor of four for an increase in ν_F by a factor of two) and there is no clear physical explanation for such a difference to occur. Instead, it suggests that the spectral break frequency being identified by the fitting procedure is not in fact the Fresnel frequency (at least for some radio frequencies).

This issue of whether the spectral break frequency can be identified with the Fresnel frequency therefore needs to be considered in any implementation of the fitting procedure proposed by Braden (2022) to ensure that the results obtained can be interpreted

accurately. Carrano and Rino (2016) do propose a spectral fitting method that accounts for these issues, but it requires complex and computationally expensive integrals and was not attempted here.

8.3.5 Simultaneous Measurements from the Swarm Satellites

To provide context for the observations made by LOFAR, in-situ satellite measurements can be used to compare direct observations of the plasma irregularities to the scintillation observed by LOFAR. The Swarm satellites (Friis-Christensen et al., 2008) are a constellation of three satellites flying in near polar orbits which drift slowly in local solar time. During the LOFAR observations considered here, two of the satellites (A and C) were flying at approximately 480 km in altitude separated by approximately one degree in longitude and half a degree in latitude, while Swarm B was flying at a slightly higher altitude of roughly 530 km. They each carry the same set of instruments, which includes a Langmuir probe providing measurements of plasma density and temperature with a cadence of 2 Hz (corresponding to a spatial scale of ~ 4 km) (Knudsen et al., 2017).

Using the measurements made by the Langmuir probe, GNSS receiver and magnetic field instruments, a range of measures of ionospheric variability have been defined and are provided in the IPIR (Ionospheric Plasma IRregularities) data product (Jin et al., 2022b) at a time resolution of 1 s. These include the absolute electron density residual after median filtering (ΔN_e , note that this is different to the definition given in equation (6) of Jin et al. (2022b) which suggests it is provided as a signed residual (Jin, personal communication)), estimated electron density gradients across 20, 50 and 100 km scales and the rate of density index (RODI: the standard deviation over some time window (10, 20 or 40 s) of the estimated rate of change of density based on the differences in density between adjacent measurement points). These measures enable

ionospheric irregularities to be investigated over a range of length scales which are comparable to or somewhat greater than the Fresnel scale at LOFAR frequencies.

There are some limitations to the comparison between LOFAR and the Swarm satellites, most obviously that it requires a satellite to pass over LOFAR during an observation. For the case studies considered here, the first event had no suitable overpasses from Swarm whereas the second event had an overpass from Swarm A and C at around 01:30 UT. Another limitation is that the orbital configuration of Swarm means that they are primarily sensitive to irregularities in the latitudinal direction (i.e. along their orbit) whereas LOFAR during U-shape events is primarily sensitive to irregularities in the longitudinal direction (more precisely: perpendicular to the line of sight). This means that it is not possible to make a one-to-one comparison of the length scales of irregularities detected by the two techniques, as the irregularity distribution is likely to be anisotropic with respect to geographic and/or geomagnetic East and North directions. Comparison of data from Swarm A and C can provide some information on the longitudinal structure of any irregularities, but this is limited by the fact that the separation between the satellites is far larger than the scale of irregularities that LOFAR is sensitive to and that the satellites are not at the same latitude (geomagnetic or geographic) at the same time, creating issues distinguishing temporal and spatial variations.

To provide context for the Swarm observations and how they may relate to the occurrence of U-shape features in the LOFAR data, an observation which did not contain a U-shape and for which a Swarm overpass occurred is used as comparison. This observation was taken on the 7th January 2022 (observation ID L842442) under quiet geomagnetic conditions (Kp 0) and no scintillation was detected by LOFAR. The Swarm overpass in this case occurred at around 23:30 UT, approximately the time at which Cygnus A is due North as viewed from the LOFAR core and hence expected to be the

midpoint of any U-shape if one had been present. The comparison of Swarm data between these two cases may give an indication of why U-shapes occur in some cases and not others, such as differences in latitudinal extent and amplitude of any irregularities. However, any definite statement linking the Swarm observations to the occurrence or absence of U-shape features would require more than two examples to be considered.

Figures 8.26 and 8.27 show the electron density and the magnitude of small scale variations in the electron density as a function of latitude for a single pass of Swarm C over Europe. The two passes occurred at slightly different times of day, but both were well within the time at which the LOFAR observing geometry was suitable for U-shape features to be observed. The first clear difference between the two observations is that the electron densities observed by Swarm were significantly higher when a U-shape was observed (Figure 8.26) than when no U-shape was observed (Figure 8.27), particularly at the highest latitudes (note the different vertical scales between the two plots). The elevated densities at high latitudes ($> 60^\circ\text{N}$) are also accompanied by significant irregularities in the electron density as quantified by δN_e . In the quiet case, the most extensive irregularities are in the mid-latitudes between roughly $45\text{--}60^\circ\text{N}$ in the equatorward wall of the mid-latitude trough. These irregularities are lower in magnitude than those in the auroral region in Figure 8.26, and also would need to extend well into the bottomside ionosphere to altitudes of ~ 200 km or lower in order to intersect the LOFAR line of sight. In both cases, both Swarm A and C observed the same distribution of irregularities with latitude in the mid-latitude and auroral regions, suggesting that the structures observed are not narrowly confined in longitude.

Previous theoretical (e.g. Aarons, 1982) and observational (e.g. Jenner et al., 2020) studies have highlighted the importance of both significant plasma density gradients (leading to high irregularity growth rates) and sufficiently high background plasma density in order for scintillation to occur. Jenner et al. (2020) report cases in which no




Figure 8.26. The measured electron density (black solid line) and the magnitude of the median filtered electron density residuals (delta_Ne, red dots, multiplied by 10 for ease of visibility) as a function of latitude from Swarm C for the observation on 21st January 2023 (the second event discussed in the previous sections). Also shown (black dashed line) is the electron density measured by Swarm A, shifted upwards by 200,000 $\text{e}^- \text{cm}^{-2}$ for clarity. delta_Ne for Swarm A is not shown for clarity but is consistent with that derived from Swarm C. The background colours are determined by the IPIR ionospheric region flags from Swarm C and the corresponding region is labelled at the top of the plot. The median filter for delta_Ne uses a 10s window, corresponding to a spatial scale of 75 km. Both satellites passed from south to north, with Swarm C at an average longitude of 6.25°E and Swarm A at 7.16°E between 40°N and 70°N. The data shown covers 01:21-01:42 UT.




Figure 8.27. The measured electron density (black solid line) and the magnitude of the median filtered electron density residuals (delta_Ne, red dots, multiplied by 10 for ease of visibility) as a function of latitude from Swarm C for the observation on 7th January 2022 (the observation for which no U-shape was observed). Also shown (black dashed line) is the electron density measured by Swarm A, shifted upwards by 100,000 $\text{e}^- \text{cm}^{-2}$ for clarity. delta_Ne for Swarm A is not shown for clarity but is consistent with that derived from Swarm C. The background colours are determined by the IPIR ionospheric region flags from Swarm C and the corresponding region is labelled at the top of the plot. The median filter for delta_Ne uses a 10s window, corresponding to a spatial scale of 75 km. Both satellites passed from north to south, with Swarm C at an average longitude of 0.18°E and Swarm A at 0.34°E between 40°N and 70°N. The data shown covers 23:18-23:36 UT.

scintillation was observed at the boundary of a polar hole, despite steep density gradients which should have been favourable for irregularity development, and attributed this to the extremely low densities within the polar hole. Naturally, from just two examples it is not possible to determine whether this condition also applies to these U-shape features, but the Swarm data would provide a possibility of investigating this if these comparisons were extended to all LOFAR observations of suitable observing geometry with a Swarm overpass.

8.4 Discussion

The U-shape scintillation features seem to be the first example to be identified of a consistent, recurring scintillation feature in LOFAR data. While the analysis presented here is necessarily limited by the impracticality of downloading and processing the raw LOFAR data from so many observations, nevertheless some reasonably confident conclusions can be drawn about the origin of these features from the quicklook data.

First and most obviously, these features are exclusively associated with Northward lines of sight at low elevations. While it is possible that low elevations alone are sufficient for the observed scintillation behaviour, the data currently available does not include any sources which reach low elevations in other directions. This is due to the sources which are suitably intense and compact for these type of observations, Cygnus A and Cassiopeia A, being at declinations which do not provide low elevation lines of sight in other directions.

The association with enhanced geomagnetic activity may support the idea that Northward lines of sight in particular are necessary for the observed phenomenon. Enhanced geomagnetic activity causes the auroral oval and mid-latitude trough to move equatorwards (Yang et al., 2015; Zhang and Paxton, 2008), which may bring them into

the field of view of LOFAR. If the irregularities responsible for the U-shape scintillation are within one of these regions, this equatorward expansion could explain the preference for enhanced geomagnetic activity. Another possible link is that the irregularities may be more favourably generated under enhanced geomagnetic activity specifically, rather than simply being moved into the field of view as a result of the activity.

The possible association with the auroral oval or the mid-latitude trough is also consistent with the observed preference for winter nighttime over summer daytime shown in Figure 8.4. The auroral oval is extended further equatorward at night than during the day (Zhang and Paxton, 2008), whereas the mid-latitude trough is typically only present at night (Krakowski et al., 2009; Yang et al., 2015). This explanation might suggest that the seasonal variation is simply a combination of the dependence on time of day and observing geometry and the season itself has no particular influence on the occurrence of U-shape scintillation, as was suggested based on Figure 8.4. However, the mid-latitude trough has been reported to be more pronounced in winter than summer (Krakowski et al., 2009), which may also factor into the preference for winter/nighttime conditions.

The detailed examination of two observations detailed in Section 8.3 provides some further information on the details of U-shape scintillation. While the results do not necessarily generalise to all U-shape events, it provides a first chance to test possible explanations. The clearest result is the strong field-aligned nature of the irregularities as shown in Figures 8.20 and 8.21. Another seemingly consistent feature is the enhancement of the Fresnel frequency at radio frequencies around the scintillation cutoff shown in Figures 8.14 and 8.15 (with the caveat mentioned above that the estimation may be identifying a spectral break not directly associated with the Fresnel frequency and hence the identified behaviour may not strictly be a Fresnel frequency enhancement). In the absence of a corresponding velocity enhancement as shown in Figures 8.24 and 8.25

however, the interpretation of this Fresnel frequency behaviour is currently unclear.

Unfortunately, due to the anisotropy and the effects of baseline foreshortening at low elevation, it is not possible to estimate the full velocity of the irregularities from the cross-correlations. This limits the ability to compare the estimated velocities from LOFAR to those measured using other instruments such as the SuperDARN radars (Chisham et al., 2007; Greenwald et al., 1995) in order to contextualise the results and possibly identify the likely location of the irregularities responsible for the scintillation more precisely than was possible here.

The observations provided by Swarm, while limited to only two nights, suggest that irregularities in the auroral region (or possibly the poleward wall of the mid-latitude trough) are more likely than irregularities in the trough minimum or equatorward wall to be responsible for the observed scintillation. Figure 8.26 shows negligible irregularities in the trough and substantial irregularities in the auroral region north of 60° in latitude when U-shaped scintillation was observed by LOFAR, while Figure 8.27 shows irregularities in the equatorward wall of the mid-latitude trough which fail to cause any such scintillation. If the mid-latitude trough was at higher latitudes during a given event, it is possible that this would bring irregularities in the equatorward wall into LOFAR's line of sight and hence give rise to scintillation, but this would only be expected to occur under quiet geomagnetic conditions and more often towards dusk and dawn than around midnight, which is inconsistent with the occurrence rates of U-shape scintillation shown in Figures 8.4 and 8.5.

If the observed scintillation is indeed associated with irregularities in the auroral ionosphere, a natural question becomes which instability process(es) (see Section 2.5.2) are responsible for the scintillation observed by LOFAR. In the absence of simultaneous measurements of all relevant parameters to the various instability growth rates, it is not possible to make any conclusive statements with the data available, but certain conclu-

sions can be drawn. Firstly, the Farley-Buneman instability is an unlikely candidate as it is primarily associated with the auroral electrojets flowing in the E-region. In order for the LOFAR line of sight to intersect with the auroral E-region ionosphere, the auroral oval would have to expand to $\sim 55 - 60^\circ$ geographic latitude which is highly unlikely except during severe geomagnetic storms and therefore cannot account for the observed scintillation which is frequently observed even under geomagnetically quiet conditions. The highly field-aligned nature of the irregularities also makes it less likely that the current convective instability is a dominant mechanism in generating the irregularities, as it requires a finite component of the irregularity wavevector along the magnetic field in order for the field-aligned current to contribute to charge separation.

Given that neither the TGI nor the Perkins instability have been reported to be significant in the auroral ionosphere, the last major candidates of those discussed in Section 2.5.2 are the GDI and KHI. Although Swarm provides measurements of the electric and magnetic fields and electron density gradients (on 100 km scale along track from the IPIR product) no attempt is made here to estimate irregularity growth rates to determine whether one or both of these processes is consistent with the observed irregularities. Such a comparison would be limited as the measurements of the irregularities (δN_e on length scales $< \sim 75$ km) are derived from the same measurement of electron density as the gradient measurements and may therefore correlate simply because they are in part sensitive to the same fluctuations. The density gradient is also limited by only providing the along track component of the gradient rather than the full gradient perpendicular to the field, although this could in principle be mitigated by combining measurements from Swarm A and C to estimate the 2D horizontal gradient. However, this analysis is not attempted here as it was beyond the scope of what was possible with the data and time available.

Future work would benefit from a more detailed statistical picture of the occurrence

of U-shape scintillation, beyond the relatively simple characterisation of whether or not an event occurred on a given day that is used here. This could be enabled by the work currently underway at ASTRON to pre-process the archive of ionospheric scintillation observations to produce a data product with RFI removed and downsampled in time to 1 s resolution. As part of this processing the S_4 index is also calculated and provided as a separate product. This downsampling reduces the data volume by approximately a factor of 100, and in addition to this the processed files are split by source and station unlike the raw data files used here. This processing would therefore make it feasible to download many multi-hour observations, addressing a limitation of the work presented in this chapter. Comparing the S_4 index across multiple events could allow the variability of the extent of the ‘U’ in both frequency and time to be assessed and for any possible relationship to driving factors such as Kp or season/time of day. The processing of the data would also allow the statistics to be extended to before October 2021. As well as improving the reliability of the existing estimates of occurrence rates, this may permit an investigation of possible solar cycle dependence of these features.

Another potential extension of this work would be to extend the comparisons with Swarm to all observations with suitable observing geometry for U-shaped features which have an overpass from one or more Swarm satellites. By building up a larger sample of observations, it may be possible to identify the requirements for U-shape scintillation to occur, and whether it is primarily dependent on the latitudinal extent of irregularities, irregularity amplitude, the magnitude of the background density and density gradients or some combination of these and other features, as was discussed briefly in Section 8.3.5. This would also enable stronger conclusions to be drawn on which region of the ionosphere the irregularities occur within than was possible here using only two case studies. Additionally, if more Swarm overpasses were included it may become more practical to investigate irregularity growth rates and confirm or deny the suggestion

made here that some combination of the GDI and/or KHI is likely the underlying mechanism forming the irregularities.

Chapter 9

Conclusions: LOFAR as an Ionospheric Observatory

Despite not being intended as such, LOFAR is one of the most sensitive instruments currently available for studying mid-latitude ionospheric variability. It has been used to study various phenomena, such as TIDs (e.g. Chapters 5, 6 and 7 of this thesis; Dorrian et al., 2023; Fallows et al., 2020), irregularities causing scintillation (e.g. Chapter 8 of this thesis; Flisek et al., 2023; Forte et al., 2022), and structures associated with sporadic-E layers (e.g. Trigg et al., 2024; Wood et al., 2024). In many cases, the structures observed using LOFAR have been shown to be difficult or impossible to detect using other common techniques for studying the ionosphere, such as GNSS (e.g. Chapter 6 of this thesis; Flisek et al., 2023). However, this is not to say that LOFAR is strictly superior to other instruments. For example, it is unable to determine large-scale background ionospheric conditions, and observing time is limited by the fact that it is primarily an astronomical rather than ionospheric instrument. The observations that can be made for ionospheric research are also limited by the positions in the sky of radio sources suitable for a given type of measurement (e.g. Chapter 8 of this thesis;

Wild and Roberts, 1956b).

TIDs are a particular phenomenon for which LOFAR can provide unique insights. Its ability to identify multiple simultaneous TIDs (e.g. Chapter 7 of this thesis; Fallows et al., 2020), extremely high sensitivity (~ 1 mTECu, e.g. Chapters 5, 6 and 7 of this thesis; Mevius et al., 2016) and ability to resolve short wavelengths (e.g. Chapters 5, 6 and 7 of this thesis) set it apart. TIDs have been identified with LOFAR using two distinct observational modes: single station observations (e.g. Chapter 5 of this thesis; Dorrian et al., 2023; Fallows et al., 2020) and combined interferometric observations (e.g. Chapters 6 and 7 of this thesis). Unfortunately, LOFAR is not currently capable of making both of these types of observation simultaneously, although the ongoing upgrade to LOFAR 2.0 (e.g. ASTRON, 2023; Edler et al., 2021) may enable this.

Chapter 7 shows the population of TIDs on a statistical basis, and establishes a processing method based on the work in Chapter 6 that could easily be extended to other LOFAR observations. The relationship identified between wave period and altitude at which the climatological wind best describes the propagation directions suggests a fundamental aspect of AGW propagation in the thermosphere which has not been reported elsewhere. In addition, the work in Chapter 5 demonstrates the possibility of using scintillation observations to identify TIDs at the shortest wavelengths available to LOFAR. While this is only a single case study, an automated method for identifying similar features in dynamic spectra and matching them to modelled spectra could be developed from this to provide an extension to the statistics of TIDs above LOFAR. The phase screen model developed in Chapter 5 has also proven useful in investigating other features in LOFAR data not associated with TIDs (e.g. Chapter 8 of this thesis; Trigg et al., 2024; Wood et al., 2024), providing a means of estimating length scales and amplitudes of structures responsible for the observed features.

The U-shape features identified in Chapter 8 are the first reported example of a

consistently occurring class of feature in the LOFAR scintillation observations. They show clear relationships with geomagnetic activity and time of day and/or season. It is shown that the irregularities responsible are highly field-aligned and these are tentatively suggested to exist within the auroral ionosphere, although the precise mechanism causing the frequency cutoff in scintillation requires more study. As part of this work, methods for processing LOFAR ionospheric scintillation observations were developed which could be used in larger statistical studies.

Existing ionospheric work with LOFAR primarily focuses on case studies rather than statistical investigations (e.g. Chapter 5 of this thesis; Dorrian et al., 2023; Fallows et al., 2020; Flisek et al., 2023; Trigg et al., 2024; Wood et al., 2024). These are clearly extremely useful, as they highlight the capabilities of LOFAR in terms of sensitivity (e.g. Trigg et al., 2024) and range of features that can be detected, as well as providing an opportunity to develop methods for data analysis (e.g. Fallows et al., 2020; Forte et al., 2022). However, case studies cannot determine the overall character of the ionospheric structures detected by LOFAR, and tend to bias in favour of unusual or eye-catching features (e.g. Chapter 5 of this thesis; Dorrian et al., 2023; Wood et al., 2024). To make the most use of LOFAR as an ionospheric instrument, it will be necessary to apply the understanding and analysis methods obtained from these case studies to large-scale statistical studies. Thus far, statistical analyses have largely been restricted to studying calibration solutions rather than single station observations (e.g. Chapter 7 of this thesis; de Gasperin et al., 2018; Mevius et al., 2016).

The main obstacle to such statistical analysis of LOFAR data has been the volume of data and difficulty of access to it. This was discussed in Chapter 8, where the full volume of data relevant to just the U-shape features was estimated as ~ 80 TB, preventing robust and quantitative description of the features on a statistical basis. Therefore, the work discussed in Section 8.4 that is being carried out at ASTRON to provide a

pre-processed, downsampled data product is vital to make rigorous statistical studies viable (at the time of writing only data from the earliest observing cycles has been processed). A further advantage of this pre-processing is that a standardised method is used for RFI mitigation, ensuring reproducibility of results. As well as the ionospheric scintillation observations discussed in this thesis, another source of data is the IDOLS project (Incremental Development of LOFAR for Space Weather: e.g. Zhang et al., 2023). This has used a single LOFAR station (CS032) to make continuous observations of Cassiopeia A with the LBA since mid 2023, along with solar observations during the daytime. While this is limited to only a single station, the continuous coverage prevents any potential bias in when observations are taken.

The pre-processed data from the ionospheric scintillation observations made using LOFAR will enable the work investigating U-shape features described in Chapter 8 to be extended substantially. Rather than relying on estimated timings from the quicklook plots, the variation of the cutoff frequency with time could be characterised based on the S_4 index. This would then enable quantitative comparison between events, to investigate systematic variations with geomagnetic activity, season/time of day or other factors. With a more robust characterisation of the cutoff frequency across a larger sample of events, potential explanations for the cutoff in scintillation, such as those proposed in Chapter 8 or others, could be tested more quantitatively by comparing the predicted and observed variations in cutoff frequency with time and/or observing geometry.

While the dTEC data from the calibration solutions is not so prohibitive in terms of data volume as the ionospheric scintillation observations (the 2,723 hours of calibration solutions used in Chapter 7 amount to ~ 840 GB), it is not currently made publically available. Alongside the upgrade to LOFAR 2.0, there are currently plans to make the dTEC solutions routinely available as they are produced. This would include reducing

the data volume substantially by removing additional calibration information (by a factor of approximately 400 based on the data used in Chapter 7). Making more dTEC solutions available would enable extending the work described in Chapter 7, for example to cover a wider range of local times and solar activities. If a substantial number of observations exist lasting longer than one hour, it would also be possible to consider periods greater than 27 min, although this would likely be limited by the reduced sensitivity to longer wavelength TIDs discussed in Chapter 6.

As well as data availability, statistical work will require generalisable approaches to processing of the single station data, analogous to that developed for identifying waves in the dTEC data in Chapter 7. This could include calculation of S_4 , characterisation of the power spectra, and estimations of velocity and anisotropy based on cross-correlations (see e.g. Chapter 8). Other useful quantities to estimate, exploiting the broadband observations made using LOFAR, could be the coherence bandwidth of scintillation or a measure the time lag between features observed on different frequencies.

While improved use of the data already obtained using LOFAR is necessary and important, the currently ongoing upgrade to LOFAR 2.0 (e.g. ASTRON, 2023; Edler et al., 2021) will provide improved capabilities. These include simultaneous observations using the LBA and HBA, and improved sensitivity particularly in the LBA, by improvements to the station electronics and data processing capacity (ASTRON, 2023; Edler et al., 2021). Simultaneous observations across a broader range of frequencies may enable more in depth investigation of the structures causing scintillation at LBA frequencies without corresponding scintillation at GNSS frequencies (e.g. Flisek et al., 2023) by providing observations at intermediate frequencies. The improved sensitivity with the LBA may also enable ionospheric scintillation observations to be made using a broader range of sources besides Cassiopeia A and Cygnus A, making it easier to

distinguish effects caused by season, time of day and observing geometry (e.g. Chapter 8 of this thesis; Wild and Roberts, 1956b). A further potential benefit of LOFAR 2.0 is the possibility of consistent observations of ionospheric scintillation in parallel with radio astronomy observations thanks to the increased data capacity, similar to the work carried out in the IDOLS project (e.g. Zhang et al., 2023) but using more stations.

Appendix A

LOFAR Interferometric Calibration and Derivation of dTEC

When LOFAR is used for radio astronomy observations, it operates as a large scale radio interferometer. The long baselines, especially when using the international stations, provide extremely high angular resolution. However, accurate imaging requires that unwanted corruptions of the data are removed. These arise from instrumental effects such as clock drift between stations (all core stations use the same clock, all remotes and internationals have independent clocks), bandpass of the antenna and beam shape, and ionospheric effects (de Gasperin et al., 2019). Due to the time variation of some of these effects, they are calculated separately at each integration step, typically ~ 4 s.

These effects create a combined complex gain factor at each station relative to a chosen reference station, which is what must be removed to create a successful image. However, accurately resolving this complex gain is not possible for an arbitrary pointing direction, as it requires a bright ‘calibrator’ source to give sufficient signal-to-noise ratio (SNR). As a result, alongside observations of the desired region of the sky (target field), a nearby calibrator is used to estimate this complex gain. In the LBA this can be done

in parallel for a calibrator in any location relative to the target, whereas for the HBA due to limitations of the beam forming the system can only simultaneously observe fields within $\sim 20^\circ$ of one another (de Gasperin et al., 2019). As doing this reduces the sensitivity in both fields, for HBA observations it is common to observe the calibrator in series for a short time before and after the main target to calibrate for the systematic effects which vary only very slowly with time (e.g. Shimwell et al., 2017).

The complex gain for the calibrator cannot simply be transferred directly to the target field however, as the corrupting effects are a mixture of direction-independent effects (which can be transferred) and direction-dependent effects (which in general cannot be transferred) (e.g. de Gasperin et al., 2019; Smirnov, 2011a). As a result, it is necessary to resolve the contributions of these individual effects in the calibrator field to determine the calibration which can be applied to the target.

The different effects are separated based on considering the complex gain (which is actually a 2×2 matrix corresponding to the two different linear polarisations) as a product of matrices each representing a distinct physical effect. This is known as a Jones chain (Smirnov, 2011a,b), and illustrates the importance of isolating the effects in the correct order as the matrices do not (in general) commute. Each effect has a characteristic type of matrix corresponding to it, for example ionospheric dispersive delay is simply a scalar as it acts identically on both polarisations.

The ionospheric effects considered are typically those associated with the first and second order terms of the expansion of the refractive index (neglecting absorption, see equation (3.2); the third order term becomes significant below ~ 40 MHz (de Gasperin et al., 2018)). These are phase delay and Faraday rotation respectively, and the phase delay is the effect determined purely by the difference in TEC. The ionospheric phase delay is evaluated along with the clock drift in the final calibration step (i.e. after all other effects have been calculated (de Gasperin et al., 2019)). The two effects are

distinguished by their different frequency dependence (clock effects are $\propto f$, ionospheric delay is $\propto f^{-1}$). The assumption that all other significant effects have been accounted for by this stage means that any residual corrupting effects will cause errors in the estimation of clock drift and/or ionospheric delay, but Mevius et al. (2016) have shown that the uncertainty in differential TEC derived from the HBA is < 1 mTECu.

This method of deriving TEC values from measured phase delays is analogous to the use of GNSS signals for TEC measurement mentioned in Section 3.1. However, absolute TEC cannot be calculated in the interferometric calibration as the initial phase relationship across frequencies is not known and the phase difference also has an ambiguity of $\pm 2n\pi$ (de Gasperin et al. (2018) demonstrated a method for estimating absolute TEC from LOFAR calibration solutions using differential Faraday Rotation and a model of the geomagnetic field in addition to dTEC, but this is not a direct measurement of absolute TEC). As a result, while the precision of differential TEC values calculated for interferometric calibration with LOFAR is higher than for absolute TEC derived from GNSS, the analysis of the data is more complex, as discussed in Section 6.1.

Appendix B

dTEC Noise Correlation Estimation

To quantify the correlation, for each baseline the noise spectrum is estimated by RLR with the rescaling described in Section 6.3.2. Then the wavelet spectrum is calculated, and masked to exclude all regions where the wavelet power exceeds the noise power, in order to remove as far as possible any contributions from coherent signals. This will also exclude some genuine noise contributions, but robustly excluding coherent signals from consideration is the most important factor here. For each period in the wavelet spectrum, the correlation between both the real and imaginary components of the wavelet coefficients on a given pair of baselines is calculated.

In this work the baselines are simply taken to be the ground level station to station baselines. The effective scan velocity of the lines of sight through the ionosphere is not considered. This will affect the apparent period of any disturbances, depending on their propagation direction and velocity. However, the apparent length scale (e.g. correlation length) and propagation direction are unaffected, and the correct period and hence velocity can be obtained if desired by calculating the line of sight scan velocity for an assumed ionospheric altitude.

To represent the correlations practically, a simplified model of the spatial correlation

function is required. Two options are considered: a simple exponential decay as a function of difference in baseline, and an exponential decay model assuming the noise is a combination of turbulence at both ends of the baseline. In other words, for two baselines \vec{r}_1 and \vec{r}_2 , the correlation R is assumed to be described by either

$$R = A \exp \left(\frac{-|\vec{r}_1 - \vec{r}_2|}{r_c} \right), \quad (\text{B.1})$$

or

$$R = \frac{1 + A \left(\exp \left(\frac{-|\vec{r}_1 - \vec{r}_2|}{r_c} \right) - \exp \left(\frac{-|\vec{r}_1|}{r_c} \right) - \exp \left(\frac{-|\vec{r}_2|}{r_c} \right) \right)}{\sqrt{1 - A \left(\exp \left(\frac{-|\vec{r}_1|}{r_c} \right) + \exp \left(\frac{-|\vec{r}_2|}{r_c} \right) \right) + A^2 \exp \left(\frac{-(|\vec{r}_1| + |\vec{r}_2|)}{r_c} \right)}}, \quad (\text{B.2})$$

where r_c is the correlation length and A is a scale factor representing the contribution of measurement errors which are assumed to be uncorrelated across baselines. Equation (B.2) is derived by assuming that the noise consists of equal contributions from the measurement at the chosen station and the reference, with that correlation defined according to equation (B.1). These two relationships are then fitted to the correlations for 100 1-hour observations randomly selected from 2,629 hours of observation (those remaining after the processing described in Section 7.1.1) made as part of the LOFAR LBA Survey (de Gasperin et al., 2021, ; see the supplementary material of Boyde et al. (2024) for the list of observations). This permits the testing of how stable the noise correlations are across different ionospheric conditions. In practice, it was found that the more complicated relationship described by equation (B.2) resulted in higher root mean square residuals after fitting, and therefore only the results using equation (B.1) are considered.

Appendix C

PCA Bias Correction

The PCA approach to estimating propagation direction is accurate in the idealised case only with a roughly uniform distribution of baselines in azimuth. With LOFAR this is not the case, and as a result the estimated azimuth will be biased. This bias is non-uniform in azimuth, and dependent on the reference station selected as this determines the azimuth distribution of the baselines.

In order to correct for this bias, for a given set of baselines, synthetic data corresponding to the idealised case (i.e. power proportional to the baseline projected onto the propagation direction squared) can be generated for a range of azimuths. For each input azimuth, the estimated azimuth from the PCA is then determined. An example of the relationship obtained using RS306 as a reference is shown in Figure C.1. This displays clear ‘preferred directions’ around $70-85^\circ$ and $-5-10/175-190^\circ$ where many true azimuths generate estimated azimuths in the same narrow ranges.

Using this idealised picture of the bias, it is generally possible to define a means of inverting the estimates to mitigate the effect of the bias. This is possible provided the variation is monotonic, as it is in Figure C.1. In this case, a correction function can be defined as a linear interpolation of the relationship between true azimuth and estimated




Figure C.1. The bias in the azimuth estimated from PCA using RS306 as a reference in the idealised case.

azimuth, and used to map the estimates back to something closer to their true value. Neglecting such a correction has been found to reduce reliability of the fitting based on synthetic data testing as described in Section 6.4, and in some cases prevent a viable solution from being found, even when the wave activity is clear from a visual inspection of the data.

Appendix D

Synthetic dTEC Data

For the synthetic data, the noise is replicated by a simplified model, the AR1 model discussed in Section 6.3.1 with $\alpha = 0.95$. While this is not an accurate representation of the noise in the real data, given that the fit method is not dependent on the spectral shape it should be suitable for this purpose. The random samples (z_i in equation (6.6)) are generated with correlation between baselines to replicate the correlation in the data (assumed $r_c = 13.2$ km, $A = 0.95$). This means that the assumed correlation used to generate $\underline{\underline{C}}$ in equation (6.11) is exactly matched to the true correlation for the synthetic data. This provides another reason that the performance of the fit method on the synthetic data should be better than its performance on real data where the correlations are only approximate.

One further possible characteristic of the noise in the real data that is not replicated here is that it may be directly driven by the waves themselves. For example, the density gradients associated with the wave can favour the development of the GDI and resulting turbulent cascade (e.g. Lin et al., 2016). This would mean that the noise would be enhanced in the presence of a wave relative to the background. However, this is not investigated further here.

The synthetic data generation uses 1 hour of data per set of parameters. This is based on the 1 hour duration of the data from the LBA survey (de Gasperin et al., 2021) which was used to estimate the correlation length in Section 6.3.3 and will be used further in the following chapter. The wave parameters are generated using Sobol sequences (Saltelli et al., 2010; Sobol and Levitan, 1999) to generate representative sampling across the parameter space with minimum clustering. The parameters generated are: wavelength Λ , between 100 – 625 km; azimuth, between 0 – 360°; target median SNR, between 5 – 50; and a ‘period scale’, between 0 and 1. The period scale is used to ensure that the wave velocity is physically reasonable, with a maximum velocity set at 400 m s^{-1} . Based on the generated wavelength, the minimum period is then defined as the period which would give the maximum velocity, and then the period scale determines the wave period linearly between this minimum period and the maximum period set by the wavelet cone of influence, which is $\sim 26 \text{ min}$ for a 1 hour observation. The result of this generation scheme is that wavelength is sampled effectively uniformly, while periods are biased towards the higher end of the available range and amplitudes are roughly proportional to wavelength with the proportionality constant determined by the target SNR. The minimum wavelength of 100 km is set to ensure no additional phase discontinuities occur within the size of the array. While the fit method can in theory account for these, it will likely impact performance and so is not considered further in this analysis.

Appendix E

Wave Event Characterisation

The method presented in Chapter 6 calculates the wave properties from the calibration solutions for a specified time and wave period. Given that these cannot be manually defined for analysing a large data set as considered in Chapter 7, it was necessary to define a sampling approach for selecting the times and periods. This was achieved by calculating the SNR of the wavelet spectrum for all baselines between every combination of remote stations and CS002 (52.915°N, 6.870°E; other core stations were neglected here to avoid biasing the results towards the very short baselines within the core) and masking the wavelet spectrum to leave only regions where the median SNR was greater than 3 (corresponding to the 95% confidence level for identifying a signal in the wavelet spectrum: Torrence and Compo, 1998). The spectrum was also masked for all periods below 1 minute and for all points within the CoI. Including the region below 1 minute period was found in test cases not to provide any reliable fits and so was excluded from the full analysis to save computational time. Within the remaining unmasked regions, for each period the wavelet spectrum was sampled at a time step of one tenth the width of the CoI, which defines the time range across which the wavelet coefficients at a given period are correlated. The width of the CoI Δt for a given period T is given

by (Torrence and Compo, 1998)

$$\Delta t = \frac{4\sqrt{2}T}{5}. \quad (\text{E.1})$$

Samples were taken every tenth of a CoI width to intentionally over-sample at this stage, ensuring that there would be a sample sufficiently close to the time of peak amplitude for any given wave activity.

For some of the sampled times and periods, despite the significant power measured by the wavelet transform, there may not actually be wave activity or it may not be adequately described by the simple plane wave model. In order to remove these fits from further analysis, fits were flagged if they had a reduced chi-squared higher than 5, relative uncertainty in wavelength or amplitude of greater than 50%, estimated measured amplitude on a 30 km baseline below 1 mTECu or estimated wavelength greater than 1000 km. The minimum amplitude was based on the known precision of LOFAR dTEC data (Mevius et al., 2016) and the maximum wavelength was set to twice the upper limit estimated in the Chapter 6 based on synthetic dTEC values. The limits on amplitude and wavelength were deliberately chosen to allow most fits through at this stage, as tighter limits could be imposed later in the analysis if necessary.

Once the solutions had been calculated and masked, the next step was to identify individual wave ‘events’ from within them. A given wave will show up at multiple times and periods, both due to its own finite duration and range of frequencies and the non-orthogonality of the wavelet basis causing it to be further spread in the wavelet spectrum. In order to determine groups of solutions corresponding to a single wave event, a version of the DBSCAN clustering algorithm (Ester et al., 1996) was applied. This algorithm identifies clusters within data based on the density of points within the N-dimensional space, and also identifies and excludes noise from the identified clusters. It can operate with an arbitrary distance metric, and requires two further parameters to be specified. One of these is the neighbourhood size ϵ and the other is the number

of points required within a points neighbourhood to form a cluster *MinPts*.

Points from a single wave event can be expected to be both close in period and time (i.e. close in the wavelet spectrum) as well as having similar fitted wave parameters, specifically azimuth of propagation and phase velocity which should be reasonably stable both with time and across periods. Based on this concept, two distance metrics $r_{spec;i,j}$ and $r_{fit;i,j}$ between a given pair of points i and j were defined as

$$r_{spec;i,j} = \sqrt{\frac{(\log_2(T_i) - \log_2(T_j))^2}{\epsilon_T} + \frac{(t_i - t_j)^2}{<\Delta t>_{i,j}^2 \epsilon_t}}, \quad (\text{E.2})$$

and

$$r_{fit;i,j} = \sqrt{\frac{(v_i - v_j)^2}{(\sigma_{v;i}^2 + \sigma_{v;j}^2) \epsilon_v} + \frac{(\theta_i - \theta_j)^2}{(\sigma_{\theta;i}^2 + \sigma_{\theta;j}^2) \epsilon_\theta}}, \quad (\text{E.3})$$

where t is the time within the observation, $<\Delta t>_{i,j}$ is the mean CoI width at periods T_i and T_j calculated from equation (E.1), v is the estimated phase velocity, θ is the estimated propagation azimuth, σ_v and σ_θ are the respective uncertainties and the ϵ_x parameters are scales for the various distance components. The overall distance $r_{i,j}$ was then taken as the maximum of $r_{spec;i,j}$ and $r_{fit;i,j}$. As discrete wave events within a given observation were generally clearly distinguished visually, the ϵ_x parameters were then determined manually along with the threshold minimum number of points to form a cluster, which will be discussed below. This was achieved using a small number of observations as test cases to determine values for the ϵ_x and *MinPts* which produced clusters consistent with a manual visual inspection. Given the freedom of these ϵ_x parameters within the distance metric, the global neighbourhood size parameter ϵ could simply be set to 1 in this implementation. This was because ϵ defines the threshold value of $r_{i,j}$ required for two points to be considered neighbours, and any given effective size could be obtained holding $\epsilon = 1$ and scaling the ϵ_x parameters.

Given the two part definition of the distance metric in terms of r_{spec} and r_{fit} and the




Figure E.1. The maximum possible number of neighbours for different locations in the wavelet spectrum. This assumes median SNR > 3 at all points and $r_{fit;i,j} < 1$ for all i and j . The shaded and cross-hatched region is the cone of influence.

method of sampling the wavelet spectrum space described above, it is possible to define the maximum possible number of neighbours for a given point in the wavelet spectrum space, which is determined by the scale factors ϵ_T and ϵ_t . There is no obvious criteria for defining ‘correct’ values for these parameters, so in practice the values were selected based on those which gave reasonable results in test cases, in that regions of obviously consistent wave activity were grouped into clusters without clearly distinct regions being combined together. These suitable values for ϵ_T and ϵ_t were found to be 0.25 and 0.75 respectively, although in practice the results are not especially sensitive to moderate (i.e. up to a factor of ~ 1.5) changes in either of these parameters. The maximum number of neighbours was found by sampling assuming that the median SNR is greater than 3 for all times and periods and then calculating the number of neighbours for each

sampled point assuming that $r_{fit;i,j}$ is less than 1 for all pairs of points. The result of this is shown in Figure E.1, and it illustrates why using a fixed value of *MinPts* was found to be unsuitable in this case. At periods above roughly 16 minutes, the maximum number of neighbours drops substantially due to the presence of the cone of influence. If *MinPts* was set low enough to detect clusters in this region it was found to also detect spurious clusters especially at the lowest periods. To remedy this, the DBSCAN algorithm was adjusted to determine *MinPts* individually for each point considered. The maximum number of neighbours shown in Figure E.1 was averaged at each period to obtain a representative value. A point was then considered a core point in DBSCAN if its number of neighbours was 60% or more of the maximum number for its period. For the scale factors ϵ_v and ϵ_θ , these were both set to $\sqrt{2}$ to correspond to a ‘1 sigma’ difference in the fit values.

Appendix F

Short Period Waves: Possible Plasmaspheric Structures

One notable aspect of data obtained from radio interferometers observing natural radio sources compared to other techniques for observing TIDs is their capability to detect plasmaspheric as well as ionospheric disturbances. GNSS satellite orbits are typically too low to observe most of the plasmasphere, and while signals transmitted by satellites in higher orbits such as geostationary satellites will be affected by plasmaspheric structure it is difficult to disentangle any plasmaspheric impacts from the effect of the ionosphere as the ionospheric variations will typically be of much higher magnitude. However, due to the motion of the line of sight to an astronomical source relative to the plasmasphere having a relatively high effective velocity compared to that in the ionosphere, the Doppler shift effect can shift the apparent period of plasmaspheric disturbances to short periods where ionospheric disturbances are minimal (Jacobson and Erickson, 1993). Previous observations using the VLA (Thompson et al., 1980) found disturbances in the short period range ($< \sim 5$ min) with a strong preference for propagation towards magnetic East (Jacobson and Erickson, 1992, 1993). Later investigation

showed that these structures are field-aligned plasmaspheric structures which actually tend to move westwards, with the apparent eastward propagation being induced by the motion of the line of sight (Hoogeveen and Jacobson, 1997).

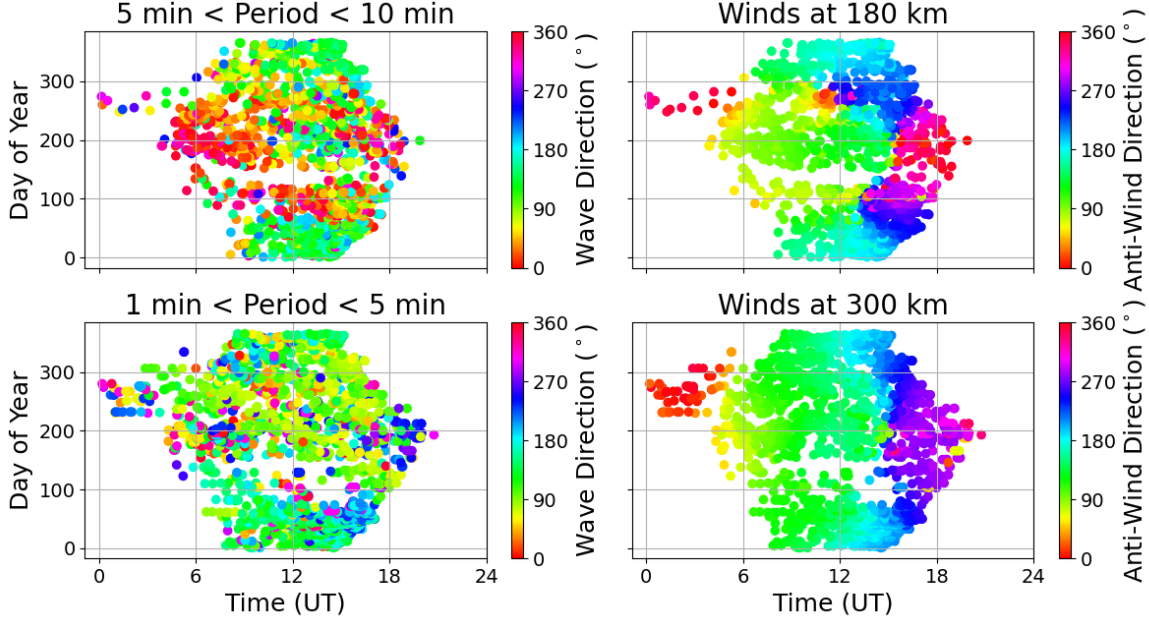


Figure F.1. The left column shows observed wave propagation directions and the right column shows HWM14 modelled anti-wind directions for the optimum altitude calculated from Figure 7.9 at the corresponding times. Each row corresponds to a different period range.

To investigate the possibility of field-aligned plasmaspheric structures being included in the short period waves shown in Figure F.1, for each wave event the source azimuth and elevation were used to calculate the projection of the magnetic field into the ground plane. This was carried out for various altitudes along the line of sight to determine if there is an altitude for which the geomagnetic field is perpendicular to the observed propagation, following the approach of Hoogeveen and Jacobson (1997).

To calculate the orientation of the ‘magnetic field shadow’ (the apparent orientation of field aligned structures when projected into the LOFAR ground plane) for a given observation, first the co-ordinates of the line of sight at a given altitude were calculated following the method given by Dorrian et al. (2023). Once the geographic

latitude and longitude corresponding to the specified altitude were known, the magnetic field orientation (inclination and declination) at that location were obtained from the International Geomagnetic Reference Field (IGRF: Alken et al., 2021). This provided the magnetic field vector in the local Cartesian co-ordinates (East, North and vertical). In order to compare this to the line of sight direction, defined in terms of source azimuth and elevation as viewed from LOFAR, the magnetic field vector must therefore be converted into the local Cartesian co-ordinates of LOFAR. For the purposes of these calculations, the location of LOFAR was taken to be that of station CS002, the most central of the core stations.

The conversion from the local Cartesian co-ordinates in which the magnetic field is defined (at a latitude and longitude of λ_i, ϕ_i) and those of LOFAR (latitude and longitude of λ_0, ϕ_0) is non-trivial to carry out directly. For simplicity, a universal co-ordinate system was defined in which the z co-ordinate points along the Earth's rotation axis towards the North pole, the x co-ordinate points through 0° longitude and the y co-ordinate points through 90°E longitude. The transformation to and from any local Cartesian system to this universal system was then defined allowing transformations between any two local Cartesian systems in a two step process. The co-ordinate vectors at (λ_i, ϕ_i) were defined as \hat{E}', \hat{N}' and \hat{Z}' respectively, and those at (λ_0, ϕ_0) as \hat{E}, \hat{N} and \hat{Z} respectively. The co-ordinate vectors in the universal co-ordinate system were \hat{x}, \hat{y} and \hat{z} respectively.

The transformation from (λ_i, ϕ_i) to the universal co-ordinate system can then be expressed as:

$$\begin{aligned}\hat{x} &= -\sin(\phi_i)\hat{E}' + \cos(\phi_i)\left[\cos(\lambda_i)\hat{Z}' - \sin(\lambda_i)\hat{N}'\right], \\ \hat{y} &= \cos(\phi_i)\hat{E}' + \sin(\phi_i)\left[\cos(\lambda_i)\hat{Z}' - \sin(\lambda_i)\hat{N}'\right], \\ \hat{z} &= \cos(\lambda_i)\hat{N}' + \sin(\lambda_i)\hat{Z}'.\end{aligned}\tag{F.1}$$

The transformation from the universal co-ordinate system to (λ_0, ϕ_0) is then:

$$\begin{aligned}\hat{\vec{E}} &= \cos(\phi_0)\hat{\vec{y}} - \sin(\phi_0)\hat{\vec{x}}, \\ \hat{\vec{N}} &= \cos(\lambda_0)\hat{\vec{z}} - \sin(\lambda_0) \left[\cos(\phi_0)\hat{\vec{x}} + \sin(\phi_0)\hat{\vec{y}} \right], \\ \hat{\vec{Z}} &= \sin(\lambda_0)\hat{\vec{z}} + \cos(\lambda_0) \left[\cos(\phi_0)\hat{\vec{x}} + \sin(\phi_0)\hat{\vec{y}} \right].\end{aligned}\tag{F.2}$$

Once the magnetic field vector was transformed into the co-ordinate system at (λ_0, ϕ_0) its ‘shadow’ was then calculated given the orientation of the line of sight. This ‘shadow’ is the vector purely in the horizontal (E-N) plane which has the same projection onto the plane perpendicular to the line of sight as the magnetic field vector itself. To calculate this the magnetic field vector \vec{B} was added to a linear multiple of the line of sight vector \vec{l} such that the resulting vector \vec{B}' had no component in $\hat{\vec{Z}}$. This can be expressed mathematically as

$$\vec{B}' = \vec{B} - \frac{B_Z}{l_Z}\vec{l},\tag{F.3}$$

where B_Z and l_Z are the components along $\hat{\vec{Z}}$ of \vec{B} and \vec{l} respectively.

To identify potential field-aligned structures, this calculation was initially carried out for altitudes from $0.1R_E$ to $10.1R_E$ in steps of $1R_E$ (R_E denotes the radius of the Earth) for each wave event. From these, either the wave was classified as non-field-aligned (no altitudes for which \vec{B}' was close to perpendicular to the wave propagation) or Newton-Raphson iteration was used to determine the altitude at which the structure would align with the field (using a tenth of the estimated azimuth uncertainty of the wave as a convergence criterion, the Monte Carlo randomly generated azimuths used a tenth of the mean of the estimated azimuth uncertainties which was approximately 0.5°).

Given that the IGRF is a model of the core field, it does not include contributions from current systems in the ionosphere or magnetosphere which will become more sig-

nificant at greater altitudes. As a result, its accuracy at the highest altitudes considered here is doubtful. However, it will be shown below that the altitudes of alignment are predominantly much lower than $10R_E$ ($< \sim 2R_E$, see Figure F.2). At these altitudes the IGRF is a more reliable model of the field than at $10R_E$, and hence restricting the altitude range further than was done here would in fact strengthen the conclusions drawn rather than weaken them.

Naturally, even if no such field-aligned waves exist in the data some waves would happen to align with the magnetic field regardless, purely by chance. To test if this is the case, a Monte Carlo analysis was also carried out where 20 random directions were generated corresponding to each wave event and each compared to see if they were field-aligned. This Monte Carlo analysis used three different methods for generating random directions: uniform (i.e. completely random) directions, normal distribution around the HWM14 anti-wind direction at 180 km with 30° standard deviation (to represent F-region propagation), and normal distribution around the HWM14 anti-wind direction at 100 km with 30° standard deviation (to represent E-region propagation).

In terms of the number of field-aligned events, for periods of 1 – 5 min there were 1,028 out of a total of 2,131. From the Monte Carlo results the expected values were estimated as 855, 760 and 756 for the uniform, 180 km and 100 km distributions respectively. This indicates a small but consistent increase in the number of field-aligned waves relative to what might be expected to occur by chance. As well as the number of events, the altitudes at which the field-alignment was found for the real data compared to the Monte Carlo results (shown in the top-left panel of Figure F.2) supports the idea that these are field-aligned waves in the plasmasphere. Relative to the Monte Carlo results, the altitudes of field-alignment for the real data are skewed lower, generally at altitudes below $1R_E$, which is expected given that at lower altitudes plasmaspheric densities will be higher and hence perturbation amplitudes are more likely to be high

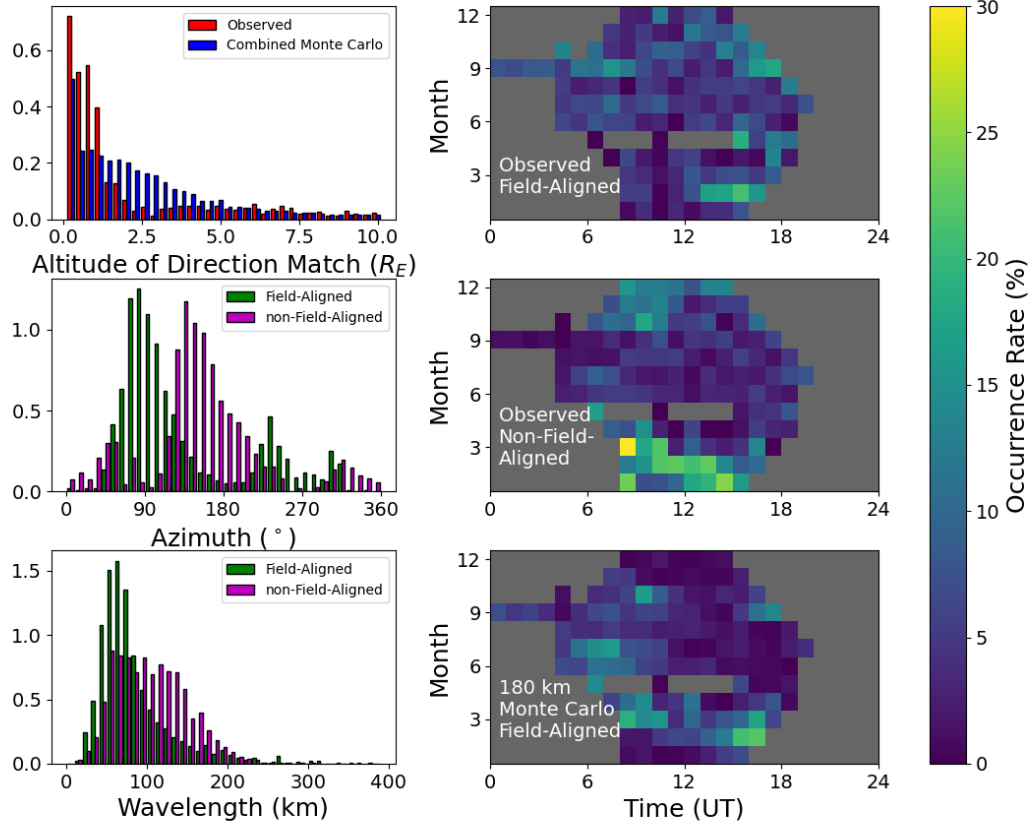


Figure F.2. The results of examining the field-alignment of the observed waves with periods between 1 – 5 min. The horizontal axis of the top-left panel shows the altitude (in units of Earth radii R_E) at which the propagation direction matched the field orientation for the observed waves (red) and the combination of all three Monte Carlo tests (the Monte Carlo results are combined for clarity as they were all consistent with one another). The middle-left panel shows the observed azimuth of propagation for field-aligned (green) and non-field-aligned (magenta) waves, while the bottom-left panel compares observed wavelengths for the same groups. Each histogram is independently normalised, and the vertical axis gives the relative probability in arbitrary units. The right hand panels show occurrence rates as a function of time and month for the observed field-aligned waves (top-right), observed non-field-aligned waves (middle-right) and Monte Carlo results using 180 km HWM14 winds (bottom-right).

enough to be detected. This emphasises the point that if the altitude range considered were further restricted due to the limitations of using the IGRF at large distances from the Earth, more of the Monte Carlo results would be excluded than the observed waves, strengthening the inference that the observations contain real field-aligned plasmaspheric structure.

Comparing the wavelengths of the field-aligned and non-field-aligned waves in the data (bottom-left panel of Figure F.2) also suggests that they come from different populations, with field-aligned waves having markedly shorter wavelengths. This wavelength difference indicates that the plasmaspheric field-aligned disturbances cannot explain the stronger high wavelength tail for the waves at 1–5 min periods that was noted in Figure 7.4. Finally, the occurrence of field-aligned events in time and season (top-right panel of Figure F.2) shows a relatively uniform occurrence of field-aligned events across observing conditions, whereas the non-field aligned events cluster at certain times (middle-right panel of Figure F.2). For the Monte Carlo results, using HWM14 winds to constrain expected wave directions causes very clear clustering in time and season for the field-aligned events (when winds are predominantly zonal, as illustrated in the bottom-right panel of Figure F.2) which is not observed in the real data. The same analysis as shown in Figure F.2 was also carried out for periods of 5–10 min (not shown) but did not show any evidence of an equivalent population in that period range, with fewer events than predicted by the Monte Carlo simulations and no preference for lower altitudes or significant bias in wave parameters between field-aligned and non-field aligned events.

These results suggest that at least for periods below 5 min there is a population of field-aligned plasmaspheric structures contributing to the waves observed with LO-FAR, broadly consistent with those observed by Hoogeveen and Jacobson (1997). The ratio of field-aligned events to non-field aligned events is higher than reported from the VLA, with Hoogeveen and Jacobson (1997) reporting 4,581 of 14,017 (32.7%) com-

pared to 1,028 of 2,131 (48.2%) in the data from LOFAR, possibly due to the different geographical locations and different sources observed. The wavelengths and velocities of the field-aligned disturbances are also broadly similar between the two datasets. However, the non-field-aligned waves show clear differences, with strong preference for certain azimuths in the LOFAR data (see middle-left panel of Figure F.2) compared to relatively uniform azimuth distribution reported in Hoogeveen and Jacobson (1997) (see their Figure 3). Their field-aligned waves also constitute a very distinct peak in the azimuth distribution, whereas the field-aligned waves seen here are much less prominent in the overall azimuth distribution. This may be due to different field geometry probed by LOFAR compared to the VLA due to their different locations and different sources observed, which may spread out the range of azimuths at which field-alignment is possible for the LOFAR observations compared to those of Hoogeveen and Jacobson (1997) and hence suppress the sharp peak around magnetic East.

While this indicates an explanation for some of the observed short period waves, the majority of them are still unexplained. Even with the field-aligned waves removed, there is no apparent match with the climatological neutral winds at any altitude, although the magnitude of the standard deviation is somewhat reduced compared to Figure 7.9 (a minimum of $\sim 75^\circ$ rather than $\sim 85^\circ$ in Figure 7.9). This may indicate that there is a further subdivision of the waves into different altitude regimes, possibly associated with secondary or tertiary AGWs created by the breaking of primary AGWs (e.g. Vadas and Fritts, 2002; Vadas et al., 2003), meaning that no single altitude can adequately explain their behaviour. Alternatively, the remaining waves may still be plasmaspheric but simply not field-aligned, which would explain their existence at periods which are unlikely for AGWs in the thermosphere but would require an alternative generation mechanism to be proposed.

Appendix G

Additional U-Shape Phase Screen Modelling

The irregularity spectrum used for the phase screen modelling was defined by an inner and outer scale and a power law variation between these, and set to zero outside these limits. A phase screen realisation was then generated by multiplying each wavelength component of this spectrum by a normally distributed real and imaginary variable to provide a spectrum with randomised phase and chi-squared distributed power around the idealised spectrum.

Given that the frequency dependence of the scintillation was of particular interest, it was important to quantify the frequency dependence of the instrumental effects. To achieve this, rather than taking constant values for the intensity of Cygnus A and the SEFD of LOFAR as in Chapter 5, the frequency dependence of both of these quantities was estimated. This is described by equation 8.2 and equation (1) of McKean et al. (2016), and explained in more detail in Section 8.3.3

To confirm that the absence of any cutoff in the modelled scintillation shown in Section 8.2.1 was not dependent on the model parameters used, the modelling has been

rerun varying each of the parameters in turn. In addition, the definition of the irregularity spectrum outside the range of the inner and outer scales was changed to ensure this did not have any significant effect on the results. S_4 was calculated over the full time range of the simulation, excluding only a portion at the start and end where edge effects may distort the results. This window was ~ 2300 s, which would be an extremely long time window to use in real data as the statistics of the scintillation cannot reasonably be assumed to be stationary over such a long time. However, for modelled data in which the statistics of the scintillation are known to be exactly stationary, a longer time window is appropriate, and provides a more precise estimate of the statistical average S_4 spectrum than would be obtained using a shorter time window.

The results of including finite spectral power outside the range defined by the inner and outer scales are shown in Figures G.1 and G.2. The finite power outside the spectral shape was represented by constant power at scales greater than the outer scale (equal to the power law power at the outer scale) and a steepening of the power law at scales below the inner scale by 2 (i.e. given the spectral slope of $\frac{-11}{3}$ between the inner and outer scales, the slope below the inner scale is $\frac{-17}{3}$). The spectral power between the inner and outer scales is unaltered, meaning that the additional power outside this range increases the root mean square TEC variation. For the parameters used in Figures G.1 and G.2, this corresponds to an increase in root mean square TEC variation from 10 mTECu to 17.6 mTECu. Figures G.1 and G.2 show no discernible difference between the two forms of the irregularity spectrum that were considered, and so for testing the impact of varying the other parameters the simpler form with zero power outside the inner and outer scales was used.

The effect of varying the root mean square TEC variation is shown in Figures G.3 and G.4. The root mean square TEC variation clearly has a significant impact on the scintillation behaviour, but in no cases is there any cutoff in scintillation comparable




Figure G.1. A comparison of the dynamic spectra obtained by assuming that the spectral power is zero except between the inner and outer scales (left panel) and with finite power outside this region (right panel, see text for details). All parameters used are the same as those in Figure 8.6, and the source elevation is set to 10° .




Figure G.2. A comparison of S_4 as a function of frequency obtained by assuming that the spectral power is zero except between the inner and outer scales (black solid line) and with finite power outside this region (red dashed line, see text for details). All parameters used are the same as those in Figure 8.6, and the source elevation is set to 10° .




Figure G.3. A comparison of the dynamic spectra obtained by varying the amplitude of the TEC variations, quantified by the root mean square variation. All other parameters used are the same as those in Figure 8.6, and the source elevation is set to 10° .




Figure G.4. A comparison of S_4 as a function of frequency obtained by varying the amplitude of the TEC variations, quantified by the root mean square variation. All other parameters used are the same as those in Figure 8.6, and the source elevation is set to 10° .

to that observed in the real data. Figure G.4 does show a decrease in S_4 towards the lower frequencies at the higher values of root mean square TEC variation considered here, but this is a smooth gradual decrease rather than the sharp transition observed in the U-shape scintillation events. It is also striking that the S_4 index starts to decrease with increasing root mean square TEC which corresponds to increasing scintillation strength. This effect has been discussed by several previous authors (e.g. Carrano and Rino, 2016; Gochelashvily and Shishov, 1971; Singleton, 1970) and arises due to strong focusing causing a maximum of S_4 which is greater than the limiting value as scintillation strength tends to infinity. This also explains the observed decrease in S_4 towards lower frequencies in Figure G.4, as the scintillation strength at lower frequencies for a given screen is higher.

As with the tests shown above, Figures G.5 and G.6 show that varying the spectral index of the irregularities does not give rise to any cutoff in the scintillation with frequency. In this case the spectral index was varied while keeping the root mean square TEC variation constant, meaning that a steeper spectral slope corresponds to more variation on longer scales and less on shorter scales. The main difference apparent in Figure G.3 is the longer timescales characterising the scintillation for the steepest spectral index (bottom right panel) due to the increased significance of the larger spatial scales in this case. The S_4 index in Figure G.4 shows that the scintillation at the lower frequencies becomes stronger with increasing (i.e. more negative) spectral slope, but even for the shallower slopes the decrease towards the lower frequencies is smooth and gradual and does not resemble the sharp cutoff observed in the U-shape scintillation events.

The effect of varying the outer scale of the irregularities was tested and the results are shown in Figures G.7 and G.8. For this case, rather than holding the root mean square TEC perturbation constant, the power at the inner scale was held constant.




Figure G.5. A comparison of the dynamic spectra obtained by varying the spectral slope of the irregularities. All other parameters used are the same as those in Figure 8.6, and the source elevation is set to 10° .




Figure G.6. A comparison of S_4 as a function of frequency obtained by varying the spectral slope of the irregularities. All other parameters used are the same as those in Figure 8.6, and the source elevation is set to 10° .

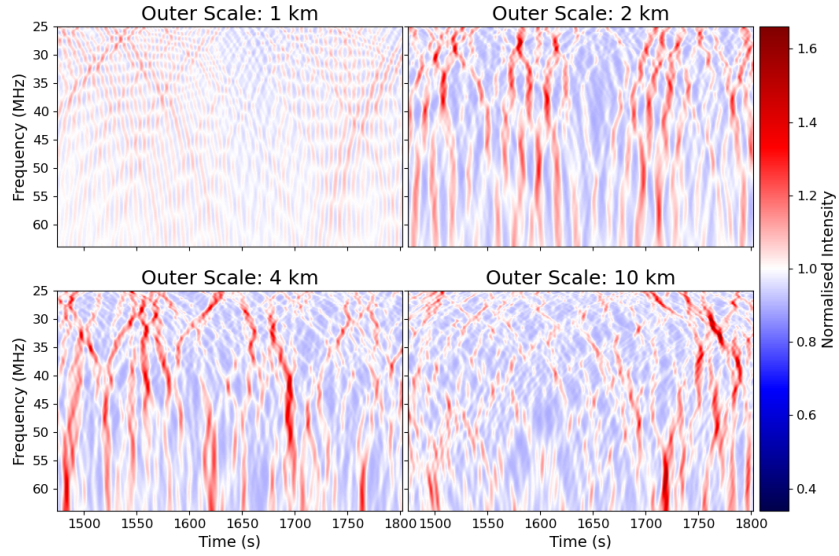


Figure G.7. The dynamic spectra obtained by varying the outer scale of the irregularities. All other parameters are the same as those in Figure 8.6 (although root mean square TEC variation is not held constant, see text for details), and elevation is set to 10° .




Figure G.8. S_4 as a function of frequency obtained by varying the outer scale of the irregularities. All other parameters are the same as those in Figure 8.6 (although root mean square TEC variation is not held constant, see text for details), and elevation is set to 10° .

Given that the spectral index was also constant, this corresponds to a fixed power law relationship being assumed and simply varying the outer scale at which the power law is truncated. Decreasing the outer scale below 4 km causes a general decrease in scintillation power as shown in Figure G.8, due to the lack of irregularities around the Fresnel scale. Increasing it above this value has a relatively negligible impact as expected given that structure on this scale does not contribute to intensity scintillation as explained in Section 3.3.1. As with the parameters considered previously, there is no value which leads to a cutoff in scintillation at the lower frequencies as is observed in the U-shape scintillation events.

Figures G.9 and G.10 show the results of varying the inner scale between 5 – 500 m, which has no discernible impact on the scintillation observed, as in any case these values are significantly below the Fresnel scale. As with varying the outer scale above, rather than holding the root mean square TEC perturbation constant, the power at the outer scale was held constant. Given that the spectral index was also constant, this corresponds to a fixed power law relationship being assumed and simply varying the inner scale at which the power law is truncated.

In general, no combination of irregularity parameters was found that could even qualitatively replicate the scintillation cutoff observed in the U-shape scintillation events. This indicates that whatever the explanation may be, the underlying physics cannot be represented by a single one-dimensional phase screen affecting all frequencies.




Figure G.9. The dynamic spectra obtained by varying the inner scale of the irregularities. All other parameters are the same as those in Figure 8.6 (although root mean square TEC variation is not held constant, see text for details), and elevation is set to 10° .




Figure G.10. S_4 as a function of frequency obtained by varying the inner scale of the irregularities. All other parameters are the same as those in Figure 8.6 (although root mean square TEC variation is not held constant, see text for details), and elevation is set to 10° .

Appendix H

Cross-Correlation Fitting

The assumed shape of the cross-correlation peak $C(\vec{\zeta}, \tau)$ can be expressed as

$$C(\vec{\zeta}, \tau) = \rho\left(\left[\vec{\zeta}, \tau\right]^T \underline{\underline{Q'}} \left[\vec{\zeta}, \tau\right]\right), \quad (\text{H.1})$$

where $\vec{\zeta}$ is the spatial baseline, τ is the time delay, ρ is some arbitrary function and $\underline{\underline{Q'}}$ is the matrix determining the shape of the quadratic relationship, which is symmetric by definition (equation (17) of Grzesiak et al., 2022). For short baselines and small time delays, this can be approximated as

$$C(\vec{\zeta}, \tau) \approx q_{11}\zeta_x^2 + q_{22}\zeta_y^2 + q_{33}\tau^2 + 2q_{12}\zeta_x\zeta_y + 2q_{13}\zeta_x\tau + 2q_{23}\zeta_y\tau + 1, \quad (\text{H.2})$$

where ζ_x and ζ_y are the components of the vector $\vec{\zeta}$ and the q_{ij} are the components of the matrix $\underline{\underline{Q'}}$ (equation (22) of Grzesiak et al., 2022). By considering the behaviour in the plasma rest frame where the correlation is defined by the matrix $\underline{\underline{Q}}$ and defining $\underline{\underline{Q'}}$

as a Galilean transform of $\underline{\underline{Q}}$ by the plasma velocity \vec{v} , it can be shown that

$$\vec{v} = \begin{pmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{pmatrix}^{-1} \cdot \begin{pmatrix} q_{13} \\ q_{23} \end{pmatrix}. \quad (\text{H.3})$$

(note that $q_{12} \equiv q_{21}$) (equation (20) of Grzesiak et al., 2022). While equation (H.3) shows that the temporal decorrelation term q_{33} is not used to calculate the velocity \vec{v} , it is nonetheless vital that it is included in the fit process. If it is not, the fit would assume that the peak correlation value remains 1 for all time delays τ and simply changes location, corresponding to completely frozen flow. As the irregularities responsible for scintillation typically evolve significantly as they pass across the LOFAR stations, neglecting this would significantly degrade the accuracy of the derived velocity.

The orientation of the irregularities can be determined using the spatial component of the matrix $\underline{\underline{Q'}}$. The orientation of the major or minor axis of the anisotropy in terms of an angle θ relative to the y -axis can be expressed as

$$\theta = \frac{1}{2} \tan^{-1} \left(\frac{2q_{12}}{q_{22} - q_{11}} \right). \quad (\text{H.4})$$

By considering the variation of $C(\vec{\zeta}, \tau)$ along this direction compared to the perpendicular direction, it can therefore be determined whether this is the major or minor axis. For the purposes of further discussion, the major axis will be used to define the anisotropy orientation.

Typically, in order to reduce the plasma motion to two spatial dimensions, it is assumed that the motion is entirely horizontal (e.g. Chapter 5 of this thesis; Fallows et al., 2020; Grzesiak et al., 2022). As a result, ζ_x and ζ_y typically therefore correspond to geographic East and North respectively. However, the scintillation measurements are effectively insensitive to motions along the line of sight (assuming the irregularity screen

itself is not moving, if this is not the case then motions of the screen could change the distance to the irregularities significantly and therefore be observable), meaning that the spatial velocity and anisotropy measured is actually in the plane of sky (i.e. normal to the line of sight). At high elevations, the difference between these two orientations is negligible, but at the low elevations considered here it is considerable. It may be suitable to use horizontal co-ordinates if the irregularities in question can be well approximated by a thin horizontal phase screen, that is that their vertical extent is significantly less than their horizontal extent. Otherwise, it is vital to consider measurements initially in the plane of sky and then convert the velocity to horizontal if appropriate.

To define the plane of sky here, the z axis is assumed to align with the line of sight, positive towards the source. The x axis is then defined as the horizontal direction normal to z , such that it points in a clockwise direction from the observer. Completing the right-handed co-ordinate set, the y axis is then normal to both of these, with positive y oriented in the downwards sense. Angles within this plane are then defined relative to the y axis as in equation (H.4), ensuring that angles of $\pm 90^\circ$ have a clear physical interpretation as corresponding to the horizontal direction within the plane of sky. The foreshortening of baselines relative to the ground plane occurs entirely in the y -direction, meaning that quasi-horizontal structures (i.e. those with negligible vertical extent) there will be an apparent anisotropy with its major axis aligned with the x -direction.

Given that the fit using equation (H.2) is specifically to the peak of the autocorrelation, and not a representation of its full structure, it is necessary to define a condition for which points are included in the fit. There is no clear means of determining an ‘objectively reasonable’ definition of which points constitute the peak, and so it is important to identify any effect the choice of definition has on the resulting estimates of velocity and anisotropy. While a restrictive definition of the peak will ensure that no

secondary peaks or other significant deviations from the assumed model are included, it may also provide insufficient datapoints to adequately constrain the fit parameters.

To test the effect of the choice of peak definition, the fitting process is here repeated with a range of different criteria, ranging from strict to more relaxed. As well as identifying the influence of this choice on the derived values, the variation in estimated velocities and anisotropies provides a possible means of estimating their uncertainties. The peak definition is specified initially by the range of time delays which are included. The strictest value chosen for this is $|\tau| \leq 0.5$ s and the most relaxed is $|\tau| \leq 2.5$ s, stepping between these with intervals of 0.5 s. At each time step τ_i , the maximum cross-correlation $C_{max}(\vec{\zeta}, \tau_i)$ is calculated, and only baselines satisfying

$$C(\vec{\zeta}, \tau_i) \geq \max\left(0.75 \times C_{max}(\vec{\zeta}, \tau_i), 0.5\right) \quad (\text{H.5})$$

are included. The floor value of 0.5 is included to ensure no spurious correlations can be included in the fitting process, particularly in cases where temporal decorrelation is extremely rapid. A further limitation of excluding baselines longer than 10 km in the plane of sky was included to avoid issues with secondary cross-correlation peaks being included in the fitting in some cases.

Correlations between all combinations of core and remote stations were calculated (although in practice due to the limitation on baseline length some of these would never be included in the fitting process). The correlations are calculated between individual frequency channels, with 20 channels selected evenly spanning the observing frequency range to give a representative sample while minimising computational requirements. Rather than filtering all dynamic spectra for RFI as described in Section 4.2, to reduce computational expense the single channel data is low-pass filtered with a cutoff frequency of 0.5 Hz to mitigate the impact of RFI and instrumental noise, as was done

by Fallows et al. (2020). The cutoff frequency of 0.5 Hz is also roughly consistent with the estimated frequencies at which noise dominates the signal from Section 8.3.2, which ranged between $\sim 0.2 - 1$ Hz. The time intervals used to calculate the cross-correlations are also identical to those used in the power spectrum analysis in Section 8.3.2 to enable direct comparison of the results.




Figure H.1. The spatial distribution of cross-correlations at zero time lag for the first event discussed in Chapter 8 (1st April 2022) at 21:00 UT. The left panel shows all baselines, while the right panel shows a zoomed in section to highlight the anisotropic structure. Both panels use an equal aspect ratio to highlight the foreshortening of baselines along the y -direction in the plane of sky. The right panel also includes the magnetic field orientation (red line).

One key limitation that arises in this analysis is the limited baseline coverage especially due to the foreshortening of baselines in the y direction. Figures H.1 and H.2 illustrate this, and the consequent difficulty in quantifying the degree of anisotropy (anisotropy ratio) for the observed U-shape features. In both cases, the left panel shows the full distribution of baselines in the plane of sky coloured according to the




Figure H.2. The spatial distribution of cross-correlations at zero time lag for the second event discussed in Chapter 8 (21st January 2023) at 02:16 UT. The left panel shows all baselines, while the right panel shows a zoomed in section to highlight the anisotropic structure. Both panels use an equal aspect ratio to highlight the foreshortening of baselines along the y -direction in the plane of sky. The right panel also includes the magnetic field orientation (red line).

cross-correlation at zero time lag, while the left panel shows a zoomed in version of this concentrating on baselines of less than 10 km. The enhanced correlations aligned with the magnetic field (red line) extend across the full range of available baselines without any discernible decay, meaning the correlation length along the magnetic field direction is unconstrained by these observations.

With the existing LOFAR network, the only possibility to resolve this would be to include one or more international stations in the cross-correlation analysis which provide long baselines in the foreshortened direction. The most suitable candidate for this would likely be station FR606 in France (47.38°N , 2.19°E). However, given

that both the baselines and magnetic field projection rotate in the plane of sky as the sources apparent position evolves, in the best case scenario this would only give baselines oriented along the magnetic field at a very limited range of times for a given event. Also, due to the lower latitude of FR606 Cygnus A actually passes below the horizon briefly as it passes north, making U-shape scintillation events less commonly observed from this station than the stations in the Netherlands and at similar or more northerly latitudes across Europe. Even if a U-shape was observed by FR606, it is quite possible that the scintillation would be originating from a different geographic region and not along the same field lines as the scintillation observed from the Netherlands. This would preclude using the cross-correlations of the signals from this station to the Netherlands as a measure of the extent of the irregularities along a given field line.

Bibliography

- Aarons, J. (1973). A descriptive model of F layer high-latitude irregularities as shown by scintillation observations. *Journal of Geophysical Research*, 78(31):7441–7450.
- Aarons, J. (1982). Global morphology of ionospheric scintillations. *Proceedings of the IEEE*, 70(4):360–378.
- Afraimovich, E., Edemskiy, I., Voeykov, S., Yasyukevich, Y. V., and Zhivetiev, I. (2009). Spatio-temporal structure of the wave packets generated by the solar terminator. *Advances in Space Research*, 44(7):824–835.
- Afraimovich, E., Ishin, A., Tinin, M., Yasyukevich, Y., and Jin, S. (2011). First evidence of anisotropy of GPS phase slips caused by the mid-latitude field-aligned ionospheric irregularities. *Advances in Space Research*, 47(10):1674–1680. GNSS Remote Sensing-2.
- Alken, P., Thébaud, E., Beggan, C. D., Amit, H., Aubert, J., Baerenzung, J., Bondar, T. N., Brown, W. J., Califf, S., Chambodut, A., Chulliat, A., Cox, G. A., Finlay, C. C., Fournier, A., Gillet, N., Grayver, A., Hammer, M. D., Holschneider, M., Huder, L., Hulot, G., Jager, T., Kloss, C., Korte, M., Kuang, W., Kuvshinov, A., Langlais, B., Léger, J.-M., Lesur, V., Livermore, P. W., Lowes, F. J., Macmillan, S., Magnes, W., Manda, M., Marsal, S., Matzka, J., Metman, M. C., Minami, T., Morschhauser, A., Mound, J. E., Nair, M., Nakano, S., Olsen, N., Pavón-Carrasco, F. J., Petrov, V. G., Ropp, G., Rother, M., Sabaka, T. J., Sanchez, S., Saturnino, D., Schnepf, N. R., Shen, X., Stolle, C., Tangborn, A., Tøffner-Clausen, L., Toh, H., Torta, J. M., Varner, J., Vervelidou, F., Vigneron, P., Wardinski, I., Wicht, J., Woods, A., Yang, Y., Zeren, Z., and Zhou, B. (2021). International Geomagnetic Reference Field: the thirteenth generation. *Earth, Planets and Space*, 73(1):1–25.
- Arikan, F., Nayir, H., Sezen, U., and Arikan, O. (2008). Estimation of single station interfrequency receiver bias using GPS-TEC. *Radio Science*, 43(4).
- ASTRON (2023). Major upgrade of International LOFAR Telescope approved. <https://www.astron.nl/major-upgrade-of-international-lofar-telescope-approved/>. [Online; accessed 31/01/2025].
- Auchère, F., Froment, C., Bocchialini, K., Buchlin, E., and Solomon, J. (2016). On the Fourier and wavelet analysis of coronal time series. *The Astrophysical Journal*, 825(2):110.
- Baars, J. W. M. and Kärcher, H. J. (2017). Seventy years of radio telescope design and construction. *URSI Radio Science Bulletin*, 2017(362):15–38.
- Baker, J. B., Clauer, C. R., Ridley, A. J., Papitashvili, V. O., Brittnacher, M. J., and Newell, P. T. (2000). The nightside poleward boundary of the auroral oval as seen by DMSP and the Ultraviolet Imager. *Journal of Geophysical Research: Space Physics*, 105(A9):21267–21280.

- Banville, S. and Langley, R. B. (2013). Mitigating the impact of ionospheric cycle slips in GNSS observations. *Journal of Geodesy*, 87(2):179–193.
- Baskaradas, J. A., Bianchi, S., Pietrella, M., Pezzopane, M., Sciacca, U., and Zuccheretti, E. (2014). Description of ionospheric disturbances observed by vertical ionospheric sounding at 3 MHz. *Annals of Geophysics*, 1:A0187.
- Basu, S., Basu, S., MacKenzie, E., Coley, W. R., Sharber, J. R., and Hoegy, W. R. (1990). Plasma structuring by the gradient drift instability at high latitudes and comparison with velocity shear driven processes. *Journal of Geophysical Research: Space Physics*, 95(A6):7799–7818.
- Basu, S., Groves, K., Basu, S., and Sultan, P. (2002). Specification and forecasting of scintillations in communication/navigation links: current status and future plans. *Journal of Atmospheric and Solar-Terrestrial Physics*, 64(16):1745–1754.
- Basu, S., MacKenzie, E., Basu, S., Carlson, H. C., Hardy, D. A., Rich, F. J., and Livingston, R. C. (1983). Coordinated measurements of low-energy electron precipitation and scintillations/TEC in the auroral oval. *Radio Science*, 18(6):1151–1165.
- Basu, S., Weber, E., Bullett, T., Keskinen, M., MacKenzie, E., Doherty, P., Sheehan, R., Kuenzler, H., Ning, P., and Bongiolatti, J. (1998). Characteristics of plasma structuring in the cusp/cleft region at Svalbard. *Radio Science*, 33(6):1885–1899.
- Berngardt, O. I., Zolotukhina, N. A., and Oinats, A. V. (2015). Observations of field-aligned ionospheric irregularities during quiet and disturbed conditions with EKB radar: first results. *Earth, Planets and Space*, 67:1–13.
- Beser, K., Mevius, M., Grzesiak, M., and Rothkaehl, H. (2022). Detection of periodic disturbances in LOFAR calibration solutions. *Remote Sensing*, 14(7):1719.
- Bianchi, C., Baskaradas, J., Pezzopane, M., Pietrella, M., Sciacca, U., and Zuccheretti, E. (2013). Fading in the HF ionospheric channel and the role of irregularities. *Advances in Space Research*, 52(3):403–411.
- Bibl, K. and Reinisch, B. W. (1978). The universal digital ionosonde. *Radio Science*, 13(3):519–530.
- Błaszkiwicz, L. P., Flisek, P., Kotulak, K., Krankowski, A., Lewandowski, W., Kijak, J., and Froń, A. (2020). Finding the ionospheric fluctuations reflection in the pulsar signals’ characteristics observed with LOFAR. *Sensors*, 21(1):51.
- Bolton, J. G. and Stanley, G. J. (1948). Variable source of radio frequency radiation in the constellation of Cygnus. *Nature*, pages 312–313.
- Boström, R. (1964). A model of the auroral electrojets. *Journal of Geophysical Research*, 69(23):4983–4999.
- Boyde, B., Wood, A., Dorrian, G., Fallows, R. A., Themens, D., Mielich, J., Elvidge, S., Mevius, M., Zucca, P., Dabrowski, B., Krankowski, A., Vocks, C., and Bisi, M. (2022). Lensing from small-scale travelling ionospheric disturbances observed using LOFAR. *Journal of Space Weather and Space Climate*, 12:34.
- Boyde, B., Wood, A., Dorrian, G., Sweijen, F., de Gasperin, F., Mevius, M., Beser, K., and Themens, D. (2024). Wavelet analysis of differential TEC measurements obtained using LOFAR. *Radio Science*, 59(4):e2023RS007871.

- Braden, E. J. (2022). A statistical analysis of the Fresnel frequency of ionospheric scintillation in the high latitude region using a piecewise fitting technique. Master’s thesis, University of New Brunswick.
- Briggs, B. and Parkin, I. (1963). On the variation of radio star and satellite scintillations with zenith angle. *Journal of Atmospheric and Terrestrial Physics*, 25(6):339–366.
- Briskin, W. F., Macquart, J.-P., Gao, J.-J., Rickett, B., Coles, W., Deller, A., Tingay, S., and West, C. (2009). 100 μ s resolution VLBI imaging of anisotropic interstellar scattering toward pulsar B0834+ 06. *The Astrophysical Journal*, 708(1):232.
- Buneman, O. (1963). Excitation of field aligned sound waves by electron streams. *Physical Review Letters*, 10:285–287.
- Buss, S., Hertzog, A., Hostettler, C., Bui, T. B., Lüthi, D., and Wernli, H. (2004). Analysis of a jet stream induced gravity wave associated with an observed ice cloud over Greenland. *Atmospheric Chemistry and Physics*, 4(5):1183–1200.
- Calais, E. and Minster, J. B. (1998). GPS, earthquakes, the ionosphere, and the Space Shuttle. *Physics of the Earth and Planetary Interiors*, 105(3-4):167–181.
- Cannon, P. S. et al. (2013). Extreme space weather: Impacts on engineered systems. Technical report, Royal Academy of Engineering, London, UK.
- Carbone, D., van der Horst, A. J., Wijers, R. A. M. J., Swinbank, J. D., Rowlinson, A., Broderick, J. W., Cendes, Y. N., Stewart, A. J., Bell, M. E., Breton, R. P., Corbel, S., Eisloffel, J., Fender, R. P., Grießmeier, J.-M., Hessels, J. W. T., Jonker, P., Kramer, M., Law, C. J., Miller-Jones, J. C. A., Pietka, M., Scheers, L. H. A., Stappers, B. W., van Leeuwen, J., Wijnands, R., Wise, M., and Zarka, P. (2016). New methods to constrain the radio transient rate: results from a survey of four fields with LOFAR. *Monthly Notices of the Royal Astronomical Society*, 459(3):3161–3174.
- Carilli, C., Perley, R., Dreher, J., and Leahy, J. (1991). Multifrequency radio observations of Cygnus A-spectral aging in powerful radio galaxies. *The Astrophysical Journal*, 383:554–573.
- Carlson Jr, H. C., Oksavik, K., Moen, J., and Pedersen, T. (2004). Ionospheric patch formation: Direct measurements of the origin of a polar cap patch. *Geophysical Research Letters*, 31(8).
- Carrano, C. S., Retterer, J. M., Groves, K. M., Crowley, G., Duly, T. M., and Hunton, D. E. (2020). Wave-optics analysis of HF propagation through traveling ionospheric disturbances and developing plasma bubbles. In *2020 XXXIIIrd General Assembly and Scientific Symposium of the International Union of Radio Science*, pages 1–4. IEEE.
- Carrano, C. S. and Rino, C. L. (2016). A theory of scintillation for two-component power law irregularity spectra: Overview and numerical results. *Radio Science*, 51(6):789–813.
- Carrano, C. S., Valladares, C. E., and Groves, K. M. (2012). Latitudinal and local time variation of ionospheric turbulence parameters during the conjugate point equatorial experiment in Brazil. *International Journal of Geophysics*, 2012(1):103963.
- Chao-Song, H. and Jun, L. (1991). Weak nonlinear theory of the ionospheric response to atmospheric gravity waves in the F-region. *Journal of Atmospheric and Terrestrial Physics*, 53(10):903–908.
- Chapman, S. (1931). The absorption and dissociative or ionizing effect of monochromatic radiation in an atmosphere on a rotating Earth. *Proceedings of the Physical Society*, 43(1):26.

- Chen, J., Ren, X., Zhang, X., Zhang, J., and Huang, L. (2020). Assessment and validation of three ionospheric models (IRI-2016, NeQuick2, and IGS-GIM) from 2002 to 2018. *Space Weather*, 18(6):e2019SW002422.
- Chisham, G., Lester, M., Milan, S. E., Freeman, M. P., Bristow, W. A., Grocott, A., McWilliams, K. A., Ruohoniemi, J. M., Yeoman, T. K., Dyson, P. L., Greenwald, R. A., Kikuchi, T., Pinnock, M., Rash, J. P. S., Sato, N., Sofko, G. J., Villain, J.-P., and Walker, A. D. M. (2007). A decade of the Super Dual Auroral Radar Network (SuperDARN): Scientific achievements, new techniques and future directions. *Surveys in Geophysics*, 28:33–109.
- Chum, J., Podolská, K., Rusz, J., Baše, J., and Tedoradze, N. (2021). Statistical investigation of gravity wave characteristics in the ionosphere. *Earth, Planets and Space*, 73:1–16.
- Chum, J., Šindelářová, T., Laštovička, J., Hruška, F., Burešová, D., and Baše, J. (2010). Horizontal velocities and propagation directions of gravity waves in the ionosphere over the Czech Republic. *Journal of Geophysical Research: Space Physics*, 115(A11).
- Clegg, A. W., Fey, A. L., and Lazio, T. J. W. (1998). The Gaussian plasma lens in astrophysics: refraction. *The Astrophysical Journal*, 496(1):253.
- Cohen, A. S. and Röttgering, H. J. A. (2009). Probing fine-scale ionospheric structure with the Very Large Array radio telescope. *The Astronomical Journal*, 138(2):439.
- Coker, C., Thonnard, S. E., Dymond, K. F., Lazio, T. J. W., Makela, J. J., and Loughmiller, P. J. (2009). Simultaneous radio interferometer and optical observations of ionospheric structure at the Very Large Array. *Radio Science*, 44(1).
- Cordes, J. M., Rickett, B. J., Stinebring, D. R., and Coles, W. A. (2006). Theory of parabolic arcs in interstellar scintillation spectra. *The Astrophysical Journal*, 637(1):346.
- Cornwell, T. J. and Wilkinson, P. N. (1981). A new method for making maps with unstable radio interferometers. *Monthly Notices of the Royal Astronomical Society*, 196(4):1067–1086.
- Cosgrove, R. B. and Tsunoda, R. T. (2002). A direction-dependent instability of sporadic-E layers in the nighttime midlatitude ionosphere. *Geophysical Research Letters*, 29(18):11–1–11–4.
- Cosgrove, R. B. and Tsunoda, R. T. (2004). Instability of the E-F coupled nighttime midlatitude ionosphere. *Journal of Geophysical Research: Space Physics*, 109(A4).
- Cosgrove, R. B., Tsunoda, R. T., Fukao, S., and Yamamoto, M. (2004). Coupling of the Perkins instability and the sporadic E layer instability derived from physical arguments. *Journal of Geophysical Research: Space Physics*, 109(A6).
- Coster, A. J., Goncharenko, L., Zhang, S.-R., Erickson, P. J., Rideout, W., and Vierinen, J. (2017). GNSS observations of ionospheric variations during the 21 August 2017 solar eclipse. *Geophysical Research Letters*, 44(24):12–41.
- Cowling, D., Webb, H., and Yeh, K. (1971). Group rays of internal gravity waves in a wind-stratified atmosphere. *Journal of Geophysical Research*, 76(1):213–220.
- Crowley, G., Jones, T., and Dudeney, J. (1987). Comparison of short period TID morphologies in Antarctica during geomagnetically quiet and active intervals. *Journal of Atmospheric and Terrestrial Physics*, 49(11):1155–1162.

- Crowley, G. and Rodrigues, F. S. (2012). Characteristics of traveling ionospheric disturbances observed by the TIDDBIT sounder. *Radio Science*, 47(4).
- Dagg, M. (1957). Radio-star ridges. *Journal of Atmospheric and Terrestrial Physics*, 11(2):118–127.
- de Gasperin, F., Brunetti, G., Brügger, M., van Weeren, R., Williams, W. L., Botteon, A., Cuciti, V., Dijkema, T. J., Edler, H., Iacobelli, M., Kang, H., Offringa, A., Orrú, E., Pizzo, R., Rafferty, D., Röttgering, H., and Shimwell, T. (2020a). Reaching thermal noise at ultra-low radio frequencies - Toothbrush radio relic downstream of the shock front. *Astronomy & Astrophysics*, 642:A85.
- de Gasperin, F., Dijkema, T. J., Drabent, A., Mevius, M., Rafferty, D., van Weeren, R., Brügger, M., Callingham, J. R., Emig, K. L., Heald, G., Intema, H. T., Morabito, L. K., Offringa, A. R., Oonk, R., Orrú, E., Röttgering, H., Sabater, J., Shimwell, T., Shulevski, A., and Williams, W. (2019). Systematic effects in LOFAR data: A unified calibration strategy. *Astronomy & Astrophysics*, 622:A5.
- de Gasperin, F., Mevius, M., Rafferty, D., Intema, H., and Fallows, R. (2018). The effect of the ionosphere on ultra-low-frequency radio-interferometric observations. *Astronomy & Astrophysics*, 615:A179.
- de Gasperin, F., Vink, J., McKean, J. P., Asgekar, A., Avruch, I., Bentum, M. J., Blaauw, R., Bonafede, A., Broderick, J. W., Brügger, M., Breitling, F., Brouw, W. N., Butcher, H. R., Ciardi, B., Cuciti, V., de Vos, M., Duscha, S., Eislöffel, J., Engels, D., Fallows, R. A., Franzen, T. M. O., Garrett, M. A., Gunst, A. W., Hörandel, J., Heald, G., Hoeft, M., Iacobelli, M., Koopmans, L. V. E., Krankowski, A., Maat, P., Mann, G., Mevius, M., Miley, G., Morganti, R., Nelles, A., Norden, M. J., Offringa, A. R., Orrú, E., Paas, H., Pandey, V. N., Pandey-Pommier, M., Pekal, R., Pizzo, R., Reich, W., Rowlinson, A., Rottgering, H. J. A., Schwarz, D. J., Shulevski, A., Smirnov, O., Sobey, C., Soida, M., Steinmetz, M., Tagger, M., Toribio, M. C., van Ardenne, A., van der Horst, A. J., van Haarlem, M. P., van Weeren, R. J., Vocks, C., Wucknitz, O., Zarka, P., and Zucca, P. (2020b). Cassiopeia A, Cygnus A, Taurus A, and Virgo A at ultra-low radio frequencies. *Astronomy & Astrophysics*, 635:A150.
- de Gasperin, F., Williams, W. L., Best, P., Brügger, M., Brunetti, G., Cuciti, V., Dijkema, T. J., Hardcastle, M. J., Norden, M. J., Offringa, A., Shimwell, T., van Weeren, R., Bomans, D., Bonafede, A., Botteon, A., Callingham, J. R., Cassano, R., Chyzy, K. T., Emig, K. L., Edler, H., Haverkorn, M., Heald, G., Heesen, V., Iacobelli, M., Intema, H. T., Kadler, M., Malek, K., Mevius, M., Miley, G., Mingo, B., Morabito, L. K., Sabater, J., Morganti, R., Orrú, E., Pizzo, R., Prandoni, I., Shulevski, A., Tasse, C., Vaccari, M., Zarka, P., and Röttgering, H. (2021). The LOFAR LBA sky survey-I. survey description and preliminary data release. *Astronomy & Astrophysics*, 648:A104.
- de Larquier, S., Eltrass, A., Mahmoudian, A., Ruohoniemi, J. M., Baker, J. B. H., Scales, W. A., Erickson, P. J., and Greenwald, R. A. (2014). Investigation of the temperature gradient instability as the source of midlatitude quiet time decameter-scale ionospheric irregularities: 1. observations. *Journal of Geophysical Research: Space Physics*, 119(6):4872–4881.
- Dewdney, P. E., Hall, P. J., Schilizzi, R. T., and Lazio, T. J. L. W. (2009). The Square Kilometre Array. *Proceedings of the IEEE*, 97(8):1482–1496.
- Ding, F., Wan, W., Li, Q., Zhang, R., Song, Q., Ning, B., Liu, L., Zhao, B., and Xiong, B. (2014). Comparative climatological study of large-scale traveling ionospheric disturbances over North America and China in 2011–2012. *Journal of Geophysical Research: Space Physics*, 119(1):519–529.

- Ding, F., Wan, W., Ning, B., and Wang, M. (2007). Large-scale traveling ionospheric disturbances observed by GPS total electron content during the magnetic storm of 29-30 October 2003. *Journal of Geophysical Research: Space Physics*, 112(A6).
- Ding, F., Wan, W., Xu, G., Yu, T., Yang, G., and Wang, J.-S. (2011). Climatology of medium-scale traveling ionospheric disturbances observed by a GPS network in central China. *Journal of Geophysical Research: Space Physics*, 116(A9).
- Ding, F., Yuan, H., Wan, W., Reid, I. M., and Wothke, J. M. (2004). Occurrence characteristics of medium-scale gravity waves observed in OH and OI nightglow over Adelaide (34.5°S, 138.5°E). *Journal of Geophysical Research: Atmospheres*, 109(D14).
- Ding, M., Tong, P., Wei, Y., and Yu, L. (2021). Multiple phase screen modeling of HF wave field scintillations caused by the irregularities in inhomogeneous media. *Radio Science*, 56(4):e2020RS007239.
- Dorrian, G., Fallows, R., Wood, A., Themens, D. R., Boyde, B., Krankowski, A., Bisi, M., Dabrowski, B., and Vocks, C. (2023). LOFAR observations of substructure within a traveling ionospheric disturbance at mid-latitude. *Space Weather*, 21(1):e2022SW003198.
- Drob, D. P., Emmert, J. T., Meriwether, J. W., Makela, J. J., Doornbos, E., Conde, M., Hernandez, G., Noto, J., Zawdie, K. A., McDonald, S. E., Huba, J. D., and Klenzing, J. H. (2015). An update to the Horizontal Wind Model (HWM): The quiet time thermosphere. *Earth and Space Science*, 2(7):301–319.
- Dungey, J. W. (1961). Interplanetary magnetic field and the auroral zones. *Physical Review Letters*, 6(2):47.
- Ecklund, W. L., Carter, D. A., and Balsley, B. B. (1981). Gradient drift irregularities in mid-latitude sporadic E. *Journal of Geophysical Research: Space Physics*, 86(A2):858–862.
- Edler, H. W., de Gasperin, F., and Rafferty, D. (2021). Investigating ionospheric calibration for LOFAR 2.0 with simulated observations. *Astronomy & Astrophysics*, 652:A37.
- Eltrass, A., Scales, W., Erickson, P., Ruohoniemi, J., and Baker, J. (2016). Investigation of the role of plasma wave cascading processes in the formation of midlatitude irregularities utilizing GPS and radar observations. *Radio Science*, 51(6):836–851.
- Eltrass, A. and Scales, W. A. (2014). Nonlinear evolution of the temperature gradient instability in the midlatitude ionosphere. *Journal of Geophysical Research: Space Physics*, 119(9):7889–7901.
- Elvidge, S. and Angling, M. J. (2019). Using the local ensemble transform Kalman filter for upper atmospheric modelling. *Journal of Space Weather and Space Climate*, 9:A30.
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD-96*, pages 226–231.
- Fallows, R., Bisi, M. M., Forte, B., Ulich, T., Konovalenko, A., Mann, G., and Vocks, C. (2016). Separating nightside interplanetary and ionospheric scintillation with LOFAR. *The Astrophysical Journal Letters*, 828(1):L7.

- Fallows, R. A., Coles, W. A., McKay-Bukowski, D., Vierinen, J., Virtanen, I. I., Postila, M., Ulich, T., Enell, C.-F., Kero, A., Iinatti, T., Lehtinen, M., Orispää, M., Raita, T., Roininen, L., Turunen, E., Brentjens, M., Ebbendorf, N., Gerbers, M., Grit, T., Gruppen, P., Meulman, H., Norden, M. J., de Reijer, J.-P., Schoenmakers, A., and Stuurwold, K. (2014). Broadband meter-wavelength observations of ionospheric scintillation. *Journal of Geophysical Research: Space Physics*, 119(12):10544–10560.
- Fallows, R. A., Forte, B., Astin, I., Allbrook, T., Arnold, A., Wood, A., Dorrian, G., Mevius, M., Rothkaehl, H., Matyjasiak, B., Krankowski, A., Anderson, J. M., Asgekar, A., Avruch, I. M., Bentum, M., Bisi, M. M., Butcher, H. R., Ciardi, B., Dabrowski, B., Damstra, S., de Gasperin, F., Duschka, S., Eislöffel, J., Franzen, T. M., Garrett, M. A., Griesmeier, J.-M., Gunst, A. W., Hoeft, M., Hörandel, J. R., Iacobelli, M., Intema, H. T., Koopmans, L. V., Maat, P., Mann, G., Nelles, A., Paas, H., Pandey, V. N., Reich, W., Rowlinson, A., Ruiter, M., Schwarz, D. J., Serylak, M., Shulevski, A., Smirnov, O. M., Soida, M., Steinmetz, M., Thoudam, S., Toribio, M. C., van Ardenne, A., van Bommel, I. M., van der Wiel, M. H., van Haarlem, M. P., Vermeulen, R. C., Vocks, C., Wijers, R. A., Wucknitz, O., Zarka, P., and Zucca, P. (2020). A LOFAR observation of ionospheric scintillation from two simultaneous travelling ionospheric disturbances. *Journal of Space Weather and Space Climate*, 10:10.
- Fallows, R. A., Iwai, K., Jackson, B. V., Zhang, P., Bisi, M. M., and Zucca, P. (2023). Application of novel interplanetary scintillation visualisations using LOFAR: A case study of merged CMEs from September 2017. *Advances in Space Research*, 72:5311–5327.
- Farley, D. T. (1963). A plasma instability resulting in field-aligned irregularities in the ionosphere. *Journal of Geophysical Research*, 68(22):6083–6097.
- Fedorenko, A. K. and Kryuchkov, E. I. (2011). Distribution of medium-scale acoustic gravity waves in polar regions according to satellite measurement data. *Geomagnetism and Aeronomy*, 51:520–533.
- Fedorenko, Y. P., Fedorenko, V., and Lysenko, V. (2011). Parameters of the medium-scale traveling ionospheric disturbances model deduced from measurements. *Geomagnetism and Aeronomy*, 51(1):88–104.
- Fejer, B. G. and Kelley, M. C. (1980). Ionospheric irregularities. *Reviews of Geophysics*, 18(2):401–454.
- Fišer, J., Chum, J., and Liu, J.-Y. (2017). Medium-scale traveling ionospheric disturbances over Taiwan observed with HF Doppler sounding. *Earth, Planets and Space*, 69:1–10.
- Flisek, P., Forte, B., Fallows, R., Kotulak, K., Krankowski, A., Bisi, M., Mevius, M., Froń, A., Tiburzi, C., Soida, M., Śmierciak, B., Grzesiak, M., Matyjasiak, B., Pożoga, M., Dąbrowski, B., Mann, G., Vocks, C., Zucca, P., and Błaszczewicz, L. (2023). Towards the possibility to combine LOFAR and GNSS measurements to sense ionospheric irregularities. *Journal of Space Weather and Space Climate*, 13:27.
- Flock, W. L. (1987). *Propagation effects on satellite systems at frequencies below 10 GHz: A handbook for satellite systems design*. NASA Reference Publication.
- Forbes, J. M. (2007). Dynamics of the thermosphere. *Journal of the Meteorological Society of Japan*, 85:193–213.
- Forte, B. (2008). Refractive scattering evidence from multifrequency scintillation spectra observed at auroral latitudes. *Radio Science*, 43(2).

- Forte, B., Fallows, R. A., Bisi, M. M., Zhang, J., Krankowski, A., Dabrowski, B., Rothkaehl, H., and Vocks, C. (2022). Interpretation of radio wave scintillation observed through LOFAR radio telescopes. *The Astrophysical Journal Supplement Series*, 263(2):36.
- Fox, J. L. (1993). Dissociative recombination in planetary ionospheres. In *Dissociative Recombination: Theory, Experiment, and Applications*, pages 219–242. Springer.
- Friedman, J. P. (1966). Propagation of internal gravity waves in a thermally stratified atmosphere. *Journal of Geophysical Research*, 71(4):1033–1054.
- Friis-Christensen, E., Lühr, H., Knudsen, D., and Haagmans, R. (2008). Swarm – an Earth observation mission investigating geospace. *Advances in Space Research*, 41(1):210–216.
- Frissell, N. A., Baker, J., Ruohoniemi, J. M., Gerrard, A. J., Miller, E. S., Marini, J. P., West, M. L., and Bristow, W. A. (2014). Climatology of medium-scale traveling ionospheric disturbances observed by the midlatitude Blackstone SuperDARN radar. *Journal of Geophysical Research: Space Physics*, 119(9):7679–7697.
- Frissell, N. A., Baker, J. B. H., Ruohoniemi, J. M., Greenwald, R. A., Gerrard, A. J., Miller, E. S., and West, M. L. (2016). Sources and characteristics of medium-scale traveling ionospheric disturbances observed by high-frequency radars in the North American sector. *Journal of Geophysical Research: Space Physics*, 121(4):3722–3739.
- Fritts, D. C. and Alexander, M. J. (2003). Gravity wave dynamics and effects in the middle atmosphere. *Reviews of Geophysics*, 41(1).
- Fukao, S., Kelley, M. C., Shirakawa, T., Takami, T., Yamamoto, M., Tsuda, T., and Kato, S. (1991). Turbulent upwelling of the mid-latitude ionosphere: 1. observational results by the MU radar. *Journal of Geophysical Research: Space Physics*, 96(A3):3725–3746.
- Galushko, V., Paznukhov, V., Yampolski, Y., and Foster, J. (1998). Incoherent scatter radar observations of AGW/TID events generated by the moving solar terminator. *Annales Geophysicae*, 16(7):821–827.
- Garcia, F. J., Kelley, M. C., Makela, J. J., and Huang, C.-S. (2000). Airglow observations of mesoscale low-velocity traveling ionospheric disturbances at midlatitudes. *Journal of Geophysical Research: Space Physics*, 105(A8):18407–18415.
- Gebre-Egziabher, D. and Gleason, S. (2009). *GNSS applications and methods*. Artech House.
- Gjerloev, J. W. (2012). The SuperMAG data processing technique. *Journal of Geophysical Research: Space Physics*, 117(A9).
- Gochelashvily, K. and Shishov, V. (1971). Laser beam scintillation beyond a turbulent layer. *Optica Acta: International Journal of Optics*, 18(4):313–320.
- Goodman, J., Ballard, J., and Sharp, E. (1997). A long-term investigation of the HF communication channel over middle-and high-latitude paths. *Radio Science*, 32(4):1705–1715.
- Greenwald, R. A., Baker, K. B., Dudeney, J. R., Pinnock, M., Jones, T. B., Thomas, E. C., Villain, J. P., Cerisier, J. C., Senior, C., Hanuise, C., Hunsucker, R. D., Sofko, G., Koehler, J., Nielsen, E., Pellinen, R., Walker, A. D. M., Sato, N., and Yamagishi, H. (1995). DARN/SuperDARN: A global view of the dynamics of high-latitude convection. *Space Science Reviews*, 71(1-4):761–796.

- Greenwald, R. A., Oksavik, K., Erickson, P. J., Lind, F. D., Ruohoniemi, J. M., Baker, J. B. H., and Gjerloev, J. W. (2006). Identification of the temperature gradient instability as the source of decameter-scale ionospheric irregularities on plasmopause field lines. *Geophysical Research Letters*, 33(18).
- Grzesiak, M., Pożoga, M., Matyjasiak, B., Przepiórka, D., Beser, K., Tomasik, L., Rothkaehl, H., and Ciechowska, H. (2022). Determining ionospheric drift and anisotropy of irregularities from LOFAR core measurements: Testing hypotheses behind estimation. *Remote Sensing*, 14(18):4655.
- Habarulema, J. B., Katamzi, Z. T., and McKinnell, L.-A. (2013). Estimating the propagation characteristics of large-scale traveling ionospheric disturbances using ground-based and satellite data. *Journal of Geophysical Research: Space Physics*, 118(12):7768–7782.
- Haldoupis, C. (2012). Midlatitude sporadic E: a typical paradigm of atmosphere-ionosphere coupling. *Space Science Reviews*, 168:441–461.
- Hamza, A. M. (1999). Perkins instability revisited. *Journal of Geophysical Research: Space Physics*, 104(A10):22567–22575.
- Hargreaves, J. K. (1992). *The solar-terrestrial environment: an introduction to geospace-the science of the terrestrial upper atmosphere, ionosphere, and magnetosphere*. Cambridge University Press.
- Harris, F. J. (1978). On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1):51–83.
- Hathaway, D. H. (2015). The solar cycle. *Living Reviews in Solar Physics*, 12:1–87.
- Headrick, J. M. and Skolnik, M. I. (1974). Over-the-horizon radar in the HF band. *Proceedings of the IEEE*, 62(6):664–673.
- Heald, G. H., Pizzo, R. F., Orrú, E., Breton, R. P., Carbone, D., Ferrari, C., Hardcastle, M. J., Jurusik, W., Macario, G., Mulcahy, D., Rafferty, D., Asgekar, A., Brentjens, M., Fallows, R. A., Frieswijk, W., Toribio, M. C., Adebahr, B., Arts, M., Bell, M. R., Bonafede, A., Bray, J., Broderick, J., Cantwell, T., Carroll, P., Cendes, Y., Clarke, A. O., Croston, J., Daiboo, S., de Gasperin, F., Gregson, J., Harwood, J., Hassall, T., Heesen, V., Horneffer, A., van der Horst, A. J., Iacobelli, M., Jelić, V., Jones, D., Kant, D., Kokotanekov, G., Martin, P., McKean, J. P., Morabito, L. K., Nikiel-Wroczyński, B., Offringa, A., Pandey, V. N., Pandey-Pommier, M., Pietka, M., Pratley, L., Riseley, C., Rowlinson, A., Sabater, J., Scaife, A. M. M., Scheers, L. H. A., Sendlinger, K., Shulevski, A., Sipior, M., Sobey, C., Stewart, A. J., Stroe, A., Swinbank, J., Tasse, C., Trüstedt, J., Varenius, E., van Velzen, S., Vilchez, N., van Weeren, R. J., Wijnholds, S., Williams, W. L., de Bruyn, A. G., Nijboer, R., Wise, M., Alexov, A., Anderson, J., Avruch, I. M., Beck, R., Bell, M. E., van Bemmell, I., Bentum, M. J., Bernardi, G., Best, P., Breitling, F., Brouw, W. N., Brüggen, M., Butcher, H. R., Ciardi, B., Conway, J. E., de Geus, E., de Jong, A., de Vos, M., Deller, A., Dettmar, R.-J., Duscha, S., Eislöffel, J., Engels, D., Falcke, H., Fender, R., Garrett, M. A., Grießmeier, J., Gunst, A. W., Hamaker, J. P., Hessels, J. W. T., Hoeft, M., Hörandel, J., Holties, H. A., Intema, H., Jackson, N. J., Jütte, E., Karastergiou, A., Klijn, W. F. A., Kondratiev, V. I., Koopmans, L. V. E., Kuniyoshi, M., Kuper, G., Law, C., van Leeuwen, J., Loose, M., Maat, P., Markoff, S., McFadden, R., McKay-Bukowski, D., Mevius, M., Miller-Jones, J. C. A., Morganti, R., Munk, H., Nelles, A., Noordam, J. E., Norden, M. J., Paas, H., Polatidis, A. G., Reich, W., Renting, A., Röttgering, H., Schoenmakers, A., Schwarz, D., Sluman, J., Smirnov, O., Stappers, B. W., Steinmetz, M., Tagger, M., Tang, Y., ter Veen, S., Thoudam, S., Vermeulen, R., Vocks, C., Vogt, C., Wijers, R. A. M. J., Wucknitz, O., Yatawatta, S., and Zarka, P. (2015). The LOFAR multifrequency snapshot sky survey (MSSS)-I. survey description and first results. *Astronomy & Astrophysics*, 582:A123.

- Hecht, J. H., Kovalam, S., May, P. T., Mills, G., Vincent, R. A., Walterscheid, R. L., and Woithe, J. (2004). Airglow imager observations of atmospheric gravity waves at Alice Springs and Adelaide, Australia during the Darwin Area Wave Experiment (DAWEX). *Journal of Geophysical Research: Atmospheres*, 109(D20).
- Helmboldt, J. and Hurley-Walker, N. (2020). Ionospheric irregularities observed during the GLEAM survey. *Radio Science*, 55(10):1–15.
- Helmboldt, J., Lazio, T., and Dymond, K. (2012a). A new technique for spectral analysis of ionospheric TEC fluctuations observed with the Very Large Array VHF system: From QP echoes to MSTIDs. *Radio Science*, 47(04):1–21.
- Helmboldt, J., Lazio, T., Intema, H., and Dymond, K. (2012b). High-precision measurements of ionospheric TEC gradients with the Very Large Array VHF system. *Radio Science*, 47(06):1–13.
- Helmboldt, J., Markowski, B., Bonanno, D., Clarke, T., Dowell, J., Hicks, B., Kassim, N., and Taylor, G. (2021). The deployable low-band ionosphere and transient experiment. *Radio Science*, 56(7):1–29.
- Helmboldt, J. and Zabolotin, N. (2022). An observed trend between mid-latitudes km-scale irregularities and medium-scale traveling ionospheric disturbances. *Radio Science*, 57(5):e2021RS007396.
- Helmboldt, J. F. (2012). Insights into the nature of northwest-to-southeast aligned ionospheric wavefronts from contemporaneous Very Large Array and ionosonde observations. *Journal of Geophysical Research: Space Physics*, 117(A7).
- Helmboldt, J. F. (2023). A method for estimating the spatial coherence of mid-latitude skywave propagation based on transionospheric scintillations at 35 MHz. *Radio Science*, 58(1):e2022RS007630.
- Helmboldt, J. F. and Intema, H. T. (2012). Very large array observations of disturbed ion flow from the plasmasphere to the nighttime ionosphere. *Radio Science*, 47(6).
- Helmboldt, J. F., Lane, W. M., and Cotton, W. D. (2012c). Climatology of midlatitude ionospheric disturbances from the Very Large Array Low-frequency Sky Survey. *Radio Science*, 47(5).
- Herman, J. R. (1966). Spread F and ionospheric F-region irregularities. *Reviews of Geophysics*, 4(2):255–299.
- Hey, J. S., Parsons, S. J., and Phillips, J. W. (1946). Fluctuations in cosmic radiation at radio-frequencies. *Nature*, page 234.
- Hines, C. O. (1960). Internal atmospheric gravity waves at ionospheric heights. *Canadian Journal of Physics*, 38(11):1441–1481.
- Hocke, K. and Igarashi, K. (2003). Wave-optical simulation of the oblique HF radio field. *Radio Science*, 38(3).
- Hocke, K. and Schlegel, K. (1996). A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995. *Annales Geophysicae*, 14(9):917.
- Hodos, T. J., Nava, O. A., Dao, E. V., and Emmons, D. J. (2022). Global sporadic-E occurrence rate climatology using GPS radio occultation and ionosonde data. *Journal of Geophysical Research: Space Physics*, 127(12):e2022JA030795.
- Hoogeveen, G. and Jacobson, A. (1997). Improved analysis of plasmasphere motion using the VLA radio interferometer. *Annales Geophysicae*, 15(2):236–245.

- Hooke, W. H. (1968). Ionospheric irregularities produced by internal atmospheric gravity waves. *Journal of Atmospheric and Terrestrial Physics*, 30(5):795–823.
- Hooke, W. H. (1970a). Ionospheric response to an isotropic spectrum of internal gravity waves. *Planetary and Space Science*, 18(12):1793–1797.
- Hooke, W. H. (1970b). The ionospheric response to internal gravity waves: 1. the F2 region response. *Journal of Geophysical Research*, 75(28):5535–5544.
- Hooke, W. H. (1970c). Ionospheric response to internal gravity waves: 2. lower F region response. *Journal of Geophysical Research*, 75(34):7229–7238.
- Hothi, I., Chapman, E., Pritchard, J. R., Mertens, F. G., Koopmans, L. V. E., Ciardi, B., Gehlot, B. K., Ghara, R., Ghosh, A., Giri, S. K., Iliev, I. T., Jelić, V., and Zaroubi, S. (2021). Comparing foreground removal techniques for recovery of the LOFAR-EoR 21 cm power spectrum. *Monthly Notices of the Royal Astronomical Society*, 500(2):2264–2277.
- Huang, C.-S., Miller, C. A., and Kelley, M. C. (1994). Basic properties and gravity wave initiation of the midlatitude F region instability. *Radio Science*, 29(1):395–405.
- Huba, J. D., Ossakow, S. L., Satyanarayana, P., and Guzdar, P. N. (1983). Linear theory of the $E \times B$ instability with an inhomogeneous electric field. *Journal of Geophysical Research: Space Physics*, 88(A1):425–434.
- Hudson, M. K. and Kelley, M. C. (1976). The temperature gradient drift instability at the equatorward edge of the ionospheric plasma trough. *Journal of Geophysical Research*, 81(22):3913–3918.
- Hunsucker, R. D. (1982). Atmospheric gravity waves generated in the high-latitude ionosphere: A review. *Reviews of Geophysics*, 20(2):293–315.
- Immel, T., Sagawa, E., England, S., Henderson, S., Hagan, M., Mende, S., Frey, H., Swenson, C., and Paxton, L. (2006). Control of equatorial ionospheric morphology by atmospheric tides. *Geophysical Research Letters*, 33(15).
- Inchin, P. A., Snively, J. B., Zettergren, M. D., Komjathy, A., Verkhoglyadova, O. P., and Tulasi Ram, S. (2020). Modeling of ionospheric responses to atmospheric acoustic and gravity waves driven by the 2015 Nepal M_w 7.8 Gorkha earthquake. *Journal of Geophysical Research: Space Physics*, 125(4):e2019JA027200.
- Ingham, M. F. (1971). The light of the night sky and the interplanetary medium. *Reports on Progress in Physics*, 34(3):875.
- Intema, H. T., van der Tol, S., Cotton, W. D., Cohen, A. S., van Bemmelen, I. M., and Röttgering, H. J. A. (2009). Ionospheric calibration of low frequency radio interferometric observations using the peeling scheme - I. method description and first results. *Astronomy & Astrophysics*, 501(3):1185–1205.
- Ivanova, V., Kurkin, V., Polekh, N., Chistyakova, L., Brynko, I., Chuyev, V., Dumbrava, Z., and Poddelskii, I. (2011). Studying large-scale traveling ionospheric disturbances according to the data of oblique-incidence sounding. *Geomagnetism and Aeronomy*, 51(8):1101–1104.
- Ivarsen, M. F., Huyghebaert, D. R., Gillies, M. D., St-Maurice, J.-P., Themens, D. R., Oppenheim, M., Gustavsson, B. J., Billett, D., Pitzel, B., Galeschuk, D., Donovan, E., and Hussey, G. C. (2024). Turbulence around auroral arcs. *Journal of Geophysical Research: Space Physics*, 129(8):e2023JA032309.

- Jacobson, A. R., Carlos, R. C., Massey, R. S., and Wu, G. (1995). Observations of traveling ionospheric disturbances with a satellite-beacon radio interferometer: Seasonal and local time behavior. *Journal of Geophysical Research: Space Physics*, 100(A2):1653–1665.
- Jacobson, A. R. and Erickson, W. C. (1992). Wavenumber-resolved observations of ionospheric waves using the Very Large Array radiotelescope. *Planetary and Space Science*, 40(4):447–455.
- Jacobson, A. R. and Erickson, W. C. (1993). Observations of electron density irregularities in the plasmasphere using the VLA radio interferometer. *Annales Geophysicae*, 11(10):869–888.
- Jansky, K. (1933). Electrical disturbances apparently of extraterrestrial origin. *Proceedings of the Institute of Radio Engineers*, 21(10):1387–1398.
- Jayachandran, P. T., Langley, R. B., MacDougall, J. W., Mushini, S. C., Pokhotelov, D., Hamza, A. M., Mann, I. R., Milling, D. K., Kale, Z. C., Chadwick, R., Kelly, T., Danskin, D. W., and Carrano, C. S. (2009). Canadian High Arctic Ionospheric Network (CHAIN). *Radio Science*, 44(1).
- Jenner, L. A., Wood, A. G., Dorrian, G. D., Oksavik, K., Yeoman, T. K., Fogg, A. R., and Coster, A. J. (2020). Plasma density gradients at the edge of polar ionospheric holes: the absence of phase scintillation. *Annales Geophysicae*, 38(2):575–590.
- Jin, S., Wang, Q., and Dardanelli, G. (2022a). A review on multi-GNSS for Earth observation and emerging applications. *Remote Sensing*, 14(16):3930.
- Jin, Y., Clausen, L. B. N., Spicher, A., Ivarsen, M. F., Zhang, Y., Miloch, W. J., and Moen, J. I. (2021). Statistical distribution of decameter scale (50 m) ionospheric irregularities at high latitudes. *Geophysical Research Letters*, 48(19):e2021GL094794.
- Jin, Y., Kotova, D., Xiong, C., Brask, S. M., Clausen, L. B. N., Kervalishvili, G., Stolle, C., and Miloch, W. J. (2022b). Ionospheric Plasma IRregularities - IPIR - data product based on data from the Swarm satellites. *Journal of Geophysical Research: Space Physics*, 127(4):e2021JA030183.
- Jin, Y., Moen, J. I., Miloch, W. J., Clausen, L. B. N., and Oksavik, K. (2016). Statistical study of the GNSS phase scintillation associated with two types of auroral blobs. *Journal of Geophysical Research: Space Physics*, 121(5):4679–4697.
- Jodalén, V., Bergsvik, T., Cannon, P. S., and Arthur, P. C. (2001). Performance of HF modems on high-latitude paths using multiple frequencies. *Radio Science*, 36(6):1687–1698.
- Johnsen, M. G. (2013). Real-time determination and monitoring of the auroral electrojet boundaries. *Journal of Space Weather and Space Climate*, 3:A28.
- Jordan, C., Murray, S., Trott, C., Wayth, R., Mitchell, D., Rahimi, M., Pindor, B., Procopio, P., and Morgan, J. (2017). Characterization of the ionosphere above the Murchison Radio Observatory using the Murchison Widefield Array. *Monthly Notices of the Royal Astronomical Society*, 471(4):3974–3987.
- Kadomtsev, B. B. (1965). *Plasma Turbulence*. Academic Press.
- Kane, T. A., Makarevich, R. A., and Devlin, J. C. (2012). HF radar observations of ionospheric backscatter during geomagnetically quiet periods. *Annales Geophysicae*, 30(1):221–233.
- Kaplan, E. D. and Hegarty, C. (2017). *Understanding GPS/GNSS: principles and applications*. Artech House.

- Katamzi, Z., Smith, N., Mitchell, C., Spalla, P., and Materassi, M. (2012). Statistical analysis of travelling ionospheric disturbances using TEC observations from geostationary satellites. *Journal of Atmospheric and Solar-Terrestrial Physics*, 74:64–80.
- Kelley, M. C. (2009). *The Earth's ionosphere: Plasma physics and electrodynamics*. Academic Press.
- Kelley, M. C. and Fukao, S. (1991). Turbulent upwelling of the mid-latitude ionosphere: 2. theoretical framework. *Journal of Geophysical Research: Space Physics*, 96(A3):3747–3753.
- Kelley, M. C., Haldoupis, C., Nicolls, M. J., Makela, J. J., Belehaki, A., Shalimov, S., and Wong, V. K. (2003). Case studies of coupling between the E and F regions during unstable sporadic-E conditions. *Journal of Geophysical Research: Space Physics*, 108(A12).
- Keskinen, M. and Ossakow, S. (1983). Theories of high-latitude ionospheric irregularities: A review. *Radio Science*, 18(06):1077–1091.
- Keskinen, M. J. (1984). Nonlinear theory of the $E \times B$ instability with an inhomogeneous electric field. *Journal of Geophysical Research: Space Physics*, 89(A6):3913–3920.
- Keskinen, M. J., Basu, S., and Basu, S. (2004). Midlatitude sub-auroral ionospheric small scale structure during a magnetic storm. *Geophysical Research Letters*, 31(9).
- Keskinen, M. J., Mitchell, H. G., Fedder, J. A., Satyanarayana, P., Zalesak, S. T., and Huba, J. D. (1988). Nonlinear evolution of the Kelvin-Helmholtz instability in the high-latitude ionosphere. *Journal of Geophysical Research: Space Physics*, 93(A1):137–152.
- Kirchengast, G. (1997). Characteristics of high-latitude TIDs from different causative mechanisms deduced by theoretical modeling. *Journal of Geophysical Research: Space Physics*, 102(A3):4597–4612.
- Kirkland, M. W. and Jacobson, A. R. (1998). Drift-parallax determination of the altitude of traveling ionospheric disturbances observed with the Los Alamos radio-beacon interferometer. *Radio Science*, 33(6):1807–1825.
- Knudsen, D. J., Burchill, J. K., Buchert, S. C., Eriksson, A. I., Gill, R., Wahlund, J.-E., Åhlen, L., Smith, M., and Moffat, B. (2017). Thermal ion imagers and Langmuir probes in the Swarm electric field instruments. *Journal of Geophysical Research: Space Physics*, 122(2):2655–2673.
- Koval, A., Chen, Y., Stanislavsky, A., Kashcheyev, A., and Zhang, Q.-H. (2018). Simulation of focusing effect of traveling ionospheric disturbances on meter-decameter solar dynamic spectra. *Journal of Geophysical Research: Space Physics*, 123(11):8940–8950.
- Koval, A., Chen, Y., Stanislavsky, A., and Zhang, Q.-H. (2017). Traveling ionospheric disturbances as huge natural lenses: Solar radio emission focusing effect. *Journal of Geophysical Research: Space Physics*, 122(9):9092–9101.
- Koval, A., Chen, Y., Tsugawa, T., Otsuka, Y., Shinbori, A., Nishioka, M., Brazhenko, A., Stanislavsky, A., Konovalenko, A., Zhang, Q.-H., Monstein, C., and Gorgutsa, R. (2019). Direct observations of traveling ionospheric disturbances as focusers of solar radiation: Spectral caustics. *The Astrophysical Journal*, 877(2):98.
- Krankowski, A., Shagimuratov, I., Ephishov, I., Krypiak-Gregorczyk, A., and Yakimova, G. (2009). The occurrence of the mid-latitude ionospheric trough in GPS-TEC measurements. *Advances in Space Research*, 43(11):1721–1731.

- Kuiack, M., Wijers, R. A., Shulevski, A., Rowlinson, A., Huizinga, F., Molenaar, G., and Prasad, P. (2021a). The AARTFAAC 60 MHz transients survey. *Monthly Notices of the Royal Astronomical Society*, 505(2):2966–2974.
- Kuiack, M. J., Wijers, R. A., Shulevski, A., and Rowlinson, A. (2021b). Apparent radio transients mapping the near-earth plasma environment. *Monthly Notices of the Royal Astronomical Society*, 504(4):4706–4715.
- Lan, J., Ning, B., Li, G., Zhu, Z., Hu, L., and Sun, W. (2018). Observation of short-period ionospheric disturbances using a portable digital ionosonde at Sanya. *Radio Science*, 53(12):1521–1532.
- Lane, T. P., Reeder, M. J., and Clark, T. L. (2001). Numerical modeling of gravity wave generation by deep tropical convection. *Journal of the Atmospheric Sciences*, 58(10):1249–1274.
- Larsen, M. F. (2000). A shear instability seeding mechanism for quasiperiodic radar echoes. *Journal of Geophysical Research: Space Physics*, 105(A11):24931–24940.
- Larsen, M. F., Yamamoto, M., Fukao, S., Tsunoda, R. T., and Saito, A. (2005). Observations of neutral winds, wind shears, and wave structure during a sporadic-E/QP event. *Annales Geophysicae*, 23(7):2369–2375.
- Laštovička, J. (2001). Nitric oxide densities and their diurnal asymmetry in the upper middle atmosphere as revealed by ionospheric measurements. *Journal of Atmospheric and Solar-Terrestrial Physics*, 63(1):21–28.
- Laughman, B., Fritts, D. C., and Lund, T. S. (2017). Tsunami-driven gravity waves in the presence of vertically varying background and tidal wind structures. *Journal of Geophysical Research: Atmospheres*, 122(10):5076–5096.
- Lawrence, R., Little, C., and Chivers, H. (1964). A survey of ionospheric effects upon Earth-space radio propagation. *Proceedings of the IEEE*, 52(1):4–27.
- Lay, E. H., Shao, X.-M., and Carrano, C. S. (2013). Variation in total electron content above large thunderstorms. *Geophysical Research Letters*, 40(10):1945–1949.
- Leick, A., Rapoport, L., and Tatarnikov, D. (2015). *GPS satellite surveying*. Wiley, fourth edition.
- Lin, F. F., Wang, C. Y., Su, C. L., Shiokawa, K., Saito, S., and Chu, Y. H. (2016). Coordinated observations of F region 3 m field-aligned plasma irregularities associated with medium-scale traveling ionospheric disturbances. *Journal of Geophysical Research: Space Physics*, 121(4):3750–3766.
- Linson, L. M. and Workman, J. B. (1970). Formation of striations in ionospheric plasma clouds. *Journal of Geophysical Research*, 75(16):3211–3219.
- Little, C. G. and Lovell, A. C. B. (1950). Origin of the fluctuations in the intensity of radio waves from galactic sources: Jodrell Bank observations. *Nature*, pages 423–424.
- Liu, X., Xu, J., Yue, J., Vadas, S. L., and Becker, E. (2019). Orographic primary and secondary gravity waves in the middle atmosphere from 16-year SABER observations. *Geophysical Research Letters*, 46(8):4512–4522.
- Liu, Y., San Liang, X., and Weisberg, R. H. (2007). Rectification of the bias in the wavelet power spectrum. *Journal of Atmospheric and Oceanic Technology*, 24(12):2093–2102.

- Liu, Y., Xiong, C., Wan, X., Lai, Y., Wang, Y., Yu, X., and Ou, M. (2021). Instability mechanisms for the F-region plasma irregularities inside the midlatitude ionospheric trough: Swarm observations. *Space Weather*, 19(7):e2021SW002785.
- Loi, S. T., Cairns, I. H., Murphy, T., Erickson, P. J., Bell, M. E., Rowlinson, A., Arora, B. S., Morgan, J., Ekers, R. D., Hurley-Walker, N., and Kaplan, D. L. (2016a). Density duct formation in the wake of a travelling ionospheric disturbance: Murchison Widefield Array observations. *Journal of Geophysical Research: Space Physics*, 121(2):1569–1586.
- Loi, S. T., Murphy, T., Bell, M. E., Kaplan, D. L., Lenc, E., Offringa, A. R., Hurley-Walker, N., Bernardi, G., Bowman, J. D., Briggs, F., Cappallo, R. J., Corey, B. E., Deshpande, A. A., Emrich, D., Gaensler, B. M., Goeke, R., Greenhill, L. J., Hazelton, B. J., Johnston-Hollitt, M., Kasper, J. C., Kratzenberg, E., Lonsdale, C. J., Lynch, M. J., McWhirter, S. R., Mitchell, D. A., Morales, M. F., Morgan, E., Oberoi, D., Ord, S. M., Prabu, T., Rogers, A. E. E., Roshi, A., Shankar, N. U., Srivani, K. S., Subrahmanyam, R., Tingay, S. J., Waterson, M., Wayth, R. B., Webster, R. L., Whitney, A. R., Williams, A., and Williams, C. L. (2015a). Quantifying ionospheric effects on time-domain astrophysics with the Murchison Widefield Array. *Monthly Notices of the Royal Astronomical Society*, 453(3):2731–2746.
- Loi, S. T., Murphy, T., Cairns, I. H., Menk, F. W., Waters, C. L., Erickson, P. J., Trott, C. M., Hurley-Walker, N., Morgan, J., Lenc, E., Offringa, A. R., Bell, M. E., Ekers, R. D., Gaensler, B. M., Lonsdale, C. J., Feng, L., Hancock, P. J., Kaplan, D. L., Bernardi, G., Bowman, J. D., Briggs, F., Cappallo, R. J., Deshpande, A. A., Greenhill, L. J., Hazelton, B. J., Johnston-Hollitt, M., McWhirter, S. R., Mitchell, D. A., Morales, M. F., Morgan, E., Oberoi, D., Ord, S. M., Prabu, T., Shankar, N. U., Srivani, K. S., Subrahmanyam, R., Tingay, S. J., Wayth, R. B., Webster, R. L., Williams, A., and Williams, C. L. (2015b). Real-time imaging of density ducts between the plasmasphere and ionosphere. *Geophysical Research Letters*, 42(10):3707–3714.
- Loi, S. T., Murphy, T., Cairns, I. H., Trott, C. M., Hurley-Walker, N., Feng, L., Hancock, P. J., and Kaplan, D. L. (2016b). A new angle for probing field-aligned irregularities with the Murchison Widefield Array. *Radio Science*, 51(6):659–679.
- Lovell, A. C. B. (1957). The Jodrell Bank radio telescope. *Nature*, pages 60–62.
- Ludwig-Barbosa, V., Rasch, J., Carlström, A., Pettersson, M. I., and Vu, V. T. (2019). GNSS radio occultation simulation using multiple phase screen orbit sampling. *IEEE Geoscience and Remote Sensing Letters*, 17(8):1323–1327.
- Mangogna, A., Swenson, G., Vargas, F., and Liu, A. (2016). A mesospheric airglow multichannel photometer and an optical method to measure mesospheric AGW intrinsic parameters. *Journal of Atmospheric and Solar-Terrestrial Physics*, 142:108–119.
- Maruyama, T., Saito, S., Yamamoto, M., and Fukao, S. (2006). Simultaneous observation of sporadic E with a rapid-run ionosonde and VHF coherent backscatter radar. *Annales Geophysicae*, 24(1):153–162.
- Matzka, J., Stolle, C., Yamazaki, Y., Bronkalla, O., and Morschhauser, A. (2021). The geomagnetic Kp index and derived indices of geomagnetic activity. *Space Weather*, 19(5):e2020SW002641.
- McCrea, I., Aikio, A., Alfonsi, L., Belova, E., Buchert, S., Clilverd, M., Engler, N., Gustavsson, B., Heinselman, C., Kero, J., Kosch, M., Lamy, H., Leyser, T., Ogawa, Y., Oksavik, K., Pellinen-Wanberg, A., Pitout, F., Rapp, M., Stanislawski, I., and Vierinen, J. (2015). The science case for the EISCAT_3D radar. *Progress in Earth and Planetary Science*, 2:1–63.

- McKay-Bukowski, D., Vierinen, J., Virtanen, I. I., Fallows, R., Postila, M., Ulich, T., Wucknitz, O., Brentjens, M., Ebbendorf, N., Enell, C.-F., Gerbers, M., Grit, T., Gruppen, P., Kero, A., Iinatti, T., Lehtinen, M., Meulman, H., Norden, M., Orispää, M., Raita, T., de Reijer, J. P., Roininen, L., Schoenmakers, A., Stuurwold, K., and Turunen, E. (2014). KAIRA: The Kilpisjärvi Atmospheric Imaging Receiver Array—system overview and first results. *IEEE Transactions on Geoscience and Remote Sensing*, 53(3):1440–1451.
- McKean, J. P., Godfrey, L. E. H., Vegetti, S., Wise, M. W., Morganti, R., Hardcastle, M. J., Rafferty, D., Anderson, J., Avruich, I. M., Beck, R., Bell, M. E., van Bemmell, I., Bentum, M. J., Bernardi, G., Best, P., Blaauw, R., Bonafede, A., Breitling, F., Broderick, J. W., Brüggem, M., Cerrigone, L., Ciardi, B., de Gasperin, F., Deller, A., Duscha, S., Engels, D., Falcke, H., Fallows, R. A., Frieswijk, W., Garrett, M. A., Grießmeier, J. M., van Haarlem, M. P., Heald, G., Hoeft, M., van der Horst, A. J., Iacobelli, M., Intema, H., Juette, E., Karastergiou, A., Kondratiev, V. I., Koopmans, L. V. E., Kuniyoshi, M., Kuper, G., van Leeuwen, J., Maat, P., Mann, G., Markoff, S., McFadden, R., McKay-Bukowski, D., Mulcahy, D. D., Munk, H., Nelles, A., Orru, E., Paas, H., Pandey-Pommier, M., Pietka, M., Pizzo, R., Polatidis, A. G., Reich, W., Röttgering, H. J. A., Rowlinson, A., Scaife, A. M. M., Serylak, M., Shulevski, A., Sluman, J., Smirnov, O., Steinmetz, M., Stewart, A., Swinbank, J., Tagger, M., Thoudam, S., Toribio, M. C., Vermeulen, R., Vocks, C., van Weeren, R. J., Wucknitz, O., Yatawatta, S., and Zarka, P. (2016). LOFAR imaging of Cygnus A – direct detection of a turnover in the hotspot radio spectra. *Monthly Notices of the Royal Astronomical Society*, 463(3):3143–3150.
- McNamara, L. F., Angling, M. J., Elvidge, S., Fridman, S. V., Hausman, M. A., Nickisch, L. J., and McKinnell, L.-A. (2013). Assimilation procedures for updating ionospheric profiles below the F2 peak. *Radio Science*, 48(2):143–157.
- Medvedev, A., Ratovsky, K., Tolstikov, M., Alsatkin, S., and Shcherbakov, A. (2015). A statistical study of internal gravity wave characteristics using the combined Irkutsk Incoherent Scatter Radar and Digisonde data. *Journal of Atmospheric and Solar-Terrestrial Physics*, 132:13–21.
- Medvedev, A. V., Ratovsky, K. G., Tolstikov, M. V., Oinats, A. V., Alsatkin, S. S., and Zhrebtssov, G. A. (2017). Relation of internal gravity wave anisotropy with neutral wind characteristics in the upper atmosphere. *Journal of Geophysical Research: Space Physics*, 122(7):7567–7580.
- Meinel, I. A. B. (1950). OH emission bands in the spectrum of the night sky. *The Astrophysical Journal*, 111:555–564.
- Mercier, R. P. (1962). Diffraction by a screen causing large random phase fluctuations. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58(2):382–400.
- Mevius, M., van der Tol, S., Pandey, V. N., Vedantham, H. K., Brentjens, M. A., de Bruyn, A. G., Abdalla, F. B., Asad, K. M. B., Bregman, J. D., Brouw, W. N., Bus, S., Chapman, E., Ciardi, B., Fernandez, E. R., Ghosh, A., Harker, G., Iliev, I. T., Jelić, V., Kazemi, S., Koopmans, L. V. E., Noordam, J. E., Offringa, A. R., Patil, A. H., van Weeren, R. J., Wijnholds, S., Yatawatta, S., and Zaroubi, S. (2016). Probing ionospheric structures using the LOFAR radio telescope. *Radio Science*, 51(7):927–941.
- Meyer-Vernet, N. (1980). On a day-time ionospheric effect on some radio intensity measurements and interferometry. *Astronomy & Astrophysics*, 84:142–147.
- Meyer-Vernet, N., Daigne, G., and Lecacheux, A. (1981). Dynamic spectra of some terrestrial ionospheric effects at decametric wavelengths-applications in other astrophysical contexts. *Astronomy & Astrophysics*, 96:296–301.

- Miller, C. A., Swartz, W. E., Kelley, M. C., Mendillo, M., Nottingham, D., Scali, J., and Reinisch, B. (1997). Electrodynamics of midlatitude spread F: 1. observations of unstable, gravity wave-induced ionospheric electric fields at tropical latitudes. *Journal of Geophysical Research: Space Physics*, 102(A6):11521–11532.
- Millman, G. H. (1967). A survey of tropospheric, ionospheric, and extra-terrestrial effects on radio propagation between Earth and space vehicles. In *Propagation Factors in Space Communications*, pages 3–58.
- Mills, B. and Thomas, A. (1951). Observations of the source of radio-frequency radiation in the constellation of Cygnus. *Australian Journal of Chemistry*, 4(2):158–171.
- Miura, A. and Pritchett, P. L. (1982). Nonlocal stability analysis of the MHD Kelvin-Helmholtz instability in a compressible plasma. *Journal of Geophysical Research: Space Physics*, 87(A9):7431–7444.
- Moffett, R. and Quegan, S. (1983). The mid-latitude trough in the electron concentration of the ionospheric F-layer: a review of observations and modelling. *Journal of Atmospheric and Terrestrial Physics*, 45(5):315–343.
- Mozer, F. (1970). Electric field mapping in the ionosphere at the equatorial plane. *Planetary and Space Science*, 18(2):259–263.
- Narayan, R. (1992). The physics of pulsar scintillation. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 341(1660):151–165.
- Negrea, C., Zabolotin, N., and Bullett, T. (2018). Seasonal variability of the midlatitude traveling ionospheric disturbances from Wallops Island, VA, Dynasonde data: Evidence of a semiannual variation. *Journal of Geophysical Research: Space Physics*, 123(6):5047–5054.
- Negrea, C., Zabolotin, N., Bullett, T., Fuller-Rowell, T., Fang, T.-W., and Codrescu, M. (2016). Characteristics of acoustic gravity waves obtained from Dynasonde data. *Journal of Geophysical Research: Space Physics*, 121(4):3665–3680.
- Nekrasov, A. and Shalimov, S. (2002). Nonlinear structures of internal gravitational waves and their effect on the ionosphere. *Cosmic Research*, 40(5):517–520.
- Nekrasov, A., Shalimov, S., Shukla, P., and Stenflo, L. (1995). Nonlinear disturbances in the ionosphere due to acoustic gravity waves. *Journal of Atmospheric and Terrestrial Physics*, 57(7):737–741.
- Newell, P. and Gjerloev, J. (2011a). Evaluation of SuperMAG auroral electrojet indices as indicators of substorms and auroral power. *Journal of Geophysical Research: Space Physics*, 116(A12).
- Newell, P. and Gjerloev, J. (2011b). Substorm and magnetosphere characteristic scales inferred from the SuperMAG auroral electrojet indices. *Journal of Geophysical Research: Space Physics*, 116(A12).
- Nicolls, M. J. and Heinselman, C. J. (2007). Three-dimensional measurements of traveling ionospheric disturbances with the Poker Flat Incoherent Scatter Radar. *Geophysical Research Letters*, 34(21).
- Nishimura, Y., Verkhoglyadova, O., Deng, Y., and Zhang, S.-R. (2021). *Cross-scale Coupling and Energy Transfer in the Magnetosphere-Ionosphere-Thermosphere System*. Elsevier.

- Nishitani, N., Ruohoniemi, J. M., Lester, M., Baker, J. B. H., Koustov, A. V., Shepherd, S. G., Chisham, G., Hori, T., Thomas, E. G., Makarevich, R. A., Marchaudon, A., Ponomarenko, P., Wild, J. A., Milan, S. E., Bristow, W. A., Devlin, J., Miller, E., Greenwald, R. A., Ogawa, T., and Kikuchi, T. (2019). Review of the accomplishments of mid-latitude Super Dual Auroral Radar Network (SuperDARN) HF radars. *Progress in Earth and Planetary Science*, 6:1–57.
- Norris, R. P. (2017). Extragalactic radio continuum surveys and the transformation of radio astronomy. *Nature Astronomy*, 1:671–678.
- Norton, R. B. and Barth, C. A. (1970). Theory of nitric oxide in the Earth’s atmosphere. *Journal of Geophysical Research*, 75(19):3903–3909.
- Oinats, A. V., Nishitani, N., Ponomarenko, P., Bergardt, O. I., and Ratovsky, K. G. (2016). Statistical characteristics of medium-scale traveling ionospheric disturbances revealed from the Hokkaido East and Ekaterinburg HF radar data. *Earth, Planets and Space*, 68(1):1–13.
- Oksman, J. and Tauriainen, A. (1971). On annual movements of the scintillation boundary of satellite signals. *Journal of Atmospheric and Terrestrial Physics*, 33(11):1727–1735.
- Oliver, W. L., Otsuka, Y., Sato, M., Takami, T., and Fukao, S. (1997). A climatology of F region gravity wave propagation over the middle and upper atmosphere radar. *Journal of Geophysical Research: Space Physics*, 102(A7):14499–14512.
- Ossakow, S. and Chaturvedi, P. (1979). Current convective instability in the diffuse aurora. *Geophysical Research Letters*, 6(4):332–334.
- Otsuka, Y., Shinbori, A., Tsugawa, T., and Nishioka, M. (2021). Solar activity dependence of medium-scale traveling ionospheric disturbances using GPS receivers in Japan. *Earth, Planets and Space*, 73:1–11.
- Otsuka, Y., Suzuki, K., Nakagawa, S., Nishioka, M., Shiokawa, K., and Tsugawa, T. (2013). GPS observations of medium-scale traveling ionospheric disturbances over Europe. *Annales Geophysicae*, 31(2):163–172.
- Panasenko, S. V., Goncharenko, L. P., Erickson, P. J., Aksonova, K. D., and Domnin, I. F. (2018). Traveling ionospheric disturbances observed by Kharkiv and Millstone Hill incoherent scatter radars near vernal equinox and summer solstice. *Journal of Atmospheric and Solar-Terrestrial Physics*, 172:10–23.
- Park, J., Lühr, H., Lee, C., Kim, Y. H., Jee, G., and Kim, J.-H. (2014). A climatology of medium-scale gravity wave activity in the midlatitude/low-latitude daytime upper thermosphere as observed by CHAMP. *Journal of Geophysical Research: Space Physics*, 119(3):2187–2196.
- Pederick, L. and Cervera, M. (2016). A directional HF noise model: Calibration and validation in the Australian region. *Radio Science*, 51(1):25–39.
- Percival, D. B. and Walden, A. T. (1993). *Parametric Spectral Estimation*, chapter 9, page 391–455. Cambridge University Press.
- Perkins, F. (1973). Spread F and ionospheric currents. *Journal of Geophysical Research*, 78(1):218–226.
- Perlango, N. J., Ridley, A. J., Cnossen, I., and Wu, C. (2018). A year-long comparison of GPS TEC and global ionosphere-thermosphere models. *Journal of Geophysical Research: Space Physics*, 123(2):1410–1428.

- Perwitasari, S., Nakamura, T., Tsugawa, T., Nishioka, M., Tomikawa, Y., Ejiri, M. K., Kogure, M., Otsuka, Y., Shinbori, A., Jin, H., and Tao, C. (2022). Propagation direction analyses of medium-scale traveling ionospheric disturbances observed over North America with GPS-TEC perturbation maps by three-dimensional spectral analysis method. *Journal of Geophysical Research: Space Physics*, 127(1):e2020JA028791.
- Pi, X., Iijima, B. A., and Lu, W. (2017). Effects of ionospheric scintillation on GNSS-based positioning. *Navigation: Journal of The Institute of Navigation*, 64(1):3–22.
- Piggott, W. R. and Rawer, K. (1972). *URSI handbook of ionogram interpretation and reduction*. URSL.
- Pignalberi, A., Pezzopane, M., and Zuccheretti, E. (2014). Sporadic E layer at mid-latitudes: average properties and influence of atmospheric tides. *Annales Geophysicae*, 32(11):1427–1440.
- Porayko, N. K., Mevius, M., Hernández-Pajares, M., Tiburzi, C., Pulido, G. O., Liu, Q., Verbiest, J. P. W., Künsemöller, J., Krishnakumar, M. A., Nielsen, A.-S. B., Brüggén, M., Graffigna, V., Dettmar, R.-J., Kramer, M., Osłowski, S., Schwarz, D. J., Shaifullah, G. M., and Wucknitz, O. (2023). Validation of global ionospheric models using long-term observations of pulsar Faraday rotation with the LOFAR radio telescope. *Journal of Geodesy*, 97:116.
- Porayko, N. K., Noutsos, A., Tiburzi, C., Verbiest, J. P. W., Horneffer, A., Künsemöller, J., Osłowski, S., Kramer, M., Schnitzeler, D. H. F. M., Anderson, J. M., Brüggén, M., Griesmeier, J.-M., Hoeft, M., Schwarz, D. J., Serylak, M., and Wucknitz, O. (2019). Testing the accuracy of the ionospheric Faraday rotation corrections through LOFAR observations of bright northern pulsars. *Monthly Notices of the Royal Astronomical Society*, 483(3):4100–4113.
- Priyadarshi, S. (2015). A review of ionospheric scintillation models. *Surveys in Geophysics*, 36:295–324.
- Pryse, S., Kersley, L., and Walker, I. (1996). Blobs and irregularities in the auroral ionosphere. *Journal of Atmospheric and Terrestrial Physics*, 58(1-4):205–215.
- Qian, L., Burns, A. G., Emery, B. A., Foster, B., Lu, G., Maute, A., Richmond, A. D., Roble, R. G., Solomon, S. C., and Wang, W. (2014). The NCAR TIE-GCM: A community model of the coupled thermosphere/ionosphere system. In *Modeling the Ionosphere–Thermosphere System*, chapter 7, pages 73–83. American Geophysical Union.
- Ratcliffe, J. A. (1959). *The magneto-ionic theory and its applications to the ionosphere*. Cambridge University Press.
- Rathod, C., Srinivasan, B., and Scales, W. (2021). Modeling the dominance of the gradient drift or Kelvin–Helmholtz instability in sheared ionospheric $E \times B$ flows. *Physics of Plasmas*, 28(5):052903.
- Reinisch, B., Huang, X., Galkin, I., Paznukhov, V., and Kozlov, A. (2005). Recent advances in real-time analysis of ionograms and ionospheric drift measurements with digisondes. *Journal of Atmospheric and Solar-Terrestrial Physics*, 67(12):1054–1062.
- Reinisch, B. W. and Galkin, I. A. (2011). Global Ionospheric Radio Observatory (GIRO). *Earth, Planets and Space*, 63(4):377–381.
- Rideout, W. and Coster, A. (2006). Automated GPS processing for global total electron content data. *GPS Solutions*, 10:219–228.
- Ridley, A., Deng, Y., and Tóth, G. (2006). The global ionosphere–thermosphere model. *Journal of Atmospheric and Solar-Terrestrial Physics*, 68(8):839–864.

- Rino, C., Gonzalez, V., and Hessing, A. (1981). Coherence bandwidth loss in transionospheric radio propagation. *Radio Science*, 16(02):245–255.
- Rino, C. L., Carrano, C. S., Groves, K. M., and Roddy, P. A. (2016). A characterization of intermediate-scale spread F structure from four years of high-resolution C/NOFS satellite data. *Radio Science*, 51(6):779–788.
- Rodger, A., Moffett, R., and Quegan, S. (1992). The role of ion drift in the formation of ionisation troughs in the mid- and high-latitude ionosphere—a review. *Journal of Atmospheric and Terrestrial Physics*, 54(1):1–30.
- Ruckstuhl, A. F., Jacobson, M. P., Field, R. W., and Dodd, J. A. (2001). Baseline subtraction using robust local regression estimation. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 68(2):179–193.
- Ryle, M. and Hewish, A. (1950). The effects of the terrestrial ionosphere on the radio waves from discrete sources in the galaxy. *Monthly Notices of the Royal Astronomical Society*, 110(4):381–394.
- Ryle, M. and Smith, F. G. (1948). A new intense source of radio-frequency radiation in the constellation of Cassiopeia. *Nature*, pages 462–463.
- Sahr, J. D. and Fejer, B. G. (1996). Auroral electrojet plasma irregularity theory and experiment: A critical review of present understanding and future directions. *Journal of Geophysical Research: Space Physics*, 101(A12):26893–26909.
- Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., and Tarantola, S. (2010). Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Computer Physics Communications*, 181(2):259–270.
- Sato, K. and Yoshiki, M. (2008). Gravity wave generation around the polar vortex in the stratosphere revealed by 3-hourly radiosonde observations at Syowa Station. *Journal of the Atmospheric Sciences*, 65(12):3719 – 3735.
- Schmidt, M. J. and Gary, S. P. (1973). Density gradients and the Farley-Buneman instability. *Journal of Geophysical Research*, 78(34):8261–8265.
- Shi, X. and Xu, Z. (2021). Plasma microlensing dynamic spectrum probing fine structures in the ionized interstellar medium. *Monthly Notices of the Royal Astronomical Society*, 506(4):6039–6051.
- Shimeis, A., Borries, C., Amory-Mazaudier, C., Fleury, R., Mahrous, A. M., Hassan, A., and Nawar, S. (2015). TEC variations along an east Euro-African chain during 5th April 2010 geomagnetic storm. *Advances in Space Research*, 55(9):2239–2247.
- Shimwell, T. W., Röttgering, H. J. A., Best, P. N., Williams, W. L., Dijkema, T. J., de Gasperin, F., Hardcastle, M. J., Heald, G. H., Hoang, D. N., Horneffer, A., Intema, H., Mahony, E. K., Mandal, S., Mechev, A. P., Morabito, L., Oonk, J. B. R., Rafferty, D., Retana-Montenegro, E., Sabater, J., Tasse, C., van Weeren, R. J., Brüggen, M., Brunetti, G., Chyzy, K. T., Conway, J. E., Haverkorn, M., Jackson, N., Jarvis, M. J., McKean, J. P., Miley, G. K., Morganti, R., White, G. J., Wise, M. W., van Bemmelen, I. M., Beck, R., Brienza, M., Bonafede, A., Calistro Rivera, G., Cassano, R., Clarke, A. O., Cseh, D., Deller, A., Drabent, A., van Driel, W., Engels, D., Falcke, H., Ferrari, C., Fröhlich, S., Garrett, M. A., Harwood, J. J., Heesen, V., Hoeft, M., Horellou, C., Israel, F. P., Kapińska, A. D., Kunert-Bajraszewska, M., McKay, D. J., Mohan, N. R., Orrú, E., Pizzo, R. F., Prandoni, I., Schwarz, D. J., Shulevski, A., Sipior, M., Smith, D. J. B., Sridhar, S. S., Steinmetz, M., Stroe, A., Varenus, E., van der Werf, P. P., Zensus, J. A., and Zwart, J. T. L. (2017). The

- LOFAR two-metre sky survey-I. survey description and preliminary data release. *Astronomy & Astrophysics*, 598:A104.
- Shiokawa, K., Otsuka, Y., Ihara, C., Ogawa, T., and Rich, F. J. (2003). Ground and satellite observations of nighttime medium-scale traveling ionospheric disturbance at midlatitude. *Journal of Geophysical Research: Space Physics*, 108(A4).
- Siddle, D., Zaalov, N., Stocker, A., and Warrington, E. M. (2004). Time of flight and direction of arrival of HF radio signals received over a path along the midlatitude trough: Theoretical considerations. *Radio Science*, 39(4):1–10.
- Singleton, D. (1970). Saturation and focusing effects in radio-star and satellite scintillations. *Journal of Atmospheric and Terrestrial Physics*, 32(2):187–208.
- Singleton, D. G. (1964). Broadband radio-star scintillations, part I. observations. *Radio Science Journal of Research*, pages 867–880.
- Sivakandan, M., Otsuka, Y., Ghosh, P., Shinagawa, H., Shinbori, A., and Miyoshi, Y. (2021). Comparison of seasonal and longitudinal variation of daytime MSTID activity using GPS observation and GAIA simulations. *Earth, Planets and Space*, 73:1–16.
- Skrutskie, M. F., Cutri, R. M., Stiening, R., Weinberg, M. D., Schneider, S., Carpenter, J. M., Beichman, C., Capps, R., Chester, T., Elias, J., Huchra, J., Liebert, J., Lonsdale, C., Monet, D. G., Price, S., Seitzer, P., Jarrett, T., Kirkpatrick, J. D., Gizis, J. E., Howard, E., Evans, T., Fowler, J., Fullmer, L., Hurt, R., Light, R., Kopan, E. L., Marsh, K. A., McCallon, H. L., Tam, R., Dyk, S. V., and Wheelock, S. (2006). The two micron all sky survey (2MASS). *The Astronomical Journal*, 131(2):1163–1183.
- Smirnov, O. M. (2011a). Revisiting the radio interferometer measurement equation-I. a full-sky Jones formalism. *Astronomy & Astrophysics*, 527:A106.
- Smirnov, O. M. (2011b). Revisiting the radio interferometer measurement equation-II. calibration and direction-dependent effects. *Astronomy & Astrophysics*, 527:A107.
- Smith, F. G. (1950). Origin of the fluctuations in the intensity of radio waves from galactic sources: Cambridge observations. *Nature*, pages 422–423.
- Snively, J. B. and Pasko, V. P. (2003). Breaking of thunderstorm-generated gravity waves as a source of short-period ducted waves at mesopause altitudes. *Geophysical Research Letters*, 30(24).
- Sobol, I. and Levitan, Y. L. (1999). A pseudo-random number generator for personal computers. *Computers & Mathematics with Applications*, 37(4-5):33–40.
- Sokolovskiy, S. V. (2001). Modeling and inverting radio occultation signals in the moist troposphere. *Radio Science*, 36(3):441–458.
- Spicher, A., Deshpande, K., Jin, Y., Oksavik, K., Zettergren, M. D., Clausen, L. B. N., Moen, J. I., Hairston, M. R., and Baddeley, L. (2020). On the production of ionospheric irregularities via Kelvin-Helmholtz instability associated with cusp flow channels. *Journal of Geophysical Research: Space Physics*, 125(6):e2019JA027734.
- Spiegel, E. A. and Veronis, G. (1960). On the Boussinesq approximation for a compressible fluid. *The Astrophysical Journal*, 131:442.

- Spiro, R., Heelis, R., and Hanson, W. (1978). Ion convection and the formation of the mid-latitude F region ionization trough. *Journal of Geophysical Research: Space Physics*, 83(A9):4255–4264.
- Spogli, L., Alfonsi, L., Romano, V., De Franceschi, G., Francisco, G. M. J., Shimabukuro, M. H., Bougard, B., and Aquino, M. (2013). Assessing the GNSS scintillation climate over Brazil under increasing solar activity. *Journal of Atmospheric and Solar-Terrestrial Physics*, 105:199–206.
- Spogli, L., Sabbagh, D., Regi, M., Cesaroni, C., Perrone, L., Alfonsi, L., di Mauro, D., Lepidi, S., Campuzano, S. A., Marchetti, D., de Santis, A., Malagnini, A., Scotto, C., Cianchini, G., Shen, X. H., Piscini, A., and Ippolito, A. (2021). Ionospheric response over Brazil to the August 2018 geomagnetic storm as probed by CSES-01 and Swarm satellites and by local ground-based observations. *Journal of Geophysical Research: Space Physics*, 126(2):e2020JA028368.
- Stewart, R. and McLean, D. (1982). Correcting low-frequency solar radio source positions for ionospheric refraction. *Publications of the Astronomical Society of Australia*, 4(4):386–389.
- Stinebring, D., McLaughlin, M., Cordes, J., Becker, K., Goodman, J. E., Kramer, M., Sheckard, J., and Smith, C. (2001). Faint scattering around pulsars: probing the interstellar medium on solar system size scales. *The Astrophysical Journal*, 549(1):L97.
- Swarup, G. (2015). Major advances in radio astronomy: Some key questions today. *Proceedings of the National Academy of Sciences, India Section A: Physical Sciences*, 85:465–481.
- Takahashi, T., Spicher, A., Di Mare, F., Rowland, D. E., Pfaff, R. F., Collier, M. R., Clausen, L. B. N., and Moen, J. I. (2022). Suppression of ionospheric irregularity due to auroral particle impact. *Journal of Geophysical Research: Space Physics*, 127(1):e2020JA028725.
- Taylor, M. J., Ryan, E., Tuan, T., and Edwards, R. (1993). Evidence of preferential directions for gravity wave propagation due to wind filtering in the middle atmosphere. *Journal of Geophysical Research: Space Physics*, 98(A4):6047–6057.
- Terra, P., Vargas, F., Brum, C. G., and Miller, E. S. (2020). Geomagnetic and solar dependency of MSTIDs occurrence rate: A climatology based on airglow observations from the Arecibo Observatory ROF. *Journal of Geophysical Research: Space Physics*, 125(7):e2019JA027770.
- Thayaparan, T., Ibrahim, Y., Polak, J., and Riddolls, R. (2018). High-frequency over-the-horizon radar in Canada. *IEEE Geoscience and Remote Sensing Letters*, 15(11):1700–1704.
- Themens, D. R., Watson, C., Žagar, N., Vasylyevych, S., Elvidge, S., McCaffrey, A., Prikryl, P., Reid, B., Wood, A., and Jayachandran, P. (2022). Global propagation of ionospheric disturbances associated with the 2022 Tonga volcanic eruption. *Geophysical Research Letters*, 49(7):e2022GL098158.
- Thompson, A. R., Clark, B., Wade, C., and Napier, P. J. (1980). The Very Large Array. *Astrophysical Journal Supplement Series*, 44:151–167.
- Tingay, S. J., Goeke, R., Bowman, J. D., Emrich, D., Ord, S. M., Mitchell, D. A., Morales, M. F., Booler, T., Crosse, B., Wayth, R. B., Lonsdale, C. J., Tremblay, S., Pallot, D., Colegate, T., Wicenc, A., Kudryavtseva, N., Arcus, W., Barnes, D., Bernardi, G., Briggs, F., Burns, S., Bunton, J. D., Cappallo, R. J., Corey, B. E., Deshpande, A., Desouza, L., Gaensler, B. M., Greenhill, L. J., Hall, P. J., Hazelton, B. J., Herne, D., Hewitt, J. N., Johnston-Hollitt, M., Kaplan, D. L., Kasper, J. C., Kincaid, B. B., Koenig, R., Kratzenberg, E., Lynch, M. J., Mckinley, B., Mcwhirter, S. R., Morgan, E., Oberoi, D., Pathikulangara, J., Prabu, T., Remillard, R. A., Rogers, A. E. E., Roshi, A., Salah, J. E., Sault, R. J., Udaya-Shankar, N., Schlagenhafer, F., Srivani, K. S., Stevens, J., Subrahmanyam, R., Waterson, M., Webster, R. L., Whitney, A. R., Williams, A., Williams, C. L.,

- and Wyithe, J. S. B. (2013). The Murchison Widefield Array: The Square Kilometre Array precursor at low radio frequencies. *Publications of the Astronomical Society of Australia*, 30:e007.
- Tokumaru, M., Fujiki, K., Iwai, K., Tyul’bashev, S., and Chashei, I. (2019). Coordinated interplanetary scintillation observations in Japan and Russia for coronal mass ejection events in early September 2017. *Solar Physics*, 294(7):1–15.
- Torrence, C. and Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, 79(1):61–78.
- Trigg, H., Dorrian, G., Boyde, B., Wood, A., Fallows, R. A., and Mevius, M. (2024). Observations of high definition symmetric quasi-periodic scintillations in the mid-latitude ionosphere with LOFAR. *Journal of Geophysical Research: Space Physics*, 129(7):e2023JA032336.
- Tsunoda, R. T. (1988). High-latitude F region irregularities: A review and synthesis. *Reviews of Geophysics*, 26(4):719–760.
- Tsunoda, R. T. and Cosgrove, R. B. (2001). Coupled electrodynamics in the nighttime midlatitude ionosphere. *Geophysical Research Letters*, 28(22):4171–4174.
- Vadas, S. L. (2007). Horizontal and vertical propagation and dissipation of gravity waves in the thermosphere from lower atmospheric and thermospheric sources. *Journal of Geophysical Research: Space Physics*, 112(A6).
- Vadas, S. L. and Fritts, D. C. (2002). The importance of spatial variability in the generation of secondary gravity waves from local body forces. *Geophysical Research Letters*, 29(20):45–1–45–4.
- Vadas, S. L. and Fritts, D. C. (2005). Thermospheric responses to gravity waves: Influences of increasing viscosity and thermal diffusivity. *Journal of Geophysical Research: Atmospheres*, 110(D15).
- Vadas, S. L., Fritts, D. C., and Alexander, M. J. (2003). Mechanism for the generation of secondary waves in wave breaking regions. *Journal of the Atmospheric Sciences*, 60(1):194–214.
- van de Hulst, H. C. (1945). Radiogolven uit het wereldruim: II. Herkomst der radiogolven (Radio waves from space). *Nederlandsch Tijdschrift voor Natuurkunde*, 11:210–221.
- van de Kamp, M., Pokhotelov, D., and Kauristie, K. (2014). TID characterised using joint effort of incoherent scatter radar and GPS. *Annales Geophysicae*, 32(12):1511–1532.
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., Heald, G., McKean, J. P., Hessels, J. W. T., de Bruyn, A. G., Nijboer, R., Swinbank, J., Fallows, R., Brentjens, M., Nelles, A., Beck, R., Falcke, H., Fender, R., Hörandel, J., Koopmans, L. V. E., Mann, G., Miley, G., Röttgering, H., Stappers, B. W., Wijers, R. A. M. J., Zaroubi, S., van den Akker, M., Alexov, A., Anderson, J., Anderson, K., van Ardenne, A., Arts, M., Asgekar, A., Avruch, I. M., Batejat, F., Bähren, L., Bell, M. E., Bell, M. R., van Bemmelen, I., Bannema, P., Bentum, M. J., Bernardi, G., Best, P., Birzan, L., Bonafede, A., Boonstra, A.-J., Braun, R., Bregman, J., Breitling, F., van de Brink, R. H., Broderick, J., Broekema, P. C., Brouw, W. N., Brügggen, M., Butcher, H. R., van Cappellen, W., Ciardi, B., Coenen, T., Conway, J., Coolen, A., Corstanje, A., Damstra, S., Davies, O., Deller, A. T., Dettmar, R.-J., van Diepen, G., Dijkstra, K., Donker, P., Doorduyn, A., Dromer, J., Drost, M., van Duin, A., Eislöffel, J., van Enst, J., Ferrari, C., Frieswijk, W., Gankema, H., Garrett, M. A., de Gasperin, F., Gerbers, M., de Geus, E., Grießmeier, J.-M., Grit, T., Gruppen, P., Hamaker, J. P., Hassall, T., Hoeft, M., Holties, H. A., Horneffer, A., van der Horst, A., van Houwelingen, A., Huijgen, A., Iacubelli, M., Intema, H., Jackson, N., Jelic, V., de Jong, A., Jette, E., Kant, D., Karastergiou, A., Koers, A., Kollen, H., Kondratiev, V. I., Kooistra, E., Koopman, Y., Koster, A., Kuniyoshi,

- M., Kramer, M., Kuper, G., Lambropoulos, P., Law, C., van Leeuwen, J., Lemaitre, J., Loose, M., Maat, P., Macario, G., Markoff, S., Masters, J., McFadden, R. A., McKay-Bukowski, D., Meijering, H., Meulman, H., Mevius, M., Middelberg, E., Millenaar, R., Miller-Jones, J. C. A., Mohan, R. N., Mol, J. D., Morawietz, J., Morganti, R., Mulcahy, D. D., Mulder, E., Munk, H., Nieuwenhuis, L., van Nieuwpoort, R., Noordam, J. E., Norden, M., Noutsos, A., Offringa, A. R., Olofsson, H., Omar, A., Orrú, E., Overeem, R., Paas, H., Pandey-Pommier, M., Pandey, V. N., Pizzo, R., Polatidis, A., Rafferty, D., Rawlings, S., Reich, W., de Reijer, J.-P., Reitsma, J., Renting, G. A., Riemers, P., Rol, E., Romein, J. W., Roosjen, J., Ruiter, M., Scaife, A., van der Schaaf, K., Scheers, B., Schellart, P., Schoenmakers, A., Schoonderbeek, G., Serylak, M., Shulevski, A., Sluman, J., Smirnov, O., Sobey, C., Spreeuw, H., Steinmetz, M., Sterks, C. G. M., Stiepel, H.-J., Stuurwold, K., Tagger, M., Tang, Y., Tasse, C., Thomas, I., Thoudam, S., Toribio, M. C., van der Tol, B., Usov, O., van Veelen, M., van der Veen, A.-J., ter Veen, S., Verbiest, J. P. W., Vermeulen, R., Vermaas, N., Vocks, C., Vogt, C., de Vos, M., van der Wal, E., van Weeren, R., Weggemans, H., Weltevrede, P., White, S., Wijnholds, S. J., Wilhelmsson, T., Wucknitz, O., Yatawatta, S., Zarka, P., Zensus, A., and van Zwieten, J. (2013). LOFAR: The LOw-Frequency ARray. *Astronomy & Astrophysics*, 556:A2.
- Vaughan, S. (2005). A simple test for periodic signals in red noise. *Astronomy & Astrophysics*, 431(1):391–403.
- Verhulst, T. G., Altadill, D., Barta, V., Belehaki, A., Buresová, D., Cesaroni, C., Galkin, I., Guerra, M., Ippolito, A., Herekakis, T., Kouba, D., Mielich, J., Segarra, A., Spogli, L., and Tsagouri, I. (2022). Multi-instrument detection in Europe of ionospheric disturbances caused by the 15 January 2022 eruption of the Hunga volcano. *Journal of Space Weather and Space Climate*, 12:35.
- Vickrey, J. F. and Kelley, M. C. (1982). The effects of a conducting E layer on classical F region cross-field plasma diffusion. *Journal of Geophysical Research: Space Physics*, 87(A6):4461–4468.
- Vincent, R. A. (2015). The dynamics of the mesosphere and lower thermosphere: a brief review. *Progress in Earth and Planetary Science*, 2(1):1–13.
- Waldock, J. and Jones, T. (1986). HF Doppler observations of medium-scale travelling ionospheric disturbances at mid-latitudes. *Journal of Atmospheric and Terrestrial Physics*, 48(3):245–260.
- Wang, C., Zhang, M., Xu, Z.-W., Chen, C., and Guo, L.-X. (2014). Cubic phase distortion and irregular degradation on SAR imaging due to the ionosphere. *IEEE Transactions on Geoscience and Remote Sensing*, 53(6):3442–3451.
- Wang, J., Shi, Y., Yang, C., and Feng, F. (2022). A review and prospects of operational frequency selecting techniques for HF radio communication. *Advances in Space Research*, 69:2989–2999.
- Warwick, J. W. (1964). Radio star scintillations from ionospheric waves. *Radio Science Journal of Research*, 68:179.
- Waszewski, A., Morgan, J., and Jordan, C. (2022). A measurement of small-scale features using ionospheric scintillation. comparison with refractive shift measurements. *Publications of the Astronomical Society of Australia*, 39:e036.
- Weber, E. J., Tsunoda, R. T., Buchau, J., Sheehan, R. E., Strickland, D. J., Whiting, W., and Moore, J. G. (1985). Coordinated measurements of auroral zone plasma enhancements. *Journal of Geophysical Research: Space Physics*, 90(A7):6497–6513.
- Whitehead, J. (1961). The formation of the sporadic-E layer in the temperate zones. *Journal of Atmospheric and Terrestrial Physics*, 20(1):49–58.

- Whitehead, J. (1989). Recent work on mid-latitude and equatorial sporadic-E. *Journal of Atmospheric and Terrestrial Physics*, 51(5):401–424.
- Wild, J. and Roberts, J. (1956a). Regions of the ionosphere responsible for radio star scintillations. *Nature*, 178:377–378.
- Wild, J. and Roberts, J. (1956b). The spectrum of radio-star scintillations and the nature of irregularities in the ionosphere. *Journal of Atmospheric and Terrestrial Physics*, 8(1):55–75.
- Wild, J., P., Sheridan, K., V., and Neylan, A., A. (1959). An investigation of the speed of the solar disturbances responsible for type III radio bursts. *Australian Journal of Physics*, pages 369–398.
- Wood, A. G., Dorrian, G. D., Boyde, B., Fallows, R. A., Themens, D. R., Mevius, M., Sprenger, T., Main, R., Pryse, S. E., and Elvidge, S. (2024). Quasi-stationary substructure within a sporadic E layer observed by the Low-Frequency Array (LOFAR). *Journal of Space Weather and Space Climate*, 14:27.
- Woodman, R. F., Yamamoto, M., and Fukao, S. (1991). Gravity wave modulation of gradient drift instabilities in mid-latitude sporadic E irregularities. *Geophysical Research Letters*, 18(7):1197–1200.
- Wright, C. J., Hindley, N. P., Alexander, M. J., Barlow, M., Hoffmann, L., Mitchell, C. N., Prata, F., Bouillon, M., Carstens, J., Clerbaux, C., Osprey, S. M., Powell, N., Randall, C. E., and Yue, J. (2022). Surface-to-space atmospheric waves from Hunga Tonga–Hunga Ha’apai eruption. *Nature*, 609(7928):741–746.
- Wright, J. W. (1990). Ionogram inversion for a tilted ionosphere. *Radio Science*, 25(6):1175–1182.
- Wu, Z., Verbiest, J. P. W., Main, R. A., Griefmeier, J.-M., Liu, Y., Osłowski, S., Mochickal Amalappat, K., Nielsen, A.-S. B., Künsemöller, J., Donner, J. Y., Tiburzi, C., Porayko, N., Serylak, M., Künkel, L., Brüggem, M., and Vocks, C. (2022). Pulsar scintillation studies with LOFAR-I. the census. *Astronomy & Astrophysics*, 663:A116.
- Wüst, S. and Bittner, M. (2006). Non-linear resonant wave–wave interaction (triad): Case studies based on rocket data and first application to satellite data. *Journal of Atmospheric and Solar-Terrestrial Physics*, 68(9):959–976.
- Xiong, C., Lühr, H., Wang, H., and Johnsen, M. G. (2014). Determining the boundaries of the auroral oval from CHAMP field-aligned current signatures - part 1. *Annales Geophysicae*, 32(6):609–622.
- Xiong, C., Stolle, C., Alken, P., and Rauberg, J. (2020). Relationship between large-scale ionospheric field-aligned currents and electron/ion precipitations: DMSP observations. *Earth, Planets and Space*, 72(1):1–22.
- Xu, S., Vadas, S. L., and Yue, J. (2024). Quiet time thermospheric gravity waves observed by GOCE and CHAMP. *Journal of Geophysical Research: Space Physics*, 129(1):e2023JA032078. e2023JA032078 2023JA032078.
- Yamamoto, M., Fukao, S., Woodman, R. F., Ogawa, T., Tsuda, T., and Kato, S. (1991). Mid-latitude E region field-aligned irregularities observed with the MU radar. *Journal of Geophysical Research: Space Physics*, 96(A9):15943–15949.
- Yang, N., Le, H., and Liu, L. (2015). Statistical analysis of ionospheric mid-latitude trough over the northern hemisphere derived from GPS total electron content data. *Earth, Planets and Space*, 67:1–11.

- Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., Labropoulos, P., Pandey, V. N., Kazemi, S., Zaroubi, S., Koopmans, L. V. E., Offringa, A. R., Jelić, V., Martinez Rubi, O., Veligatla, V., Wijnholds, S. J., Brouw, W. N., Bernardi, G., Ciardi, B., Daiboo, S., Harker, G., Mellema, G., Schaye, J., Thomas, R., Vedantham, H., Chapman, E., Abdalla, F. B., Alexov, A., Anderson, J., Avruch, I. M., Batejat, F., Bell, M. E., Bell, M. R., Bentum, M., Best, P., Bonafede, A., Bregman, J., Breitling, F., van de Brink, R. H., Broderick, J. W., Brügger, M., Conway, J., de Gasperin, F., de Geus, E., Duscha, S., Falcke, H., Fallows, R. A., Ferrari, C., Frieswijk, W., Garrett, M. A., Griessmeier, J. M., Gunst, A. W., Hassall, T. E., Hessels, J. W. T., Hoeft, M., Iacobelli, M., Juette, E., Karastergiou, A., Kondratiev, V. I., Kramer, M., Kuniyoshi, M., Kuper, G., van Leeuwen, J., Maat, P., Mann, G., McKean, J. P., Mevius, M., Mol, J. D., Munk, H., Nijboer, R., Noordam, J. E., Norden, M. J., Orru, E., Paas, H., Pandey-Pommier, M., Pizzo, R., Polatidis, A. G., Reich, W., Röttgering, H. J. A., Sluman, J., Smirnov, O., Stappers, B., Steinmetz, M., Tagger, M., Tang, Y., Tasse, C., ter Veen, S., Vermeulen, R., van Weeren, R. J., Wise, M., Wucknitz, O., and Zarka, P. (2013). Initial deep LOFAR observations of epoch of reionization windows-I. the north celestial pole. *Astronomy & Astrophysics*, 550:A136.
- Yiğit, E., Knížová, P. K., Georgieva, K., and Ward, W. (2016). A review of vertical coupling in the atmosphere-ionosphere system: Effects of waves, sudden stratospheric warmings, space weather, and of solar activity. *Journal of Atmospheric and Solar-Terrestrial Physics*, 141:1–12.
- Yin, F., Lühr, H., Park, J., and Wang, L. (2019). Comprehensive analysis of the magnetic signatures of small-scale traveling ionospheric disturbances, as observed by Swarm. *Journal of Geophysical Research: Space Physics*, 124(12):10794–10815.
- Yu, B., Xue, X., Yue, X., Yang, C., Yu, C., Dou, X., Ning, B., and Hu, L. (2019). The global climatology of the intensity of the ionospheric sporadic E layer. *Atmospheric Chemistry and Physics*, 19(6):4139–4151.
- Yue, J., Perwitasari, S., Xu, S., Hozumi, Y., Nakamura, T., Sakanoi, T., Saito, A., Miller, S. D., Straka, W., and Rong, P. (2019). Preliminary dual-satellite observations of atmospheric gravity waves in airglow. *Atmosphere*, 10(11).
- Zalizovski, A. V., Yampolski, Y. M., Mishin, E., Kashcheyev, S. B., Sopin, A. O., Koloskov, A. V., Lisachenko, V. N., and Reznichenko, A. I. (2021). Multi-position facility for HF Doppler sounding of ionospheric inhomogeneities in Ukraine. *Radio Science*, 56(10):e2021RS007303.
- Zaroubi, S., de Bruyn, A. G., Harker, G., Thomas, R. M., Labropoulos, P., Jelić, V., Koopmans, L. V. E., Brentjens, M. A., Bernardi, G., Ciardi, B., Daiboo, S., Kazemi, S., Martinez-Rubi, O., Mellema, G., Offringa, A. R., Pandey, V. N., Schaye, J., Veligatla, V., Vedantham, H., and Yatawatta, S. (2012). Imaging neutral hydrogen on large scales during the Epoch of Reionization with LOFAR. *Monthly Notices of the Royal Astronomical Society*, 425(4):2964–2973.
- Zhang, P., Offringa, A. R., Zucca, P., Kozarev, K., and Mancini, M. (2023). RFI flagging in solar and space weather low frequency radio observations. *Monthly Notices of the Royal Astronomical Society*, 521(1):630–637.
- Zhang, X. and Forbes, J. M. (2014). Lunar tide in the thermosphere and weakening of the northern polar vortex. *Geophysical Research Letters*, 41(23):8201–8207.
- Zhang, Y. and Paxton, L. (2008). An empirical Kp-dependent global auroral model based on TIMED/GUVI FUV data. *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(8-9):1231–1242.
- Zhou, Q. and Mathews, J. D. (2006). On the physical explanation of the Perkins instability. *Journal of Geophysical Research: Space Physics*, 111(A12).