Characterising Spatial and Temporal Ionospheric Variability with a Network of Oblique Angle-of-arrival and Doppler Ionosondes

by

Andrew James Heitmann

BSc (Hons), MSc (Defence)

Thesis submitted for the degree of

Doctor of Philosophy

Department of Physics

School of Physical Sciences

THE UNIVERSITY
of ADELAIDE

November 2020

Table of contents

List of figures

List of tables

Glossary

AGW	Atmospheric gravity wave
AoA	Angle-of-arrival
ART	Analytic ray tracing
ARTIST	Automatic Real-Time Ionogram Scaler with True height analysis
CCE	Clear channel evaluation
CDF	Cumulative density function
CIT	Coherent integration time
CR	Coordinate registration
CRPL	Central Radio Propagation Laboratory
CSF	Channel scattering function
DFT	Discrete Fourier transform
DORS	Digital Oblique Receiving System
DOTS	Digital Oblique Transmitting System
DPS	Digisonde Portable Sounder
DST Group	Defence Science and Technology Group, Australia (formerly DSTO)
DST-IIP	Defence Science and Technology Ionogram Image Processing
DSTO	Defence Science and Technology Organisation (now DST Group)
ELOISE	Elevation-sensitive Oblique Incidence Sounder Experiment (formerly Elevation-scanning Oblique Incidence Sounder Experiment)
EMD	Empirical mode decomposition
ESPRIT	Estimation of Signal Parameters via Rotational Invariance Techniques
EUV	Extreme ultraviolet

FAS	Frequency and Angular Sounding (technique)
FFT	Fast Fourier transform
GGIM	Gardiner-Garden Ionospheric Model
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HF	High frequency (3–30 MHz)
I/Q	In-phase and quadrature
IMF	Intrinsic mode function
IRI	International Reference Ionosphere
JORN	Jindalee Operational Radar Network
LFMCW	Linear frequency-modulated continuous-wave
LSTID	Large-scale travelling ionospheric disturbance
LWA	Long Wavelength Array
MIMO	Multiple-input multiple-output
MOF	Maximum observed frequency
MUF	Maximum usable frequency
MQP	Multi-segment quasi-parabolic (profile)
MSTID	Medium-scale travelling ionospheric disturbance
MUSIC	Multiple Signal Classification
MVDR	Minimum Variance Distortionless Response
NCAR	National Center for Atmospheric Research
NEC	Numerical Electromagnetics Code
NRL	Naval Research Laboratory

NRLMSISE-00	Naval Research Laboratory Mass Spectrometer and Incoherent Scatter Extended – 2000
NRT	Numerical ray tracing
O/X	Ordinary/extraordinary wave components
OI	Oblique incidence
OIS	Oblique incidence sounder
OTHR	Over-the-horizon radar
PHaRLAP	Provision of High-frequency Raytracing Laboratory for Propagation studies
PRIME	Portable Remote Ionospheric Monitoring Equipment
QP	Quasi-parabolic
QVIS	Quasi-vertical incidence sounder
RAAF	Royal Australian Air Force
RFI	Radio frequency interference
RMSE	Root-mean-square error
RTIM	Real-Time Ionospheric Model
SMM	Spherical mirror model
SNR	Signal-to-noise ratio
SSM	Spherically-symmetric model
TEC	Total electron content
TID	Travelling ionospheric disturbance
TMM	Tilted mirror model
TRACE	Thermospheric Radar Airglow Correlation Experiment
UHF	Ultra high frequency (0.3–3 GHz)

ULA	Uniform linear array
UT	Universal Time
VHF	Very high frequency (30–300 MHz)
VI	Vertical incidence
VIPIR	Vertical Incidence Pulsed Ionospheric Radar
VIS	Vertical incidence sounder
WGS-84	World Geodetic System 1984
WRF	Waveform repetition frequency

ELOISE site abbreviations

CLO   Cloncurry, Queensland

COO   Coondambo, South Australia

CUR   RAAF Base Curtin, Western Australia

HRO   Harts Range, Northern Territory

HUM   Humpty Doo, Northern Territory

KAL   Kalkarindji, Northern Territory

LAV   Laverton, Western Australia

LYN   Lynd River, Queensland

R2E   (JORN) Radar 2 East, Laverton, Western Australia

SCH   RAAF Base Scherger, Queensland

SHD (STH)   South Hedland, Western Australia

Abstract

Ionospheric variability exists on a broad range of scales, and routinely impacts sky-wave propagation modes of high frequency radio waves, to the detriment of radar and communication systems. In order to better understand the electron density structures associated with such variability at mid-latitudes, a network of oblique angle-of-arrival (AoA) and Doppler ionosondes were installed in central and northern Australia as part of the ELOISE campaign in 2015. This thesis analyses observations from the ELOISE AoA ionosondes, with a focus on characterising the influence of medium- to large- scale gradients and signatures of travelling ionospheric disturbances (TIDs).

Following an overview of the experiment, the design and calibration of the new ionosonde system is described. With multi-channel receivers connected to each element of two twin-arm arrays, a total of eleven AoA paths of between 900 and 2700 km were collected, including nine with interleaved Doppler measurements using a special channel scattering function (CSF) capability. On-board signal processing was developed to perform real-time clear channel evaluation and CSF scheduling, and generate the AoA ionograms and delay-Doppler images with fitted electron density profiles. In further offline analysis, peak detection and mode classification was carried out, to support reflection point mapping and tilt estimation.

Significant testing and validation of the new ionosonde before and after the experiment revealed AoA uncertainties on the scale of 0.2–0.5° in bearing and 0.4–0.9° in elevation. Having identified a low-elevation bias, models of tropospheric refraction and antenna mutual coupling effects were considered as possible correction strategies, but ultimately an empirical approach based on aggregated ionospheric returns was implemented. Small-scale (intra-dwell) ionospheric variability also has the potential to compromise results, through unresolved multi-mode mixing, and this has been investigated using a combination of spatial and temporal variability metrics derived from the CSF data.

The analysis of large quantities of F2 peak data shows persistent diurnal patterns in the oblique AoA observables that are also well-captured by a conventional data-assimilative ionospheric model, even without the benefit of AoA and Doppler inputs. Furthermore, Doppler measurements are reproduced remarkably well using just the midpoint fitted profiles. A statistical study has quantified the level of consistency between observations

and model, to provide greater confidence in the results. Many of the geophysical features can be interpreted as ionospheric gradients, as evident in the tilt estimates, and horizontally moving structures such as TIDs, using a form of Doppler-based drift analysis.

While signatures of TIDs vary considerably, two simple wave-like perturbation models have been evaluated to help classify quasi-periodic behaviour in the AoA observations, as well as understand the directional filtering effect imposed by the path geometry. In some cases, a set of TID parameters can be determined by eye, but in others automatic parameter inversion techniques may be more viable. Two such techniques were implemented but results using both real and synthetic data demonstrated some significant limitations. Finally, attempts to relate TID signatures across multiple paths shows promise, but there still appears to be a strong dependence on path geometry that is difficult to eliminate.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship and as an employee of the Defence Science and Technology Group.

Andrew James Heitmann   30 November 2020

Data statement

Data from the ELOISE trial, including measurements collected on the angle-of-arrival arrays and ionospheric parameter maps generated by Dr Robert Gardiner-Garden, are owned by the Commonwealth of Australia. Access requests will be considered on a case-by-case basis by contacting the author at andrew.heitmann@dst.defence.gov.au.

The 3D numerical ray tracing results were obtained using the HF propagation toolbox, PHaRLAP, created by Dr Manuel Cervera (manuel.cervera@dst.defence.gov.au) and available for download from the DST Group website at the following address: https://www.dst.defence.gov.au/opportunity/pharlap-provision-high-frequency-raytracing-laboratory-propagation-studies.

Acknowledgements

1   Introduction

1.1   Opening remarks

Fifty years ago, Australia embarked on an ambitious over-the-horizon radar surveillance program that would prove to become one of the great success stories of Defence science and innovation in this country. Persistent remote observations of the regional ionosphere have always been a cornerstone of this research program, with a rich data set spanning many solar cycles having now been amassed by the Defence Science and Technology Group and its predecessors. As the radar hardware became more sensitive, and signal processing algorithms improved, there has been increasing impetus to better understand the many aspects of ionospheric variability, both from a geophysical standpoint and a real-time modelling exercise, as it constitutes one of the main limiting factors for system performance. Yet there still remain many unknowns about this first frontier of space, particularly as the Sun enters another year of record low activity in 2020, and the ionosphere continues to surprise and baffle as higher fidelity measuring techniques drill down into finer levels of detail and complexity.

The ELOISE trial in 2015 was conceived as a pivotal step forward in building upon this amassed knowledge of ionospheric effects and radar expertise. Bringing together several well-established techniques for remotely observing the ionosphere, along with an unprecedented number and density of sampling points, the experiment sought to relate the signatures of mid-latitude ionospheric variability (including the ubiquitous travelling ionospheric disturbances or TIDs) between the different sensors. A core component was the deployment of two new high frequency (HF) receiving arrays in central Australia, designed to observe a network of oblique propagation paths from 900 to 2700 km with 2D angle-of-arrival (AoA) and Doppler measurements. This thesis covers the development and analysis of observations from this instrument; in particular,

focusing on how it responds to ionospheric disturbances on a range of different spatial and temporal scales, and the anisotropies in AoA that the oblique geometry introduces.

1.2   Research questions

High frequency sky-wave radars use the refractive properties of the ionosphere to detect and track targets at very long ranges, well beyond the geometric horizon. HF propagation through the ionosphere is sensitive to variability on a range of scales, driven primarily by solar forcing from above, neutral atmospheric coupling from below, and interactions with the geomagnetic field. The understanding of these effects ultimately determines the coordinate registration accuracy of sky-wave radars; that is, the ability to determine the true location of any target of interest. While there are many approaches to measuring ionospheric variability, the ground-based ionosonde (or ionospheric sounder) is by far the most common and lends itself well to synoptic-style data collection over a wide area of coverage. Although conventional ionosondes usually estimate the ionospheric electron density profile from group delay measurements alone, regular and often quasi-periodic "off-angle" perturbations from the great circle path, measurable across other HF observables as well, offer additional information about the structure of the disturbances that make up this variability.

Oblique incidence sounder networks, in particular, provide an efficient and cost-effective way to monitor large regions of the ionosphere with a high density of sampling points, due to their ability to collect soundings on all pair-wise combinations of transmit/receive sites. With the addition of both AoA and Doppler capabilities, the off-angle behaviour and dynamics in propagation characteristics caused by ionospheric disturbances can be studied in detail; that is, in terms of how the manifestation of disturbances in sounder observations (ionograms) relates to realistic models of the electron density distribution. It is on this foundation that the project considers the following research questions:

- What aspects of ionospheric variability can be identified and parameterised from a network of oblique AoA and Doppler ionosondes?

- Are these observations consistent with existing models of electron density distributions and neutral atmosphere dynamics, and moreover, what is missed by conventional sounding techniques?

- To what extent does the path geometry (i.e. length and orientation) influence the response to ionospheric disturbances, particularly those with frontal characteristics such as TIDs, and over what spatial scales are coherent disturbance signatures observed?

- What are the instrumental limitations that affect sensitivity and accuracy, and how can these effects be corrected or mitigated?

1.3   Aims and objectives

The principal aim of this PhD project was to investigate the effects of ionospheric variability on the AoA and Doppler signatures of oblique HF radio wave propagation, and evaluate techniques for inferring ionospheric dynamics and structure directly from HF observations (i.e. the inverse problem). A unique data set was acquired as part of the ELOISE trial, utilising a new AoA and Doppler ionosonde that built upon the author's previous work on DST Group's chirp-sounder systems. With a diverse network of oblique paths measured simultaneously across central and northern Australia, various scales of disturbances in the ionosphere were able to be modelled, observed and characterised, both spatially and temporally, in support of the above aim.

The motivation behind ELOISE as a whole was primarily influenced by its Defence application: to understand performance implications on the current and future Jindalee Operational Radar Network (JORN), and drive improvements to its ionospheric model. To this end, two 19-element twin-arm arrays were constructed in Laverton, Western Australia, and Coondambo, South Australia, to receive one-way sky-wave propagation paths from a collection of eight down-range transmitters, most of which were separately upgraded to operate on a faster revisit rate. By design, these paths overlap with the coverage of the JORN Radar 2 East (R2E) backscatter radar. As part of the PhD project, a substantial suite of on-board processing software was written to enable remote monitoring and collection of the AoA (bearing/elevation) and Doppler (channel

scattering function) data, and significant effort was undertaken to test and calibrate the arrays, particularly to address the accuracy at low elevation angles.

The objective of the subsequent analysis was to interpret the observations from the new ionosonde in terms of underlying variability in the bottomside ionosphere. Whereas previous research on the HF signatures of ionospheric variability has tended to focus on paths at vertical incidence, comparatively less work has been carried out at oblique incidence, despite being more representative of operational applications in radar and communications. The challenge in the oblique case is relating observables across different paths with vastly different geometries. The problem can be approached in two directions: either by constructing models of the disturbance field and ray tracing through these to synthesise the HF observables, or by inverting the observables directly to give some parameterised form of the disturbance. This project sought to explore the merits of each, with a focus on the F2-layer effects of medium- to large-scale gradients (tilts) and TIDs for which the ELOISE AoA network was best suited.

Although truth data is hard to come by in the ionosphere, a semi-independent representation of the electron density distribution, constructed by spatially mapping conventional sounder midpoint profiles over an arbitrary set of latitudes and longitudes, provided a means of testing alternative disturbance characterisation schemes based on single AoA paths. Given widespread use of the spatial mapping approach, including in the JORN Real-Time Ionospheric Model, the purpose of this comparison was to offer insights into how such a model responds to different scales of variability, and quantitative measures of its performance at predicting AoA and Doppler excursions, without actually assimilating any such data.

Of course, such ambitions are fairly broad and it is worth briefly clarifying the scope. The layout of the ELOISE experiment restricts coverage to the mid-latitude ionosphere over Australia, so naturally no direct statements can be made about the relationships with the equatorial and high-latitude regions. Furthermore, the total data collection period, spanning just a few weeks with the full set of AoA and Doppler paths, does not lend itself to a comprehensive climatological analysis; instead, the emphasis will be on case studies to illustrate the typical classes of variability and their interpretations. Finally, the spatial and temporal sampling of the ELOISE AoA paths (being several hundred km and 3.75 min, respectively) means that rapid small-scale irregularities

cannot be easily identified in the data, except within the course of a single channel scattering function dwell. Therefore, apart from a short investigation into these effects from a system limitation perspective, the bulk of this thesis will address the medium- to large-scale variability component.

1.4 Significant outcomes and contributions

The key outcomes of this work and contributions to the field are as follows:

- A highly capable experimental ionosonde has been developed for oblique incidence reception, incorporating a frequency-agile channel scattering function (Doppler) mode and 2D interferometric angle-of-arrival estimation algorithm as part of its on-board processing. While this system drew components from DST Group's existing ionosonde hardware and software, the AoA and Doppler measurement capabilities, along with the subsequent (offline) peak analysis and reflection point mapping software, were implemented and tested by the author specifically for this project. The ionosonde system takes advantage of the latest direct-digital HF receiver technology to offer simultaneous reception of up to eight paths on each element of the array. Such observations are believed to be the first of their kind in Australia, particularly in terms of the number and range of paths, and offer new contributions to the understanding of spatial and temporal variability in the mid-latitude ionosphere.

- Low-elevation array performance has been studied, including modelling the effects of mutual coupling and tropospheric refraction, to establish the limitations on AoA accuracy. An improved knowledge of such factors, and ways to correct for their influence, enables greater value to be gained from longer oblique paths of up to 3000 km (the notional limit of 1-hop propagation) using simple horizontal planar arrays. The impact of small-scale irregularities and unresolved mode mixing has also been explored, although admittedly the ELOISE data set is not ideally suited to these studies.

- An analysis framework for relating oblique AoA and Doppler observations across many disparate propagation paths to disturbances in the underlying electron density field has been created. This includes both parameterised models for generic tilted

structure and TIDs, designed to be used in conjunction with HF ray tracing, as well as implementations of direct inversion techniques, which have been extended to better cater for the oblique incidence geometry (allowing for a spherical Earth wherever possible). While finding convincing evidence of coherent disturbance signatures over the whole ELOISE coverage region proved to be somewhat elusive, substantial steps were nevertheless made towards identifying and classifying these signatures on individual paths, and better understanding the path length and aspect dependence of the HF response. This last point is an important outcome for the design of AoA ionosonde networks, with diversity of oblique path orientations being critical for TID studies in particular.

- Observed ionospheric Doppler shifts have been compared with those estimated by ray tracing through fitted midpoint electron density profiles, and the overall level of agreement was found to be very good, suggesting that the majority of perturbations in Doppler relate to the large-scale variability component which is well-sampled from one sounding to the next. In addition, there are some characteristic diurnal features that present consistently from day to day. The implication is that, to first order, Doppler may be reliably predicted from the group delay measurements of conventional sounders only.

- Peak data from the new ionosonde have been used to validate a data-assimilative ionospheric model, representative of the JORN Real-Time Ionospheric Model, in terms of its ability to synthetically reproduce the AoA and Doppler observables with only conventional sounder inputs. This is assessed in multiple ways: (1) the qualitative appearance of diurnal features, (2) the error statistics in the synthetic observables, and (3) the characterisation of midpoint gradients (tilts). All indications are that the existing modelling approach is sufficiently accurate and fit for purpose, despite a degree of spatial and temporal smoothing that suppresses the amplitude of some medium-scale disturbances. As well as being a significant statement on the utility of conventional sounders in radar coordinate registration, the results also offer a sense of the value added by such an AoA ionosonde, should these new observations be considered for model assimilation in future.

1.5   Thesis outline

The remainder of the thesis is structured as follows:

- **Chapter 2** ("Background") presents an overview of the existing literature pertaining to observations and modelling of ionospheric variability, to set the context for the subsequent analysis of ELOISE AoA soundings.

- **Chapter 3** ("Experiment design, testing and calibration") provides details on the design and development of the ELOISE oblique angle-of-arrival ionosonde array, which was carried out by the author as part of this project. Preliminary calibration activities, conducted prior to the ELOISE experiment itself, are also documented.

- **Chapter 4** ("Signal processing of ELOISE angle-of-arrival data") describes how DST Group's existing sounder software suite was expanded to support automated scheduling and on-board processing of AoA ionograms and delay-Doppler maps, as well as the offline processing of peak data and reflection point mapping that underpinned the subsequent variability analysis.

- **Chapter 5** ("Understanding instrumental effects and limitations") shows examples of some typical disturbance signatures from the ELOISE AoA array, and delves into two instrumental effects that have the potential to compromise AoA estimation accuracy; namely, tropospheric refraction and mutual coupling between antenna elements. This culminates in an empirical set of bearing and elevation bias corrections for each array.

- **Chapter 6** ("Characterisation of small-scale (intra-dwell) ionospheric variability") explores the small-scale variability, including multi-mode interference effects, that exists within the coherent integration time and is unresolved by standard processing. A series of metrics are devised and compared in terms of their ability to detect non-planarity across the array and degradation in the AoA estimates.

- **Chapter 7** ("Characterisation of medium- to large-scale ionospheric variability") looks at the day-to-day and hour-to-hour variability as it manifests in AoA and Doppler. This is carried out both in the observable space and using geometrically transformed ionospheric tilt and drift parameters. Diurnal patterns identified in aggregated observables as well as individual perturbation events are successfully

reproduced using data-assimilative models constructed from conventional sounder profiles only, providing statistical validation of such approaches.

- **Chapter 8** ("Identification and parameterisation of travelling ionospheric disturbance signatures") focuses on one particular class of medium- to large-scale disturbance, the TID, and evaluates ways of identifying and extracting the wave parameters from the delay, Doppler and AoA observables using relatively simple TID models and geometrical arguments. A framework for synthesising TID signatures is a crucial part of this study and allows the effects of varying the wave parameters to be better understood.

- **Chapter 9** ("Summary and conclusions") offers a brief overview of the key research outcomes, which are also summarised at the end of each chapter, and concludes with a discussion of future work.

2   Background

2.1   The ionosphere

Ever since Sir Edward Appleton first confirmed the existence of an outer charged layer of the Earth's atmosphere in 1924 [Appleton & Barnett, 1925; Ratcliffe, 1959; Gillmor, 1982], the ionosphere has been the subject of countless studies, using a variety of instruments, that seek to better understand its formation, structure, and interactions with other parts of the atmosphere and the Earth's magnetic field. Between heights of about 60 km and 1000 km (largely overlapping with the thermosphere), ionisation from extreme ultraviolet (EUV) and X-ray solar radiation creates a weak plasma, which forms into multiple stratified layers due to the changing neutral densities and ion chemistry at different heights [Rishbeth, 1988; Davies, 1990; Kelley, 2009; Schunk & Nagy, 2009]. The primary layers in the electron density profile are known as D, E, F1 and F2 [Chapman, 1950], in order of increasing altitude, and to some extent each varies with location (latitude), time-of-day (solar zenith angle), season, solar activity (sunspot number), and geomagnetic activity according to the balance of photoionisation, recombination and transport. Sporadic-E (Es) layers are less persistent but equally important and, along with the F2 layer, exhibit the greatest level of hour-to-hour and day-to-day variability, which makes them more difficult to predict from climatologies alone.

For high frequency (HF) radio waves (3–30 MHz), the negative refractive index gradient produced by the layers of free electrons in the bottomside of the ionosphere causes a bending of the propagation path, in many instances returning the energy to the ground where it can be received at very long distances ("over-the-horizon"). This property gives rise to a number of applications, such as radar and communications, which exploit these sky-wave modes for their operation, along with others, such as

satellite navigation systems and radio astronomy, for which the refractive effects are a nuisance requiring special compensation.

To complicate matters, the ionosphere is both dispersive and birefringent, and generally will support multiple propagating modes at any given frequency. Signal absorption also occurs in the ionosphere as a result of electron collisions with surrounding ions and neutral atmospheric constituents; this may be categorised as either non-deviative absorption, as in the D layer where collision rates are high but the refractive index is close to unity, or deviative absorption, as encountered near the ray apogee height where there is strong group retardation.

In the absence of the geomagnetic field, the HF signal is reflected when the wave frequency at vertical incidence equals the plasma frequency at that height. The plasma frequency is related to the square-root of the (free) electron density N_e by

$$ f_p = \frac{1}{2\pi} \sqrt{\frac{N_e e^2}{\varepsilon_0 m_e}} \approx 9.0\sqrt{N_e}, \tag{2.1} $$

where e is the elementary charge, ε_0 is the permittivity of free space, and m_e is the electron rest mass. The approximation on the right-hand side applies for f_p in Hz and N_e in electrons per m^3. Higher operating frequencies will thus penetrate deeper into the ionosphere until the peak of the profile is reached; above this so-called critical frequency (of the order of 10 MHz for $N_e \sim 10^{12}$ m^{-3}), all signal power is transmitted out to space at vertical incidence.

The refractive index in a cold, collisional, magnetised plasma is given by the Appleton-Hartree formula [Budden, 1985, eq. (4.48); Davies, 1990, eq. (3.8)]:

$$ \mu = \sqrt{1 - \frac{X}{1 - iZ - \frac{Y^2 \sin^2\vartheta}{2(1-X-iZ)} \pm \left(\frac{Y^4 \sin^4\vartheta}{4(1-X-iZ)^2} + Y^2\cos^2\vartheta\right)^{1/2}}}, \tag{2.2} $$

$$ X \equiv \left(f_p/f\right)^2, \tag{2.3} $$

$$Y \equiv f_g/f, \tag{2.4}$$

$$Z \equiv \nu/2\pi f, \tag{2.5}$$

where f is the operating frequency, f_g is the electron gyrofrequency, ν is the electron collision frequency (in rad/s), and ϑ is the angle between the direction of propagation and the geomagnetic field. Note that the refractive index always lies between 0 and 1, and as such, the phase speed of sky-wave propagation is faster than the speed of light, whereas the group speed must remain slower. At the point of vertical reflection, $\mu = 0$. The plus-minus sign in equation (2.2) gives rise to two characteristic components, with different speeds and polarisations, known as the ordinary (+) and extraordinary (–) or O and X waves. The ordinary wave is so-called as at vertical incidence it is unaffected by the geomagnetic field.

By neglecting collisions (i.e. $Z \approx 0$) and the geomagnetic field (i.e. $Y \ll 1$), the Appleton-Hartree formula can be reduced to the simpler expression

$$\mu = \sqrt{1 - X} = \sqrt{1 - f_p^2/f^2}, \tag{2.6}$$

which more clearly highlights the importance of the plasma frequency in determining the reflection height (where $\mu = 0$).

A typical daytime, mid-latitude plasma frequency profile, including an embedded Es layer, is shown in the left panel of **Figure 2.1**. As the D layer causes only non-deviative absorption in the HF band, with no significant refraction, it is omitted for the purpose of estimating path delays. The F2 critical frequency in this example is 8.2 MHz, meaning that at vertical incidence, frequencies of up to 8.2 MHz are reflected (as shown in the middle panel, which will be revisited in the next section). The upper frequency support is extended at oblique incidence (right panel), and in some cases, particularly during intense Es, it is possible to observe propagation into the lower VHF band (up to ~50 MHz). At night, the E layer critical frequency drops below 1 MHz and the F1 layer ceases to exist, so only Es and F2 layers remain relevant for HF propagation.
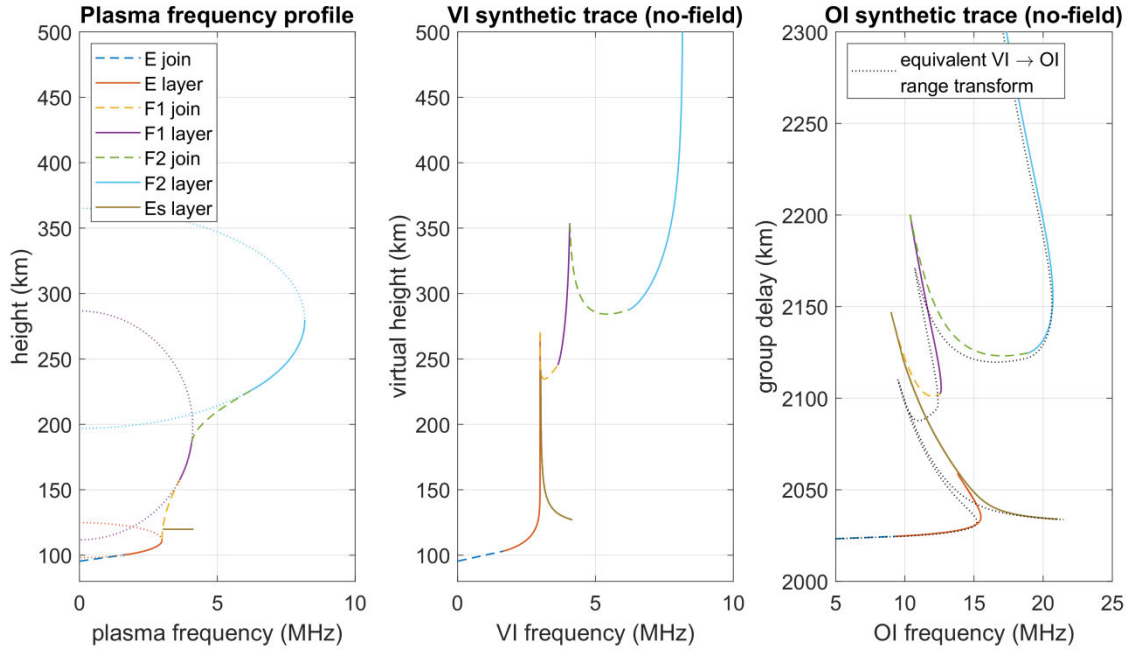
Figure 2.1. Illustrative ionospheric plasma frequency profile (left panel) and synthetic 1-hop ionogram traces for HF propagation at both vertical incidence (middle panel) and oblique incidence (right panel, for a ground range of 2000 km). The four-layer profile, colour coded by layer, is constructed as a monotonically increasing set of quasi-parabolic segments, up to the peak of the F2 layer (excluding the E-F1 valley for sake of simplicity). The synthetic traces are generated using analytic ray tracing equations, assuming spherical symmetry (i.e. no horizontal gradients) and no geomagnetic field. The black dotted line in the right (oblique incidence) panel shows the result of performing an equivalent range transform on the vertical incidence trace, revealing the limitations of this virtual mirror approximation.

As noted above, the ionosphere is always evolving, and the time rate of change of electron density is governed by the continuity equation [Davies, 1990, eq. (2.49)]

$$\frac{\partial N_e}{\partial t} = q - L - \boldsymbol{\nabla} \cdot (N_e \boldsymbol{u}_e), \tag{2.7}$$

where q and L are the electron production and chemical loss terms (i.e. photoionisation and recombination rates, respectively), and the divergence term accounts for the net loss due to transport, with $\boldsymbol{u}_e$ the mean electron ("drift") velocity. In the lower D, E and F1 regions, the system tends to be in quasi-static equilibrium, with production and loss

terms approximately equal and opposite, and the electron density changing only very slowly (i.e. $\partial N_e / \partial t \approx 0$). However, in the F2 region, transport due to diffusion, neutral winds (atmospheric tides), and electromagnetic E×B drift (giving rise to the equatorial plasma fountain) becomes much more significant and cannot be neglected. Indeed, the transport term becomes critical in any treatment of ionospheric disturbances, including modelling the behaviour of travelling ionospheric disturbances (TIDs). During extreme solar events, emissions of X-rays and energetic particles interact with the ionosphere and geomagnetic field to modify the terms in the continuity equation, causing potentially rapid changes and ionospheric irregularities.

Based on the continuity equation and some simplifying assumptions (i.e. a single solar wavelength ionising a single-species isothermal atmosphere), Chapman theory [Chapman, 1931; Rishbeth & Garriott, 1969, ch. 3.2; Davies, 1990, ch. 2.9] predicts the theoretical shape of the ionospheric layers, which for computational purposes are often approximated by parabolas or quasi-parabolas [Croft & Hoogasian, 1968]. Under this formulation, each layer is parameterised by a peak (critical) frequency, peak height, and semi-thickness. A sufficiently flexible ionospheric profile can therefore be constructed from a limited number of parameters, as in **Figure 2.1**, with smooth join segments (dashed lines) inserted between layers to condition the profile for ray tracing (e.g. Hill [1979]; Dyson & Bennett [1988]; Gardiner-Garden et al. [2018]).

2.2   High frequency ionospheric soundings

Ground-based high frequency (HF) ionospheric sounders (or ionosondes) have long been used to measure the vertical structure of the bottomside ionosphere, dating back to the pioneering experiments of Breit & Tuve [1925; 1926] (see also Bibl [1998] for a historical overview). In essence, an ionosonde is a low-power radar which records the one-way signal amplitude (or path loss) and propagation delay to and from the ionosphere as a function of frequency [Davies, 1990, ch. 4; Hunsucker, 1991, ch. 3; McNamara, 1991, ch. 5]. At any given frequency, there are typically multiple propagation modes, reaching different altitudes in the ionosphere and returning after a variable number of hops. The graphical representation of such a radio sounding is the ionogram: an image of the received power at each frequency and group delay. Group delay is usually expressed in km (i.e. $P' = cT_g$, where c is the speed of light in a

vacuum and T_g is the time delay between transmission and reception), and at vertical incidence may be divided by two to give the virtual height. Each mode maps out an "echo" trace in the ionogram, which is scaled and inverted to give the electron density profile. Synthetic renders of such traces, derived from analytic ray tracing, are shown in the middle and right panels of **Figure 2.1**; this excludes geomagnetic effects, which would otherwise cause a splitting of the traces into O/X wave components.

Although single-site vertical incidence sounders (VIS) have historically been the norm for their simplicity of operation and analysis, oblique incidence sounders (OIS) provide an alternative means of estimating the electron density profile, but at the path midpoint instead of overhead. Under a geometric ray interpretation, the propagation is mostly influenced by the ionosphere surrounding the ray apogee, so in theory an OIS path is able to probe the same disturbances as a VIS at its midpoint, or indeed another intersecting OIS path with a co-located midpoint. The two sounder geometries (OIS and VIS) can be related to first-order by a trio of fundamental equivalence relationships [Davies, 1990, ch. 6.3.1; McNamara, 1991, ch. 12.1], which strictly only hold for a plane-symmetric ionosphere in the absence of a geomagnetic field. The secant law [Pickering, 1975; Davies, 1990] states that the oblique height of reflection is equal to that of a wave incident vertically at an equivalent vertical frequency $f_v = f_o \cos \gamma$, where f_o is the oblique frequency and γ is the angle of incidence to the reflecting surface. Breit and Tuve's theorem [Breit & Tuve, 1926] states that the group delay of propagation is equal to the free-space delay of a triangular path, with the same elevation angle at the ground, which travels via an equivalent or virtual height. Martyn's theorem [Martyn, 1935] ties these concepts together, declaring that the virtual heights for oblique and vertical incidence, at frequencies f_o and f_v, respectively, are the same, thereby enabling an OIS trace to be readily transformed to a VIS trace or vice versa.

Figure 2.2 schematically shows, for a single ray over a spherical Earth, the oblique propagation geometry and the triangular virtual mirror path on which the equivalence relationships are based. Despite its limitations, the concept of mirror reflection is used repeatedly throughout this work to assist in interpreting the angle-of-arrival data. The result of transforming an entire 1-hop trace via equivalent geometry is presented in the right panel of **Figure 2.1** as a black dotted line (in this case mapping from vertical to oblique incidence). Note that the critical frequency cusp at vertical incidence effectively doubles back on itself to produce low- and high-angle rays at oblique incidence, which

14

for the same frequency, but different elevation angles, reach different virtual heights in the one layer. The nose-like feature at which these meet defines the maximum observed frequency (MOF)[1]; this is closely related to the backscatter leading edge, described in the next section. Evidently the transform is not perfect compared to the exact synthetic trace, however most of the trace attributes are reproduced, and moreover it provides a strong foundation for asserting that OIS and VIS are interchangeable in terms of their ionospheric characterisation. A number of studies have validated this equivalence by comparing OIS ionograms with a VIS located below the midpoint (e.g. Ippolito et al. [2015]; Heitmann & Gardiner-Garden [2019]).
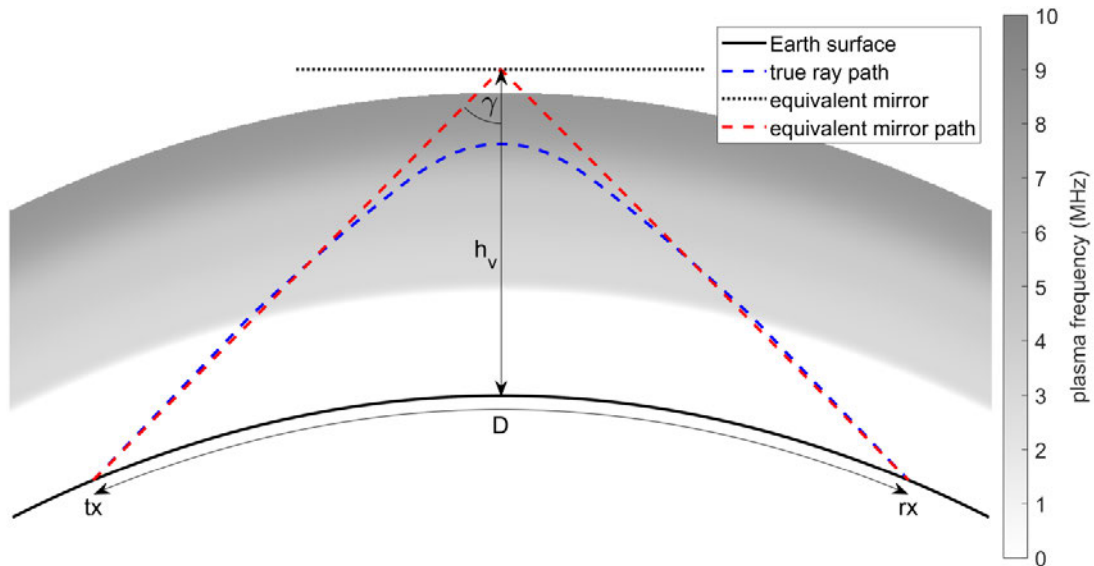
Figure 2.2. Sketch of oblique propagation geometry, from transmitter (tx) to receiver (rx), for a spherical Earth and spherically-symmetric ionosphere (same as in **Figure 2.1**). The true refracted ray path (blue dashed line) traverses the E and F1 regions before reaching its apogee in the F2 region and returning to Earth. The virtual path (red dashed line) has the same group delay as the true path if free space propagation and reflection from a mirror were assumed. The virtual height of reflection (h_v) can be used to relate the oblique and vertical incidence cases, although the equivalence relationships strictly only hold for a planar ionosphere. It is because of the sphericity that the launch and set-down angles are not quite identical between the blue and red paths, as Breit and Tuve's theorem predicts.

---

[1] Here the MOF is distinguished from the Maximum Usable Frequency (MUF), which is generally taken to be a monthly-median quantity [Piggott & Rawer, 1978].

There are several advantages to the bistatic OIS configuration. First, for large-scale ionospheric mapping purposes, their many combinations of point-to-point propagation paths permit a greater spatial sampling density than can be offered by a VIS network alone, given the same number of sites. Second, they may be used to observe the ionosphere at locations where a VIS is impractical (e.g. over oceans). Finally, where transmitters and receivers are not co-located, their design and installation is simplified considerably; that is, system-generated noise and emissions remain minimal at the receiver site and 100% duty cycle waveforms can be used to reduce peak power demands. Modern Global Navigation Satellite Systems (GNSS) make OIS network synchronisation far easier and more accurate than previous methods such as internal clocks or HF timing signals. Oblique path lengths of up to around 2000 km support unobscured observations of both 1-hop E and F propagation, and are therefore generally preferable, while extended ranges from 2000–3000 km are still of interest for 1-hop F and multi-hop propagation.

Despite ionospheric sounders having been used widely for almost a century (with much of the early history outlined by Waynick [1975]), recent advances in technology have enabled a new generation of digital ionosondes to be developed, at an unprecedented level of sensitivity, and with an increased capacity for sophisticated signal and image processing. Notable examples from the last couple of decades include the Digisonde DPS-4D [Reinisch et al., 2009], VIPIR/Dynasonde [Grubb et al., 2008; Rietveld et al., 2008], CADI [MacDougall et al., 1995], AIS-INGV [Zuccheretti et al., 2003], IRIS [Arthur et al., 1997], Cyclone [Akchurin et al., 2009], and WMISS [Gong et al., 2016]. The Defence Science and Technology (DST) Group has considerable experience in this field, with both the Digital Oblique Receiving System (DORS) OIS [Gardiner-Garden, Ayliffe, et al., 2008; Gardiner-Garden et al., 2011; Ayliffe et al., 2019] and Portable Remote Ionospheric Monitoring Equipment (PRIME) VIS [Harris et al., 2016] designed specifically to serve Australia's sky-wave radar capability, the Jindalee Operational Radar Network (JORN); this will be described further in the next section. Both DST Group ionosondes employ low-power (20 W) swept waveforms with a 100% duty cycle to maximise sensitivity, simplify demands on the amplifier, and reduce potential interference with other HF users, compared to their pulsed alternatives, and robust direct-digital waveform generator and receiver designs, to enable network-wide reception over multiple simultaneous channels. Founded on the pioneering work of

Fenwick & Barry [1966], the basic principles of chirp sounders are reviewed by Ivanov et al. [2003], with signal processing fundamentals outlined in **Appendix A**.

Whereas ionogram scaling was in the past a largely manual task, automatic scaling algorithms are now the norm. Most run in near real time on the sounder system itself, applying a consistent set of conventions [Piggott & Rawer, 1978] to characterise the propagation features and reduce the ionogram image to a parameterised representation. These algorithms generally fall into one of two main categories: (1) trace extraction or skeletonisation, followed by profile inversion, or (2) template matching techniques, which directly fit a synthesised trace or curve to extracted features or image data. A prominent example of the former is the software package ARTIST [Reinisch & Huang, 1983; Reinisch et al., 1988; Reinisch et al., 2005; Galkin & Reinisch, 2008], developed for the Lowell Digisonde [Reinisch et al., 2009], while an example of the latter is Autoscala [Pezzopane & Scotto, 2004; 2007; 2008; Scotto & Pezzopane, 2007; Scotto, 2009], developed for the AIS-INGV ionosonde [Zuccheretti et al., 2003].

The DORS OIS has its own on-board scaling software within the second category of template matching techniques known as the DST-IIP (Ionogram Image Processing) algorithm [Heitmann & Gardiner-Garden, 2019]. A sample ionogram with key features and fitted traces overlaid is shown in **Figure 2.3**. Note that O/X separation is not required for the DST-IIP algorithm, and indeed the standard DORS configuration uses only a single vertically polarised, omni-directional receive antenna [Ayliffe et al., 2019]. A solution for the automatic separation of oblique O/X returns using dual-polarised dipole antennas with two orthogonal arms has been explored [Harris et al., 2017], although the required phase offsets are elevation dependent and, as such, results can be less robust than for the VIS geometry under complex multi-mode conditions (e.g. spread-F).

Beyond measuring an amplitude-only ionogram, modern ionosondes are often equipped with additional sensing capabilities to improve their characterisation of ionospheric variability. While there exist a number of commercially available vertical incidence sounders with support for angle-of-arrival (AoA) and/or Doppler (e.g. the Lowell Digisonde [Reinisch et al., 2009] and Dynasonde [Rietveld et al., 2008]), such systems are not as common for one-way oblique incidence as they require comparatively large arrays or vertical towers (e.g. Clarke & Tibble [1978]) to overcome the reduced angular
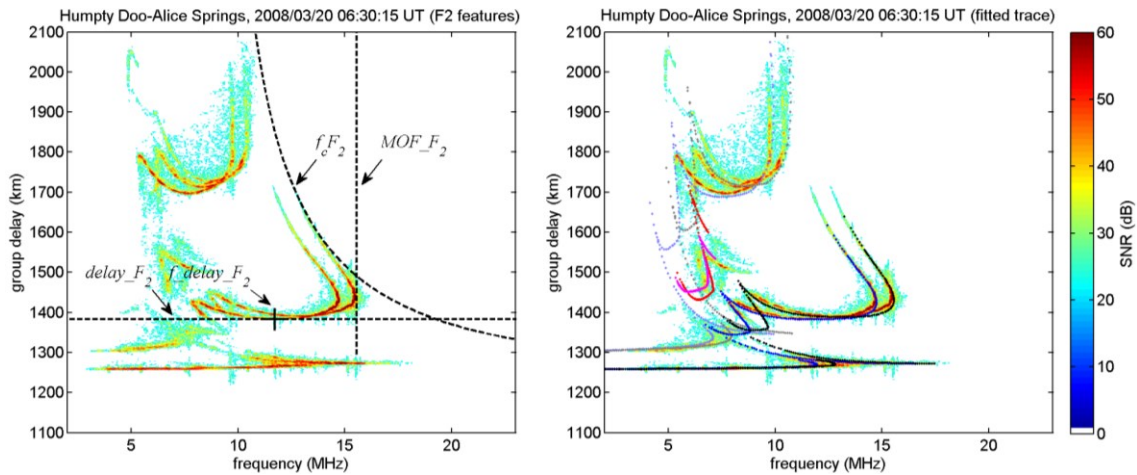
Figure 2.3. A daytime DORS ionogram from Humpty Doo to Alice Springs (1234 km) in northern Australia. The key 1-hop F2 features, extracted automatically, are overlaid in the left panel, while the fitted (synthetic) traces are overlaid in the right panel. The 1-hop trace in dark blue (O-mode) and black (X-mode) has been fitted to the extracted E, Es and F2 features, while the 2-hop traces (including Es/F combinations) in light blue, light grey, red and magenta are predicted only from the 1-hop fit, and may or may not match the image depending on horizontal gradients in the ionosphere.

resolution at low elevations. Multipath scatter due to re-radiation from local conductive structures (up to a few km away from the array) is also more detrimental for oblique AoA estimation [Ross et al., 1951], making site selection extra important. Many earlier oblique experiments (e.g. Ross et al. [1951]; Sweeney [1970]; Rice [1973]; Sherrill & Smith [1977]) focused on AoA observations over narrow frequency channels only, although several more recent studies have collected oblique AoA soundings over the full HF band (e.g. Wright & Kressman [1983]; Black et al. [1993]; Sherrill & Brown [1997]; Rogers et al. [2003]; Vertogradov et al. [2013]). The principles of modern direction finding [Gething, 1991; McNamara, 1991], including array design and digital wavefront analysis (AoA estimation), are of course pertinent to this oblique AoA sounding problem.

The ionospheric Doppler shift is an important observable for characterising horizontal drifts and the motion of TIDs. By convention, it is a measure of the time rate of change of phase path (itself an ambiguous quantity), and therefore requires one or more replicas of the sounder chirp or pulsed signal to produce a valid estimate. Observations can be

collected in the form of either wideband Doppler ionograms, typically with just one or a small number of Doppler estimates per frequency and delay cell (e.g. Wright & Pitteway [1982]; Reinisch [1986]; Parkinson et al. [1999]; Ayliffe et al. [2019]), or narrowband channel scattering functions (CSFs), capturing the full Doppler spectra but at a discrete set of frequencies only (e.g. Basler et al. [1988]; Wagner et al. [1988]; Warrington [1996]; Angling et al. [1998]; Harris & Frazer [2005]). As always, this represents a compromise between frequency resolvability, Doppler resolvability, and the total observation time. The choice ultimately depends on the required sounder cadence (with CSFs requiring a longer dwell time per frequency), spectral licensing agreements, and nature of the disturbances being studied. Instantaneously wideband, noise-like waveforms may offer the potential for concurrent CSFs at all frequencies [Barnes et al., 2010], although the interference implications on other HF users must be carefully considered. Time-interleaved measurements can alternatively be used to boost the Doppler resolution [Lynn, 2008], but for a swept waveform, this means breaking the time/frequency coherence of the signal.

2.3   Over-the-horizon radar in Australia

The Jindalee Operational Radar Network (JORN) is a network of three HF sky-wave over-the-horizon radars (OTHRs) operated by the Australian Defence Force, for monitoring aircraft and ships off the northern and western coasts of Australia. Each radar consists of two electronically steered phased arrays of vertically polarised antennas, for oblique transmission and reception at long ranges in a quasi-monostatic configuration (with transmit and receive sites separated by up to 100 km). Backscattered energy from targets of interest are resolved in apparent azimuth (steer angle), slant range (group delay), and Doppler, with the latter particularly essential for separating out the unwanted "clutter" caused by ground and sea backscatter. The principles of HF sky-wave radar are described in Headrick & Anderson [2008] and Fabrizio [2013], while the relevant ionospheric effects are addressed in Cervera et al. [in press]. More specific details about JORN itself can be found in Cameron [1995], Colegrove [2000], and Johnson et al. [2008]. A map of the JORN radar sites and coverage is shown in **Figure 2.4**.
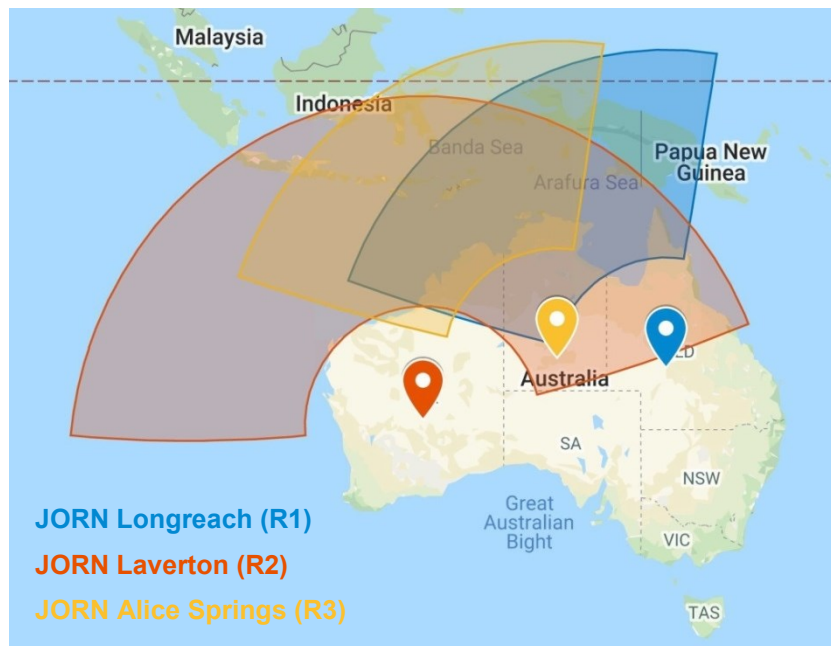
Figure 2.4. Locations and surveillance regions for each of the three JORN radars. Remote OIS and VIS assets provide ionospheric observations surrounding the 1-hop midpoints (notionally at half the range).

Because of diurnal changes in the ionospheric frequency support and varying channel characteristics, the JORN radars, like their international OTHR counterparts (including ROTHR [Headrick & Thomason, 1996] and NOSTRADAMUS [Bazin et al., 2006]), are frequency-agile systems by necessity. At any time, there exist both leading and trailing edges of backscatter propagation which dictate the minimum and maximum ranges that can be reached by each frequency [McNamara, 1991, ch. 5.4]; all shorter ranges fall in the "skip zone" and fail to be illuminated, while longer ranges at low elevation angles are highly attenuated. Among the available propagating modes, 1-hop E and Es are favoured for stable and often single-mode propagation out to 2000 km, most suitable for the detection of slow-moving ships at Doppler frequencies in and around the clutter, whereas the leading edge of 1-hop F2 offers increased range out to 3000 km and generally higher sub-clutter visibility (or clutter-to-noise ratio), most suitable for the detection of aircraft. The choice of frequency/mode does also depend on the size (radar cross section) of the target and clear channel availability.

The JORN radars are supported by a substantial number of remote OIS and VIS assets (currently forming 21 and 13 ionospheric sample points, respectively), which are used

to accurately observe and model radar propagation out to distances of 1000–3000 km [Earl & Ward, 1987; Wheadon et al., 1994]. Two-way backscatter sounders and mini-radars also complement the one-way OIS and VIS observations with measurements of the backscattered clutter power, leading/trailing edges of propagation, and ionospheric Doppler characteristics as a function of frequency and range. Together these instruments feed data into a real-time parameterised representation of the ionosphere, which serves two major roles: (1) to offer waveform parameter advice, including frequency selection, for illuminating the desired geographic region, and (2) to perform coordinate registration (CR) on target tracks, in conjunction with known reference points (such as cooperative targets, HF transponders/beacons, land-sea boundaries, and prominent terrain features). In short, coordinate registration is the mapping from radar space (steer angle and group delay) to ground coordinates (latitude and longitude); this is highly dependent on the operating frequency and propagation mode. The nature of ionospheric variability means that climatological models are inadequate to perform these key roles. In particular, unlike the HF communication problem, HF radar demands not only accurate estimates of critical frequencies, but also layer heights, in order to provide effective CR.

Although the JORN radar coverage spans geomagnetic latitudes of 5–40 °S, the ionospheric region that governs propagation only extends as far as the 1-hop midpoints, at roughly half the range. This sits on the southern shoulder of the equatorial anomaly (at ~20 °S) and is therefore mid-latitude in nature. Compared to equatorial and auroral/polar latitudes, the ionosphere at mid-latitudes is much more benign, largely avoiding the worst of the plasma instabilities and magnetospheric coupling which can lead to strong radio wave scattering from small-scale irregularities. Direct field-aligned backscatter is unlikely owing to the non-perpendicular alignment of JORN ray paths with respect to the geomagnetic field.

Commonly encountered ionospheric disturbances at mid-latitudes, such as TIDs and night-time spread-F, affect HF sky-wave radar observations in two ways:

1 By reducing the detectability of targets, due to ionospherically induced perturbations in backscattered power (focusing/defocusing effects) and Doppler. In complex multi-mode propagation environments, unresolved mode structure may lead to wave interference effects [Västberg & Lundborg, 1997; Fabrizio,

2000; Yau, 2003], causing rapid signal fading, distortion and potential azimuthal errors, while time and range dependent ionospheric Doppler perturbations also degrade the Doppler spectra of clutter in and around which targets may be detected. Changes in the AoA may also cause low-elevation rays to and from the target to become obscured by the Earth.

2   By reducing the CR accuracy, as these disturbances possess fine structure and short decorrelation scales which are often poorly captured by the sparse spatial and temporal sampling of conventional ionosonde networks. This can lead to errors in both the automatic association of target tracks on different propagation modes and CR corrections [Barnes et al., 2000; Nickisch et al., 2007], ultimately resulting in poor target localisation and the potential for duplicate (ghost) tracks.

Mitigation techniques for TIDs in OTHR data are non-trivial and generally require full spatial and temporal characterisation of the disturbance field [Jones & Reynolds, 1975; Nickisch et al., 2006; 2007]; AoA and Doppler soundings are especially relevant to such an inversion problem. Those TIDs with a horizontal propagation component transverse to the HF ray path tend to be more unfavourable, as anti-symmetric or tilted structure in the down-range (longitudinal) direction has very little impact on delay and azimuth measurements [Barnes et al., 2000; Bennett & Dyson, 2002]. This is due to the self-correcting geometry between the path segments to and from the ionosphere. Recent projects such as Net-TIDE [Reinisch et al., 2018] and HFGeo [Munton et al., 2019] have studied this problem in some detail, with data-assimilative models extended to incorporate AoA inputs in order to provide improved CR [Nickisch et al., 2016; Mitchell et al., 2017]. Signal processing techniques also have an important role to play in mitigating the fading and distortion effects (e.g. Anderson & Abramovich [1998]; Fabrizio [2000]), as well as novel radar architectures that offer directional selectivity on both transmit and receive to eliminate degraded propagation modes (e.g. Abramovich et al. [2013]; Frazer et al. [2014]).

Some of these challenges are currently being addressed as part of a mid-life upgrade of JORN which is now underway [Allison et al., 2019]. As well as replacing ageing equipment, this project will see the ionosonde network expanded to offer an even finer level of spatial sampling, to improve the modelling of medium- to large-scale disturbances. This builds on knowledge gained through DST Group trials such as the

Dense Integrated Network of Ionospheric Sounders (DINIS) [Gardiner-Garden et al., 2011], the Spatial Ionospheric Correlation Experiment (SpICE) [Harris et al., 2012], and most recently, the Elevation-sensitive Oblique Incidence Sounder Experiment (ELOISE) [Gardiner-Garden et al., 2019]. A subset of data from ELOISE forms the basis of the analysis in this thesis.

2.4   Observations and models of ionospheric variability

Ionospheric variability is broadly defined as behaviour that deviates from the long-term mean or "climatology" [Forbes et al., 2000; Rishbeth & Mendillo, 2001; Araujo-Pradere et al., 2005]. There are many sources of variability at mid-latitudes, and their effects on radio waves range from small-scale amplitude and phase scintillations, causing short-term signal fading, to large-scale perturbations in the propagation path, causing an increased likelihood of errors in direction finding and coordinate registration applications. During periods of severe degradation, when both signal and clutter are significantly spread in range and Doppler due to micro-multipath irregularities (typically associated with night-time spread-F), much of the HF band may be left unusable for voice/data communications and radar target detections. In essence, any irregularities on the scale of a wavelength or larger can have a potential impact on HF propagation and are therefore of great interest to observe and model.

It is important to distinguish between normal and abnormal variability. The former is what most data-assimilative models such as the JORN RTIM primarily seek to address, in a deterministic and often real-time capacity. It is contained in the usual day-to-day and hour-to-hour changes in the electron density distribution, driven by the solar flux and bulk plasma transport mechanisms. Longitudinal gradients associated with the solar (dawn/dusk) terminators and latitudinal gradients associated with the shoulder of the equatorial anomaly [Davies, 1990, ch. 5.1] are prominent examples at mid-latitudes. The latter (abnormal) component is more difficult to predict, often appearing as data outliers which test the robustness of models, and in some cases must rely on a statistical characterisation only. It can be triggered by isolated occurrences such as high-energy solar events (e.g. X-ray flares and coronal mass ejections), geomagnetic storms (through interactions with the solar wind and auroral particle precipitation), earthquakes and

tsunamis, and explosions. For localised sources, this can include propagating wave-like disturbances such as TIDs, which are the focus of the next section.

Having explained the operational importance of understanding ionospheric dynamics and disturbances, this section will review some of the ways in which spatial and temporal variability in the bottomside ionosphere is observed and modelled. Although networks of ground-based ionosondes are a well-established means of observing the ionosphere at low cost and with high sensitivity, various other instruments have also contributed to the vast knowledge on this subject, including many outside the HF band. Among these are powerful VHF incoherent scatter radars (e.g. Hunsucker [1991, ch. 5]; Balan et al. [2000]; Lei et al. [2005]; Chau et al. [2010]), GNSS ground-based receiver networks performing total electron content mapping/tomography (e.g. Mannucci et al. [1998]; Otsuka et al. [2002]; Jakowski et al. [2012]), GNSS radio occultation satellite missions (e.g. Schreiner et al. [1999]; Jakowski et al. [2002]; Yue et al. [2014]), in situ rocket and low Earth orbit sensors (Langmuir probes) (e.g. Brace [1998]; McNamara, Cooke, et al. [2007]), and all-sky airglow cameras (e.g. Makela [2006]; Shiokawa et al. [2009]).

Within the HF band, observations of ionospheric variability are commonly made using ground-based ionosondes (as described earlier), continuous-wave Doppler radars (e.g. Davies & Baker [1966]; Chum et al. [2012]; Crowley & Rodrigues [2012]), dedicated CSF beacon networks (e.g. Nickisch et al. [2016]; Hysell et al. [2016]), and backscatter sky-wave radars (e.g. Croft [1972]; Anderson & Lees [1988]; He et al. [2004]; Chisham et al. [2007]; Ishida et al. [2008]). Each measures some combination of signal amplitude, group delay, phase path (Doppler), AoA, and/or polarisation information, although only the ionosonde can provide unambiguous electron density profiles. Knowledge of AoA is crucial for the study of gradients and other horizontal structure. Routine gradients such as those associated with passage of the solar terminator, especially at dawn, can cause AoA deviations of up to 20° at vertical incidence [Paul, 1985] and several degrees at oblique incidence [Black et al., 1993].

There are a vast number of both physics-based and data-driven models of the ionosphere that seek to represent ionospheric propagation variability in terms of changes in the underlying electron density distribution. The data-driven models can be further split into climatological models, based on empirical functions or monthly-

median statistics, and data-assimilative models, which ingest real-time observations to support "nowcast" predictions; each class is suited to different applications. Notable examples of the former are the International Reference Ionosphere (IRI) [Bilitza, 2001; Bilitza & Reinisch, 2008; Bilitza et al., 2014; Bilitza et al., 2017], NeQuick [Radicella & Leitinger, 2001; Leitinger et al., 2005; Nava et al., 2008], Simplified Ionospheric Regional Model (SIRM) [Zolesi et al., 1993; 1996], and Advanced Stand Alone Prediction System (ASAPS) [McNamara, Decker, et al., 2007; Space Weather Services, 2020], while examples of the latter are the Global Assimilation of Ionospheric Measurements (GAIM) [Schunk et al., 2004; Scherliess et al., 2004; Scherliess et al., 2006], IRI Real-Time Assimilative Mapping (IRTAM) [Galkin et al., 2012; Bilitza et al., 2017], SIRM Updated (SIRMUP) [Zolesi et al., 2004; Tsagouri et al., 2005; Pezzopane et al., 2011], Electron Density Assimilative Model (EDAM) [Angling & Cannon, 2004; Angling & Khattatov, 2006; Angling & Jackson-Booth, 2011], GPS Ionospheric Inversion (GPSII) [Fridman et al., 2006; Fridman et al., 2012; Fridman et al., 2016], and JORN's Real-Time Ionospheric Model (RTIM) [Barnes et al., 2000; Gardiner-Garden, Heitmann, et al., 2008; Gardiner-Garden et al., 2011]. Nowcasting usually adopts a background ionosphere as a starting point for modifications, which may be either a monthly-median representation (e.g. the IRI) or a physics-based model (as in GAIM).

Evaluations of nowcast modelling performance against truth ionosonde data have consistently shown significant accuracy improvements over purely climatological models such as the IRI [McNamara et al., 2011; Pezzopane et al., 2011; McNamara et al., 2013; Elvidge et al., 2014], particularly in terms of the day-to-day variability caused by solar and geomagnetic activity (space weather) and hour-to-hour variability caused by neutral winds and TIDs. Physics-based ("first principles") models such as the Thermosphere-Ionosphere-Electrodynamics General Circulation Model (TIE-GCM) [Qian et al., 2014], Global Ionosphere Thermosphere Model (GITM) [Ridley et al., 2006], and SAMI3 [Huba et al., 2008] tend to be poorer again than the empirical approaches, both in terms of systematic biases and correlation with observations [Elvidge & Angling, 2014; Elvidge et al., 2017], making them more suited to basic research of atmospheric coupling than operational use.

Although each model adopts a slightly different parameterisation or sampling scheme, all the nowcast models listed above are primarily based on the assimilation of

measurements from ground-based ionosondes and/or GNSS Total Electron Content (TEC) receivers. GPSII is rather unique in its ability to also assimilate AoA observations. Data from other satellite sensors, backscatter leading edges, and networks of HF beacons and known reference sources may be combined through more general tomographic imaging and constrained optimisation techniques [Fridman & Nickisch, 2001; Bust & Mitchell, 2008; Hysell et al., 2016], although the ionosonde still remains one of the best sources of fine height information [Rogers et al., 2001]. Besides the global coverage of GAIM and IRTAM, most approaches are designed for regional ionospheric modelling; the JORN RTIM is one such model. The RTIM is constructed as a set of spatial maps, each of which represents a parameter from the JORN multi-segment quasi-parabolic (MQP) profile specification [Gardiner-Garden et al., 2018] on a regular latitude/longitude grid over Australia.

A core part of any ionospheric modelling effort is the ability to ray-trace through an electron density distribution, under the assumption of geometric optics. While various implementations exist (e.g. Coleman [1998]; Bennett et al. [2004]; Nickisch [2008]; Azzarone et al. [2012]; Cervera & Harris [2014]; Zawdie [2015, app. A]; Psiaki [2019]), all involve integrating the refractive index, as given for example by the Appleton-Hartree formula in equation (2.2), over some conceptual ray path, defined by a frequency and initial launch angles. Phase path P and group path P' (more strictly, path lengths) are distinguished from the geometric path S by the integrand. In the absence of collisions [Davies, 1990, ch. 1.5; Bennett et al., 2004],

$$P = \int \mu \cos\alpha \, ds, \tag{2.8}$$

$$P' = \int \mu' \cos\alpha \, ds, \tag{2.9}$$

and

$$S = \int ds, \tag{2.10}$$

where α is the angle between the wave vector and ray direction (zero for an isotropic media without the geomagnetic field), and $\mu' = \frac{\partial}{\partial f}(f\mu) \approx 1/\mu$ (neglecting collisions). Note that allowing for collisions leads to a complex refractive index and ray path, which makes calculations much more laborious [Budden & Terry, 1971]; instead, absorption effects can be modelled separately [Pederick & Cervera, 2014; Zawdie et al., 2017].

Some parametric profiles, such as the JORN MQP profile, support fast analytic ray tracing (ART) under certain approximations [Dyson & Bennett, 1988; Bennett et al., 1991], even allowing for spherical asymmetry [Norman & Cannon, 1999; Norman et al., 2012] and first-order adjustments for O/X splitting [Bennett et al., 1991; Chen et al., 1990; 1992]. However, with the increased performance of modern computing, numerical ray tracing (NRT) is now far more common, offering considerably more flexibility, without parametric constraints, and accuracy, particularly with regard to characterising horizontal gradients and geomagnetic effects. The JORN parameter advice and coordinate registration sub-systems currently run 2D NRT in operations, while 3D NRT is employed for research purposes using the PHaRLAP toolbox [Cervera & Harris, 2014; Cervera, 2019]. The latter is based on a form of the Hamiltonian (canonical) equations derived by Haselgrove & Haselgrove [1960] and Haselgrove [1963].

Synthetic ionogram traces, such as those in **Figure 2.1** and **Figure 2.3**, are an important tool for evaluating an ionospheric model, and can be generated by applying ART/NRT and a ray homing (interpolation) technique [Chen et al., 1992]. Equivalent geometry may be used as a starting point for defining a suitable fan of rays. Group delay $P' = cT_g$ (from sounder observations) and group path $P' = \int \mu' \cos\alpha \, ds$ (equation (2.9), from ray tracing) are taken to be one and the same by Fermat's principle, thus permitting a direct comparison with observations. Meanwhile, the ionospheric Doppler shift is related to the rate of change of phase path by $f_{id} = -\frac{f}{c}\frac{dP}{dt}$, and can be estimated from the ray data as a finite difference. The worth of any model can be measured in terms of its ability to synthetically reproduce the HF observables across space (i.e. different paths) and time.

Despite the limitations discussed earlier, spherical mirror models still have their place in studies of medium- to large-scale ionospheric disturbances, especially when qualitative understanding is more important than numerical exactness. The simplicity and elegance is unrivalled, and tilts of up to a few degrees can be incorporated fairly easily, for

example, by shifting the centre of the reflecting sphere so that it is no longer concentric with the Earth [Folkestad, 1968; Dyson & Bennett, 1992]. For rapid smaller-scale fluctuations associated with fading, scintillation and spread Doppler signatures, specular reflection is no longer a valid approximation, and the effects of scattering and wave interference must be considered; a simple but commonly-used technique is the thin diffracting phase screen [Ratcliffe, 1956; Basler et al., 1988; Nickisch, 1992]. Such effects can also be simulated by combining a statistical model for the irregularities with standard ray tracing [Coleman, 1996]. Owing to the vastly different scales involved, it is not unusual to focus on only one of the above variability components at a time, and in this work, both the experimental setup and models were designed first and foremost for characterising the medium- to large-scale component dominated by gradients and TIDs.

2.5 Travelling ionospheric disturbances

Travelling ionospheric disturbances are understood to be the ionospheric manifestation of atmospheric gravity waves (AGWs) in the thermosphere, where quasi-periodic perturbations in the neutral winds cause free electrons to be displaced along geomagnetic field lines (e.g. Hocke & Schlegel [1996]). Wave-like irregularities were first identified in HF ionospheric returns by Pierce & Mimno [1940], who ascribed them to regions of concave and convex curvature, and the pioneering spaced-receiver experiments of Munro [1948; 1950; 1958] provided clear evidence of their horizontal motion. It was the seminal work by Hines [1960] (also Pitteway & Hines [1963; 1965]) that established the theoretical relationship between TIDs and AGWs, and their propagation in an isothermal atmosphere. As one of the main contributors to ionospheric variability, and present almost continuously, there have been countless subsequent studies on the subject of TIDs, and these are summarised in reviews by Yeh & Liu [1974], Francis [1975], Hunsucker [1982], and Hocke & Schlegel [1996], and the compilation of annotated papers in Hines et al. [1974].

Part of the strong interest in AGWs/TIDs is their role in transferring energy and momentum between different atmospheric heights and latitudes [Laštovička, 2006], as the waves dissipate and break. While the term is used rather broadly to encapsulate waves generated from a wide variety of disparate mechanisms, among the common sources are meteorological processes (e.g. frontal systems, orographic wind flows, and

convective activity such as during severe thunderstorms) in the troposphere, and auroral dynamics (e.g. Lorentz forces and Joule heating associated with the auroral electrojet and energetic particle precipitation) in the high-latitude upper atmosphere. Isolated events such as earthquakes and explosions can also trigger AGWs/TIDs. As a general rule of thumb, medium-scale TIDs (MSTIDs, with ~15–60 min periods) tend to be associated with regional tropospheric sources within 1000–2000 km, while large-scale TIDs (LSTIDs, with ~30–180 min periods) typically propagate equatorwards from distant auroral sources at E region heights, even under geomagnetically quiet conditions [Hocke & Schlegel, 1996]. The range of horizontal wavelengths and periods that can propagate up to the F2 region is dictated by the AGW dispersion relation, dissipative effects, and background winds.

In the ionogram domain, TIDs often appear as distortions or kinks in the otherwise smooth trace, which ripple down from high to low heights as a result of the characteristic forward-tilted phase-front of the horizontally moving disturbance [Martyn, 1950; Munro & Heisler, 1956a]. Such features correspond to off-angle propagation geometries and are usually accompanied by an S-shaped perturbation in the ionospheric Doppler shift as a function of time [Davies & Baker, 1966; Georges, 1968; Chum et al., 2008]. Post-sunset TID amplitudes are typically larger, due to the increased height of the ionosphere and lower neutral atmospheric densities at these heights, and can result in "satellite" traces, distinct from the main trace altogether. Depending on the path length, it is not uncommon for reflection points to be horizontally displaced by tens or hundreds of kilometres, on the same scale as the TID wavelength, which can have serious repercussions on techniques that assume midpoint reflection [Zawdie et al., 2016].

To place the relevant research problems into context, it is useful to refer to the AGW/TID response function model in **Figure 2.5**. This expresses the response as the convolution of four filtering stages, each of which imparts its own characteristics on the power spectrum and directional pattern of the observed perturbations. From left to right, stage 1 covers the distribution of AGW sources and how they radiate away from the source locations, stage 2 addresses the influence of neutral wind filtering and dissipation as the waves propagate to greater heights, stage 3 governs the ion/neutral interactions that produce the TID within the ionosphere, and stage 4 draws in the propagation and instrumental effects. An overview of each is provided in the following paragraphs.

1. AGW sources and radiation patterns	2. Atmospheric wind filtering and dissipation	3. Ionospheric response (AGW → TID)	4. Observing technique and response
• Local tropospheric versus distant auroral sources • Electric field mapping to geomagnetic conjugate regions	• Horizontal wind blocking diagrams • AGW propagation models and dispersion relations • Wave dissipation	• Collisional model, incorporating geomagnetic field • Relate neutral wind velocities to electron density perturbations	• HF, VHF, microwave, optical, etc. • Path geometry modelled via ray tracing • Instrumental sensitivity

Figure 2.5. A simple AGW/TID response function model, adapted from Figure 8 of Hunsucker [1982]. Each of the four filtering stages expresses one aspect of the problem that must be modelled in order to fully understand how the observations relate back to the AGW sources. The dot-points underneath highlight just a few of the main factors and physical relationships that need to be considered.

The initial stages 1 and 2 rely on a detailed understanding of neutral atmospheric dynamics and, ideally, a method of gravity wave ray tracing [Waldock & Jones, 1987; Eckermann & Marks, 1997; Vadas & Fritts, 2006]. Knowledge of the source, and the degree to which it is localised or distributed, is critical to modelling the radiation pattern and long-distance propagation and growth of the wave. Mid-latitude MSTIDs, for example, can present quite different characteristics (including seasonal occurrence and peak azimuthal direction) depending on their source mechanism, with statistical studies (e.g. Shiokawa et al. [2003]; Kotake et al. [2007]; Ding et al. [2011]; Park et al. [2014]) often making the distinction between daytime, night-time, and terminator MSTIDs.

As waves ascend from the lower to the upper atmosphere, the stratosphere and thermosphere filter those AGWs that are travelling slower than the prevailing winds [Hines & Reddy, 1967; Cowling et al., 1971; Waldock & Jones, 1984], leaving only the faster and larger-scale waves at F2 heights. While AGW blocking diagrams based on this critical layer theory can be insightful [Taylor et al., 1993; Medeiros et al., 2003], they are only as good as the representation of the horizontal winds (e.g. using NRL's Horizontal Wind Model [Drob et al., 2015] or NCAR's TIE-GCM [Qian et al., 2014]), which may be rather inaccurate in the thermosphere [Dandenault & Richards, 2015]. Another element to consider here is the mapping of polarisation electric fields along

geomagnetic field lines to create so-called "electrified" TIDs at night [Saito et al., 1998; Kelley, 2011; Narayanan et al., 2018]; evidence for this can be found in airglow images and GNSS TEC values mirrored at magnetic conjugate points (e.g. Otsuka et al. [2004]; Shiokawa et al. [2005]; Valladares & Sheehan [2016]).

In addition to the primary sources in the lower atmosphere, it has been proposed that under certain conditions, a strong AGW may also break and deposit its energy and momentum into the background flow around the mesopause, via a cascade to smaller-scale waves and turbulence. This causes local body forces which excite larger-period secondary waves whose amplitudes grow rapidly with height [Vadas & Fritts, 2001; Vadas et al., 2003]. With few direct observations, it is fair to say that these aspects of the response function model remain largely based on theory.

Stage 3 in **Figure 2.5** covers the key relationship between wind perturbations in the thermosphere (the AGW) and electron density perturbations in the ionosphere (the TID). Through ion-neutral collisions and ambipolar diffusion, a component of the neutral wind motion is imparted to the ionospheric plasma along the direction of the geomagnetic field [Hines, 1960]. At F2 heights, neutral density changes are sufficiently small and ion lifetimes sufficiently long that fluctuations in the rates of photoionisation and recombination, as well as the solar flux reaching a given point, are mostly negligible compared to the dominant transport process. As such, the electron density perturbations may be calculated from the divergence alone [Hooke, 1968]. The result is a TID that has many of the same properties as the underlying AGW, including period, horizontal wavelength and velocity, and can therefore function as a passive tracer for the source wave [Francis, 1975]. Fractional TID amplitudes can be much larger than the associated AGW [Hooke, 1968; Dyson et al., 1970; Kirchengast, 1996], with electron density perturbations of 10–20% not unusual in the F2 region (particularly at night) [Thome, 1964; Hooke, 1968], which makes them readily detectable by ground-based sensors. In addition to these larger disturbances, a smaller-amplitude TID background of ~1–3% perturbations is observed almost continuously by more sensitive instruments such as incoherent scatter radars at mid-latitudes [Oliver et al., 1997; Djuth et al., 1997; Djuth et al., 2004]. It is important to note that the AGW/TID relationship is not strictly one way, with ion drag also acting, albeit weakly, to damp the AGW [Kirchengast et al., 1996].

A number of physics-based TID models (e.g. Hooke [1968]; Clark [1970]; Clark et al. [1971]; Davis [1973]; Francis [1974]; Bottone [1993]; Kirchengast [1996]) have explored the AGW/TID relationship in some detail, starting with a dispersion formula and, from this, deriving equations of motion for the plasma. In most cases a small-amplitude, linearised form is assumed. The lossless model treatment developed by Hooke [1968] is a leading example, and is briefly summarised in **Appendix E**, with a focus on its application to TID parameterisation and synthesis. Along with the above models, based on electron density perturbations driven by the neutral wind as a whole (i.e. bulk motion), several authors have also reported numerical results using coupled atmospheric-ionospheric models to resolve interactions with the individual plasma and neutral species (e.g. Millward et al. [1993]; Duly et al. [2014]; Miyoshi et al. [2017]).

A fundamental property of the ionospheric response to AGWs, present in both theory and measurements, is its directional anisotropy [Hunsucker, 1982; Hocke & Schlegel, 1996]. Classical collision theory predicts a greater TID amplitude response for medium- to large-scale waves propagating equatorward in the mid-latitude F2 region, with the ambient vertical electron density gradient, geomagnetic inclination (dip angle), and azimuthal dependence of frictional ion drag being among the most influential factors [Hines, 1968; Liu & Yeh, 1969; Hooke, 1970]. Diurnal and seasonal variations occur due to the dependence on the profile shape, with a steep vertical gradient further enhancing the equatorward response. Wave-induced variations in the solar radiation flux, which depends on the alignment between the AGW phase-fronts and solar incidence angles, and anisotropic wind filtering of AGWs (from stage 2) may also play a role.

Night-time observations in particular indicate a preferential westward and equatorward propagation direction for mid-latitude F-region structure, including the electrified class of MSTIDs [Shiokawa et al., 2003; Shiokawa et al., 2009; Ding et al., 2011]. These disturbances are found to coexist over a very wide latitude span (at least 30°), with wavefronts extending out to thousands of kilometres [Tsugawa et al., 2007; Ogawa et al., 2009]. The night-time case is believed to be caused by the combination of the Perkins instability and polarisation electric fields [Kelley & Makela, 2001], in contrast to daytime MSTIDs which conform more to classical collisional theory [Kelley & Miller, 1997].

As one of the main research interests of the ELOISE experiment, stage 4 in **Figure 2.5** focuses on how the observing instrument responds to the ionospheric perturbations, as well as the inverse problem of extracting TID characteristics from observed signatures. Whereas measurements like GNSS TEC and airglow are height-integrated quantities, HF ionospheric soundings at multiple frequencies offer the benefit of height-sensitive information about the ionosphere. Off-angle returns are a leading indicator of tilted structure caused by TIDs, and this is where AoA estimation plays a key role, along with analysis tools like HF ray tracing. The measurement technique and its geometry also sets sensitivity limits on the amplitudes, spatio-temporal scales, and propagation directions of TIDs that may be detected, and it is not a trivial task to remove these effects when comparing different instruments, or even HF oblique paths of different lengths/orientations.

The signatures of TIDs in HF data are varied but generally include the following three characteristics: (1) quasi-periodicities that are common across multiple observables (such as power, group delay, Doppler and AoA) and fall within the physical spectrum of AGWs capable of reaching ionospheric heights; (2) a descending phase progression with height (frequency), caused by the forward-tilted phase-front; and (3) spatial lags between ionospheric sample points that point to a consistent horizontal propagation direction. It is worth noting that a quasi-periodic signature alone does not necessarily imply a TID. Depending on the scale of the disturbance, a kink in the ionogram trace is not always present, and even when it is, the TID can affect the observables over a much broader part of the trace [Lynn, 2008].

The HF response can be synthesised to first-order by considering specular reflection from a simple corrugated mirror surface (e.g. Lyon [1979]), or by ray tracing through a model ionosphere consisting of a horizontally-moving electron density peak or trough (e.g. Lobb & Titheridge [1977]). A more sophisticated approach is to couple an AGW-seeded perturbation model (from stage 3, e.g. Hooke [1968]) with a ray tracing engine (e.g. Cervera & Harris [2014]). Although in theory a single TID can be characterised by just a few parameters (i.e. amplitude, period, wave vector, and phase), in most real world examples, TIDs appear as a superposition of many components, of different origins, and isolated waves are unusual [Jacobson & Carlos, 1991; Bowman, 1995; Afraimovich et al., 2003].

Automatic inversion of TID parameters is considerably more difficult and usually relies on a number of simplifying approximations, such as mirror reflection and a small-amplitude limit, to make the problem tractable. A review of such techniques can be found in Hocke & Schlegel [1996]. For conventional ionosondes without AoA, the parameter inversion typically involves some form of multi-site true height analysis, which neglects horizontal motion in the reflection points (e.g. Klostermeyer [1969]; Tedd et al. [1984]; Bowman [1992]). With only a single AoA sounder path, partial parameterisation is possible using purely geometrical equations [Bramley, 1953; Clarke, 1972, ch. 7], but spectral analysis is usually called upon in the more general multi-component case. The Frequency and Angular Sounding (FAS) technique [Beley et al., 1995; Galushko et al., 2003; Paznukhov et al., 2012] has cemented its place as one of the leading algorithms for both oblique and vertical incidence AoA sounders, with application to real-time TID monitoring in the European Net-TIDE project [Reinisch et al., 2018], while the VIPIR/Dynasonde approach [Negrea et al., 2016] appears to offer excellent results at vertical incidence, using gradient (tilt) estimates from the NeXtYZ inversion technique [Zabotin et al., 2006]. Greater model sophistication again is needed to relate the TID observations back to the source AGW parameters (i.e. reversing both stages 3 and 4 via a physics-based model of AGW/TID interactions), with an attempt by Kirchengast et al. [1995] relying on parameter iteration and extra information from incoherent scatter radar, including ion/electron temperatures, to overcome the complexity of this problem.

There is still some debate about the relationship between AGWs/TIDs and spread-F, a phenomenon that is pervasive in the post-midnight ionosphere at mid-latitude sites, particularly around sunspot minimum and more so near the winter/June solstice(s) [Bowman, 1964; Hajkowicz, 2007]. Although the term "spread-F" originated as a description for diffuse F-region echoes in night-time ionograms, it is now used much more broadly to describe enhanced scattering and scintillation of radio waves in general, and the underlying irregularities themselves, which confuses things considerably.

It has long been established that spread-F in ionograms often coincides with severe radio scintillation [Briggs, 1964; Booker, 1979]; in equatorial regions this is widely accepted to be the result of small-scale plasma bubbles driven by the generalised Rayleigh-Taylor instability [Sultan, 1996; Kelley, 2009, ch. 4.2], whereas at mid-latitudes the mechanism for scintillation is postulated to be turbulent upwelling due to

AGW seeding of the Perkins instability [Perkins, 1973; Kelley & Fukao, 1991]. For HF soundings, off-angle total reflection from irregular (wave-like) electron density surfaces appears to be the primary cause [King, 1970; Bowman, 1990; Wright et al., 1996; Earle et al., 2010], with smaller-scale field-aligned irregularities having only a more minor role [Bowman & Hajkowicz, 1991]. Satellite traces are frequently observed to be a precursor to spread-F [McNicol et al., 1956; Herman, 1966; Bowman; 1981; Bowman & Monro, 1988; Lynn et al., 2013; Paul et al., 2019], adding weight to the argument that the two effects are related. Furthermore, through the use of ray tracing, TID-seeded models of plasma bubbles have been shown to reproduce spread-F-like returns in simulated ionograms [Jiang et al., 2020].

3   Experiment design, testing and calibration

3.1   Overview of the ELOISE angle-of-arrival experiment

The Elevation-sensitive Oblique Incidence Sounder Experiment (ELOISE) was a campaign aimed at observing the HF propagation effects of small- to medium-scale ionospheric variability in the mid-latitude Australian region (geomagnetic latitudes from 30–40 °S), and assessing the performance of data-assimilative ionospheric models such as the Jindalee Operational Radar Network (JORN) Real-Time Ionospheric Model (RTIM) in representing this variability. ELOISE augmented the existing JORN network of ground-based ionosondes with additional sites and measurement capabilities and included periods of wide-area target tracking and backscatter observations on the JORN Laverton sky-wave over-the-horizon radar. The experiment as a whole, which extends beyond the scope of this thesis, was carried out by large team of scientists and engineers from the Defence Science and Technology (DST) Group, in conjunction with the Royal Australian Air Force (RAAF), Capability Acquisition and Sustainment Group (CASG), BAE Systems Australia, and Lockheed Martin Australia.

ELOISE builds on the many earlier spaced ionosonde experiments, such as those conducted by Munro [1950], Clarke [1972], and Harris et al. [2012] (SpICE) in Australia, and more recently by Munton et al. [2019] (HFGeo), Reinisch et al. [2018] (Net-TIDE), and Caton et al. [2017] (AFCAP) internationally. However, it is believed to be unique in both its size (i.e. the total number of ionospheric sample points) and its joint use of large- and small-scale measurements, in space and time, across a variety of different high-fidelity instruments. By focusing on the use of oblique incidence sounders, equipped with multi-channel direct-digital receivers, ELOISE was able to sample the ionosphere very finely without an overwhelmingly large number of sites, as each pair of transmitting and receiving sites can be observed simultaneously. Such a network would not have been feasible with vertical incidence sounders alone.

The inclusion of angle-of-arrival (AoA) estimation on many of these oblique paths further enhanced the ability to characterise spatial structures such as large-scale electron density gradients and travelling ionospheric disturbances (TIDs), while Doppler spectra generated by the channel scattering function (CSF) mode complemented this with a measure of the corresponding dynamics. Since a significant proportion of ionospheric variability occurs on a timescale smaller than the sounder revisit interval, the CSF observations were crucial for gaining insights into the full spectrum of ionospheric variability. Although TIDs are understood to be a major contributor to mid-latitude ionospheric variability, it is important to note that the scope of ELOISE was broader than a study of TIDs alone.

A full description of the aims and data types of the ELOISE experiment can be found in Gardiner-Garden et al. [2019] (also reproduced in **Appendix F.5**), but in summary it consisted of the following components:

- Two twin-arm receiving arrays near Laverton, Western Australia, and Coondambo, South Australia, collecting angle-of-arrival (AoA) ionograms and channel scattering function (CSF) delay-Doppler maps on oblique paths of between 900 and 2700 km [Heitmann et al., 2018]. Eight cooperative down-range transmitters were controlled by these systems, each radiating from a single omni-directional antenna (as depicted in Figure 2 of Ayliffe et al. [2019]). The locations of the arrays and the eleven AoA paths observed are shown in **Figure 3.1**. The great circle path lengths and bearings are specified in **Table 3.1**. **Table 3.2** summarises the two standard observation modes (ionogram and CSF) that were scheduled on these ionosonde arrays.

- A dense network of conventional oblique incidence sounder (OIS) and vertical incidence sounder (VIS) systems, providing a total of 51 and 12 1-hop ionospheric sample points, respectively, over central and northern Australia.

- A fine-scale quasi-vertical incidence sounder (QVIS) network in the region of Alice Springs, Northern Territory, providing an additional five ionospheric sample points within ~60 km of the existing VIS sample point at Alice Springs Airport.

- Nine 4–8 hr collection periods on the JORN Radar 2 East (R2E) sky-wave backscatter radar at Laverton, including track data from a network of eight transponders (i.e. down-range repeaters) and commercial air targets of opportunity,

along with spatial maps of backscattered surface clutter over the full 90° arc of coverage (shown in **Figure 3.1**).

- An all-sky airglow camera (TRACE, the Thermospheric Radar Airglow Correlation Experiment) near Alice Springs [Unewisse et al., 2016], with 557.7 nm (green) and 630 nm (red) filters to measure atomic oxygen emissions at primarily E (~96 km) and F2 (~250 km) heights, respectively [Unewisse & Cool, 2016].

The research presented in this thesis focuses on the analysis of observations from the two ELOISE AoA arrays, described in the first dot-point. In carrying out this research, the author was responsible for much of the pre-trial planning and software development that pertained to these AoA systems, including the array design, testing and calibration, CSF scheduling algorithm, and on-board signal processing. Additional hardware and software support, including during the array deployment, was provided by staff at DST Group.
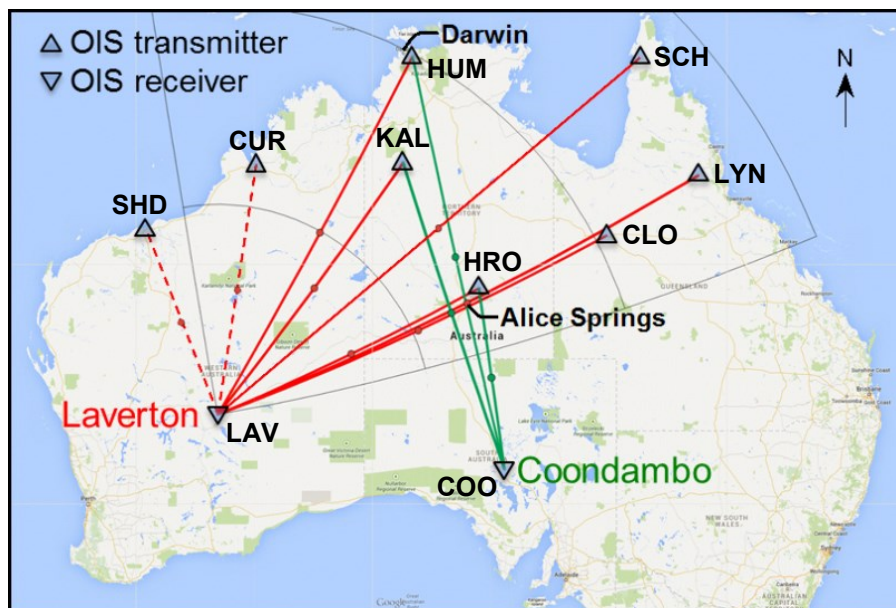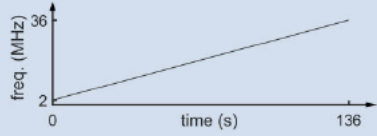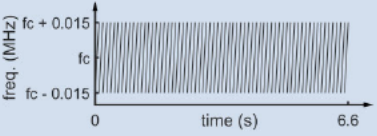
Figure 3.1. Locations of the two receiving arrays (Laverton and Coondambo) and eleven oblique AoA paths during the ELOISE campaign (as listed in **Table 3.1**). Doppler (CSF) observations were made on all but two of these paths (dashed lines). The black lines define the eastern arc of coverage for the co-located JORN R2E sky-wave radar at Laverton, showing down-range distances of 1000, 2000 and 3000 km. Local solar times at the path midpoints were between 8 and 9 hours ahead of Universal Time (UT).

Table 3.1. List of great circle path lengths and bearings for oblique reception on the ELOISE AoA arrays at Laverton (left column) and Coondambo (right column). All path bearings are azimuth angles measured clockwise with respect to True North (°T), a convention that is used throughout this thesis.

	Laverton (LAV) -28.331 °N, 122.004 °E	Coondambo (COO) -31.065 °N, 135.872 °E
South Hedland (SHD) -20.408 °N, 118.578 °E	944 km / -22 °T	–
Curtin (CUR) -17.610 °N, 123.823 °E	1202 km / 9 °T	–
Kalkarindji (KAL) -17.443 °N, 130.828 °E	1507 km / 39 °T	1593 km / -20 °T
Humpty Doo (HUM) -12.610 °N, 131.274 °E	1989 km / 31 °T	2098 km / -14 °T
Scherger (SCH) -12.658 °N, 142.087 °E	2712 km / 54 °T	–
Lynd River (LYN) -18.025 °N, 144.870 °E	2598 km / 69 °T	–
Cloncurry (CLO) -20.671 °N, 140.506 °E	2055 km / 70 °T	–
Harts Range (HRO) -22.969 °N, 134.449 °E	1383 km / 67 °T	908 km / -9 °T

Compared to the other conventional OIS, VIS and QVIS systems, the AoA arrays contribute valuable extra information about the spatial structure and dynamics of ionospheric variability, where characterised by off-angle reflections and Doppler shifts. With a range of different path lengths and orientations to the geomagnetic field, the AoA paths also have the potential to offer insights into the aspect angle dependence of disturbances that affect HF propagation (i.e. any anisotropic properties). For instance, the oblique path geometries in **Figure 3.1** include two paths (Lynd River to Laverton and Kalkarindji to Coondambo) that run roughly perpendicular to each other and have 1-hop midpoints separated by only 100 km, as well as two path triplets (Harts Range/ Cloncurry/Lynd River to Laverton and Harts Range/Kalkarindji/Humpty Doo to Coondambo) that are roughly collinear with spaced down-range midpoints.

Table 3.2. Key waveform parameters and processing characteristics for the ionogram and channel scattering function (CSF) modes of operation, which were interleaved on the ELOISE AoA array. Both observation modes were based on the use of linear frequency-modulated continuous-wave (LFMCW) or "chirp" waveforms from a common set of down-range transmitters.

	Ionogram mode	CSF mode
Waveform type	wideband single-sweep	narrowband repetitive
Waveform sketch		
Frequency band	2–36 MHz[*]	agile channel selection (from 2–36 MHz)
Sweep rate	250 kHz/s[*]	300 kHz/s[†]
Waveform repetition frequency (WRF)	–	10 Hz
Observation time	136 s[*]	6.6 s
Revisit interval	225 s[*]	225 s (for each of 3 channels)
Coherent integration time (CIT)	0.8 s[*]	6.4 s
Processing bandwidth	200 kHz (tapered)	30 kHz (tapered)[†]
Frequency step size	40 kHz	–
Group delay resolution	~1.5 km	~10 km[†]
Doppler resolution	–	~0.16 Hz
Image dimensions	801 (delay) × 848[*] (frequency)	121[†] (delay) × 64 (Doppler)
Key on-board processing steps	▪ RFI rejection and interpolation ▪ overlapped spectral processing ▪ parameterised profile fitting ▪ 2D AoA estimation	▪ background noise monitoring ▪ clear channel evaluation (CCE) ▪ delay-Doppler processing ▪ 2D AoA estimation

[*] Two of the transmitter sites, South Hedland and Curtin, were not upgraded to support CSF operations, and instead collected solely ionograms over a band of 2–45 MHz, with a 125 kHz/s sweep rate, 344 s observation time, and 450 s revisit interval. On these paths, the CIT and number of frequency cells also increased accordingly to 1.6 s and 1073, respectively.

[†] The CSF bandwidth was automatically and routinely reduced to 15 or 7.5 kHz when no suitable clear channel was available at 30 kHz. In such a case, the effective CSF sweep rate, group delay resolution, and number of delay cells would also be lowered in proportion to the bandwidth.

While the ELOISE campaign as a whole notionally ran from July to October 2015, the full set of AoA data were collected for only a subset of this, as illustrated in the availability plots of **Figure 3.2**. The top and bottom panels show ionogram and CSF availability, respectively. Missing data can occur for a number of different reasons, including power and data communication outages, transmitter and receiver hardware failures, system tests and upgrades, and (for CSF only) lack of propagation/clear channels, all of which were logged. The Coondambo system in particular also required on-site supervision, owing to its reliance on a diesel generator, so was operated on a shift basis (typically 8 hours per day) at times when DST Group staff were available.

The eleven oblique AoA paths were observed simultaneously with a revisit interval of 3.75 min in most cases. This cadence was chosen to ensure TID responses, with periods as low as 10–15 min, were adequately sampled. Scheduling was interleaved between wideband sounder sweeps (for ionograms) and narrowband repetitive waveforms (for channel scattering functions) on nine of these paths, such that during each epoch, one ionogram and three CSF dwells (on different frequencies) were recorded. The use of multiple CSF frequencies ensured that ionospheric Doppler was sampled at a variety of different heights in the F2 region. A further two paths into Laverton (from South Hedland and Curtin) were receiving AoA observations, without a CSF capability, and with a longer 7.5 min revisit interval; as such, the ionograms from these paths are reported as having 50% availability in **Figure 3.2**.

Following pre-trial maintenance and calibration activities, good-quality ionogram and CSF observations were made on both ELOISE AoA arrays from 5–17 September 2015; this will be the main focus period of subsequent chapters. Geomagnetic storm conditions were reported on three of these days: 7, 9 and 11 September (days 250, 252 and 254), the effects of which are briefly described in Gardiner-Garden et al. [2019]. The enlarged availability plot in **Figure 3.3** better highlights the data collected within this core period. The Coondambo array was disassembled shortly after this, whereas the Laverton array continued to collect data until suffering catastrophic lightning damage on 17 January 2016, albeit without access to CSF transmissions from 21 September 2015. The front-end receiver design has since been modified to provide more protection from lightning, although the Laverton system is yet to be restored.
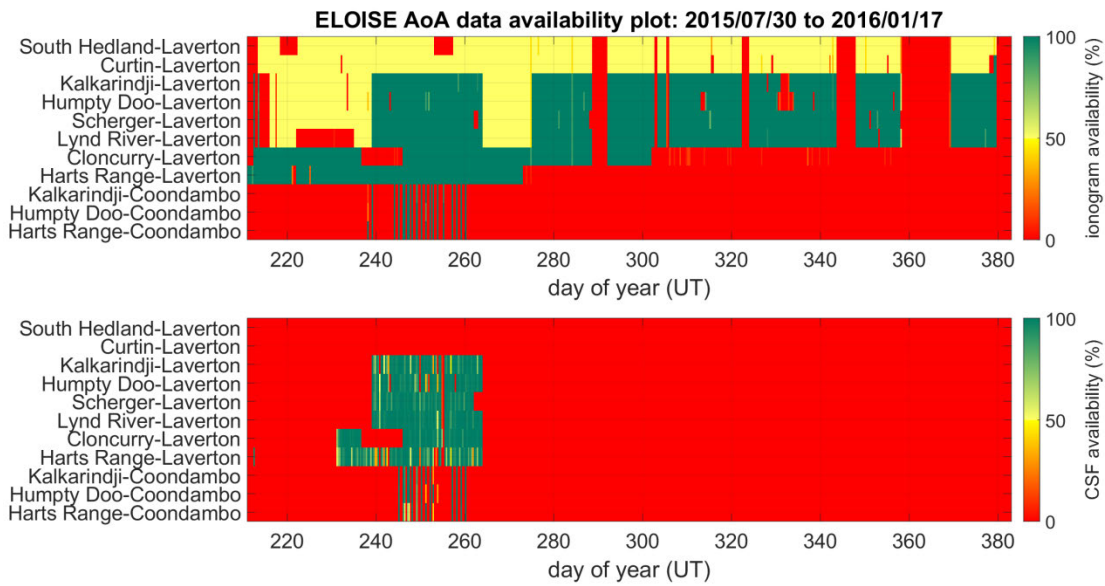
Figure 3.2. ELOISE AoA data availability, for both ionogram (top panel) and CSF (bottom panel) data types over the paths listed in **Table 3.1**. Full (100%) availability in a given hourly interval assumes one ionogram and three CSF dwells on each path every 3.75 min. Some sounders ran with a longer (7.5 min) revisit interval for certain periods, and as such, these appear as yellow (50%) in the top panel. Day-of-year is measured in UT days beginning on 1 January 2015 (day 1).
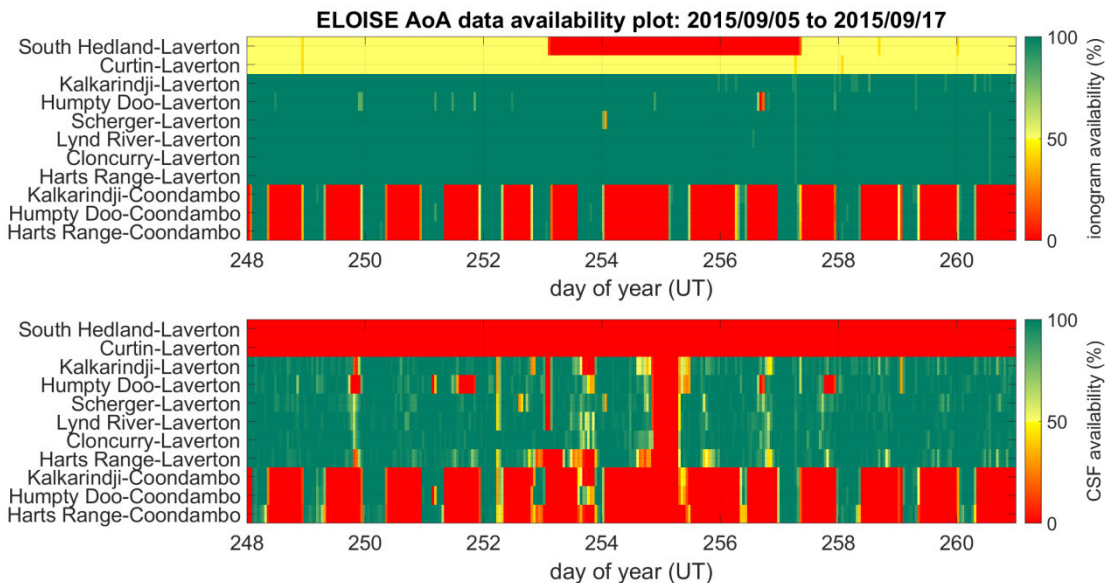
Figure 3.3. ELOISE AoA data availability, as in **Figure 3.2**, but enlarged to show just the core period from 5 to 17 September 2015 (days 248 to 260).

42

The remainder of this chapter presents an outline of the ELOISE AoA array design, deployment, testing and calibration activities, all of which occurred as a precursor to the experiment itself.

3.2 Array design

ELOISE presented the opportunity to develop a modern direction-finding ionosonde, customised for this particular experiment. Leveraging DST Group's considerable experience in ionospheric sounders (e.g. Harris et al. [2016]; Ayliffe et al. [2019]) and phased-array HF radars (e.g. Fabrizio [2013]; Frazer et al. [2014]), the ELOISE AoA system was a new design with a new set of on-board processing built upon this previous work. In contrast to older-generation direction-finding ionosondes (e.g. Clarke [1971]; Paul et al. [1974]; Wright & Pitteway [1979]; Baulch et al. [1984]; Poole & Evans [1985]; Black et al. [1995]; Vertogradov et al. [2013]), the ELOISE AoA array featured a multi-channel direct-digital receiver per element, for high-fidelity observations on multiple oblique paths simultaneously. Also, unlike many previous systems, it was designed for mid- to long-range oblique reception, primarily over a ~40° arc of azimuths with elevations from 0–20° (excluding the two paths without CSF support).

Two-dimensional phased antenna arrays are a common means for making directional measurements in radar, radio direction-finding, and radio astronomy applications. These arrays may be filled or sparse, with regular or irregular antenna spacing. Such design choices are a trade-off between resolution and gain requirements, subject to constraints on the physical array dimensions and number of antennas/receivers. For practical reasons, a horizontal planar array was chosen for ELOISE, although this provides a reduced effective aperture at low elevation angles. Vertically suspended HF arrays, including so-called curtain arrays, while possible, are costly and come with an increased potential for aeolian phase noise due to wind-induced antenna vibrations.

The ELOISE AoA array was constructed as a cross-type layout, consisting of two orthogonal uniform linear arms of ten vertically polarised broadband monopoles sharing a common antenna element (i.e. 19 elements in total). Unlike the conventional Mills Cross [Mills & Little, 1953], the common element was not at the geometric centre of the array; the asymmetry was a compromise to conform to the available cleared land and

43

accommodate another adjacent HF system at Laverton. Uniform spacing was chosen over a minimum redundancy solution to enhance sensitivity and increase robustness against wavefront distortions and degraded or poorly calibrated elements. A wire-mesh ground screen surrounded the entire installation to help stabilise the antenna impedances and improve low-elevation performance. The ground screen extended to a minimum of 10 m clearance on all sides of the array, including 30 m directly in front of the array (i.e. towards the paths of interest).

Schematics of the array layouts at Laverton and Coondambo are shown in **Figure 3.4** and **Figure 3.5**, respectively, with the ground mat dimensions indicated in **Figure 3.6**. The asterisks represent each of the 6.5 m aluminium monopoles, aligned to form one arm of length 180 m, with 20 m inter-element spacing, and another of length 90 m, with 10 m spacing. Boresight is defined to be the axis of the long arm. The arrays were deployed in May 2015 (Laverton) and August 2015 (Coondambo) by qualified technical staff from DST Group, allowing for a period of significant software development and testing to be carried out by the author ahead of the ELOISE experiment itself.
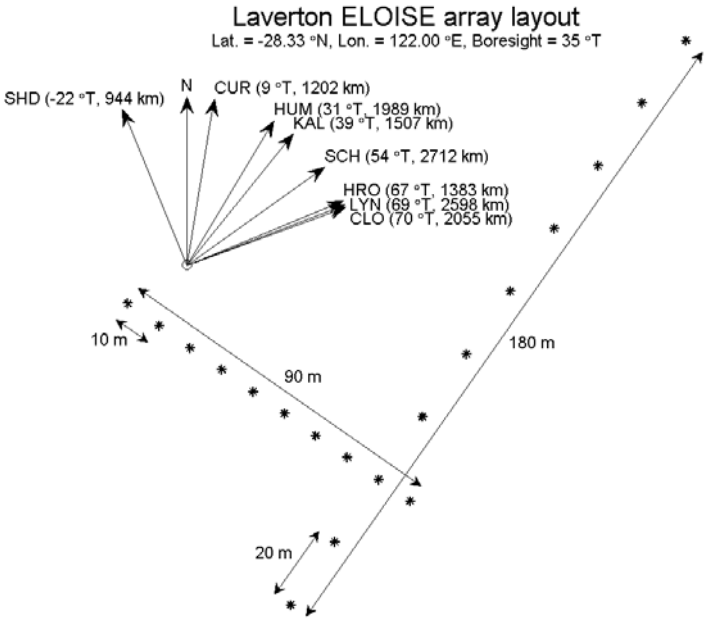
Figure 3.4. Overhead layout of the Laverton ELOISE array. The asterisks indicate the locations of the 19 monopole antenna elements, arranged in two uniform linear arms. Great circle bearings for all observed AoA paths in **Table 3.1** are shown by the inset compass arrows. Element numbering is from 1–10 on the longer ("endfire") arm, and [11, 8, 12, …, 19] on the shorter ("broadside") arm, where 8 is the common element.
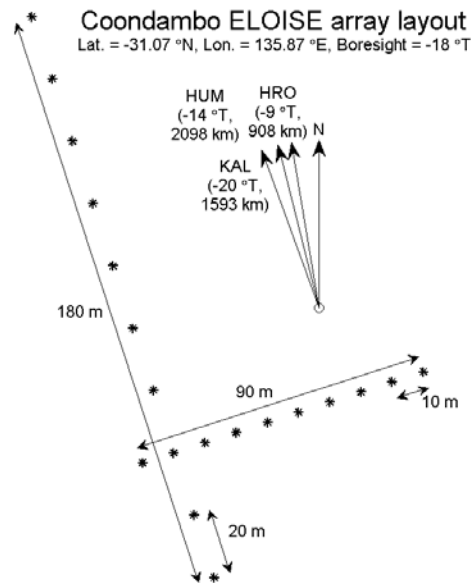
Figure 3.5. As above, but for the Coondambo ELOISE array.

A photo looking down the end of the short arm at Laverton is presented in **Figure 3.7**, while an aerial view clearly showing the two orthogonal arms is presented in **Figure 3.8**. On the paths of interest, the long and short arms notionally observe signals close to end-fire and broad-side, respectively, and provide the ability to estimate both elevation and bearing (azimuth) from both arms combined. Note that the Coondambo array is simply a mirrored and reoriented version of the Laverton array. To simplify calibration, the intent was to use a set of electrically matched feed cables at each site, although due to deployment issues there ended up being two sets of matched cables at Laverton (differing by about 4.5 m or 15 ns in electrical length).

Although such a sparse 2D array is not well suited to imaging applications, with conventional phased array beamforming offering poor side-lobe performance in general, its purpose in this case was simply to operate as a digital interferometer; that is, adopting a robust phase-fitting technique to estimate (but not fully resolve) AoA for the discrete one-way propagation in each processed frequency, delay and Doppler bin. Similar techniques (often described as Fourier or Doppler interferometry) have had a long history of use in radio direction-finding and imaging problems (e.g. Burtnyk & Wolfe [1967]; Adams et al. [1986]; Meek & Manson [1987]; Kudeki [1988]; Parkinson et al. [1997]; Holdsworth & Reid [2004]). It is worth noting that no polarisation
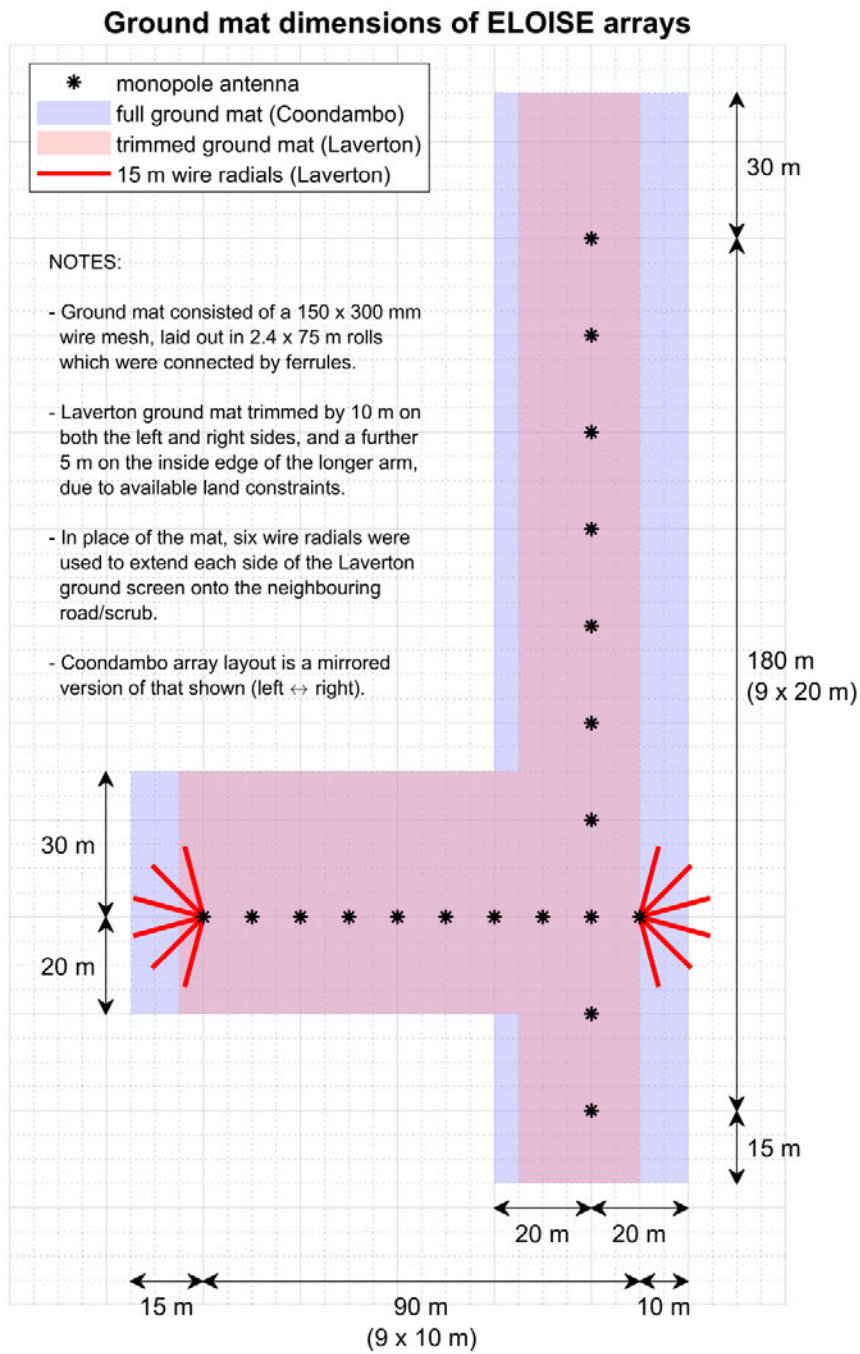
Figure 3.6. Ground mat dimensions of the ELOISE arrays, with the Laverton site modified (as shown in red) to conform to available land constraints. To improve the low-elevation performance, the wire mesh was extended to a distance of 30 m, or one wavelength at a typical operating frequency of 10 MHz, at all points in front of the array (i.e. towards the oblique paths of interest).

46
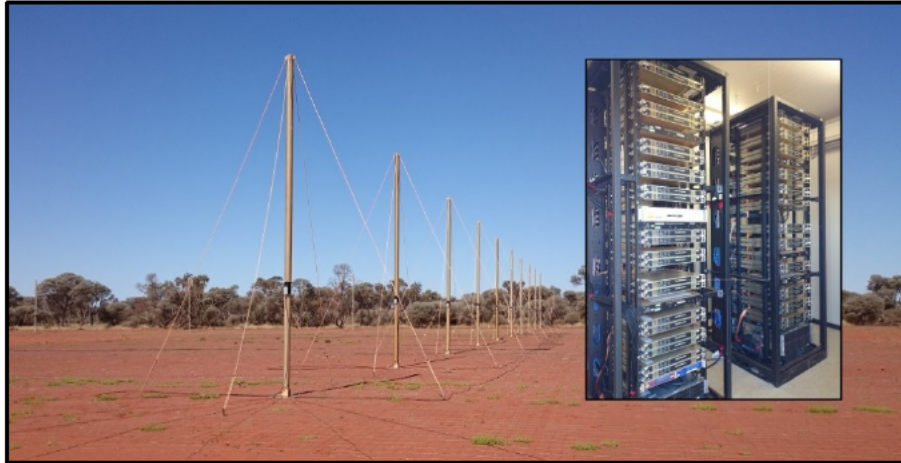
Figure 3.7. Photos of the Laverton ELOISE array, from the end of the short arm, and (inset) the digital receiver racks inside the equipment hut.
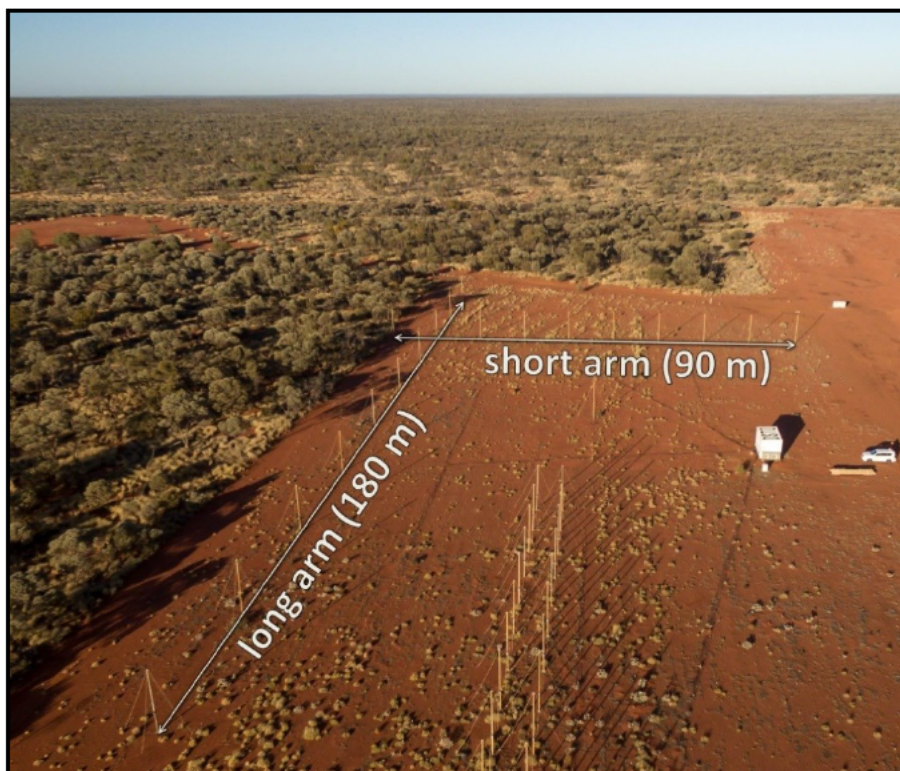
Figure 3.8. An aerial photo of the Laverton ELOISE array, with the two orthogonal arms marked. Photo supplied by Gordon Frazer.

information was available from the ELOISE array setup. The 2D imaging problem is explored further in **Appendix B**, using sample data from the ELOISE array, and the verdict is that the spatial domain offers little extra in terms of mode resolvability.

Since the antenna elements are notionally omni-directional, there is a grating lobe ambiguity problem caused by the spatial under-sampling (i.e. elements spaced by more than half a wavelength above 7.5 MHz). For an array designed to operate across more than four octaves (2–36 MHz), this is near impossible to eliminate without reducing the inter-element spacing to the point where mutual coupling becomes unacceptable. Instead, the ELOISE array layout deliberately increases the spacing as much as possible to maximise the effective aperture, while keeping the nearest ambiguities (i.e. at the upper end of the HF band) outside physical sky-wave propagation limits.

This is illustrated in **Figure 3.9**, in terms of steer angles on each arm, and **Figure 3.10**, in terms of bearing/elevation, for the Laverton array at a frequency of 30 MHz. The large "x" and "+" markers represent the nominal directions of 1-hop E and F2 propagation, respectively, for each of the eight AoA paths, and the smaller markers (connected by dashed lines) indicate the nearest ambiguities. An "exclusion" window of at least ±20° in both bearing and elevation, over the entire HF band, ensures that such ambiguities can be suitably managed using the 1-hop great circle geometry as a reference. In practice, the window is actually much wider than this at typical propagation frequencies (e.g. at least ±80° at 10 MHz). Based on DST Group's prior experience in HF radar, and subsequent results from the ELOISE experiment itself, only the most extreme disturbances could lead to AoA deviations which fall outside this unambiguous window for the oblique paths in **Table 3.1**.

Where multiple hops are present, **Figure 3.11** shows that the elevation offset from the 1-hop mirror mode at the same group delay never exceeds about 3° on the shortest Laverton path (South Hedland) or 7° on the longest path (Scherger). This is for up to 3-hops and virtual heights of up to 500 km. Combination E/F modes (i.e. the so-called M and N mode geometries depicted in Figure 4.8 of McNamara [1991]) are also expected to have elevation offsets within these ranges. The significance of this is that the largest multi-hop elevation offset is always substantially less than the smallest exclusion window of ±20° for the ELOISE array, so the use of 1-hop geometry to resolve grating lobe ambiguities is a sound strategy.
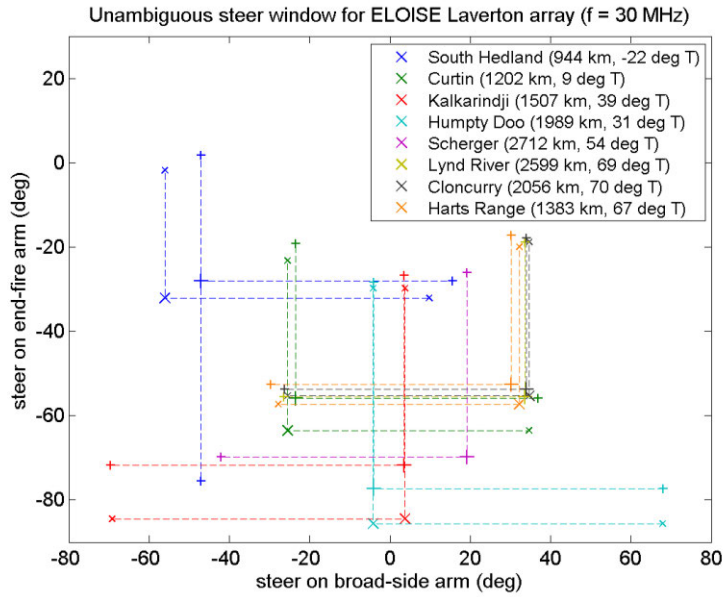
Figure 3.9. Nominal steer angles of arrival for sky-wave E and F propagation on each of the Laverton AoA paths (represented by the large "x" and "+" markers, respectively). Mirror reflection geometry is assumed, with representative 1-hop virtual heights of 100 km (E) and 300 km (F2). The nearest ambiguities at 30 MHz, caused by spatial under-sampling of the array, are shown by the smaller markers, connected to the true locations by the horizontal and vertical dashed lines. These define the limits of the unambiguous ("exclusion") window.
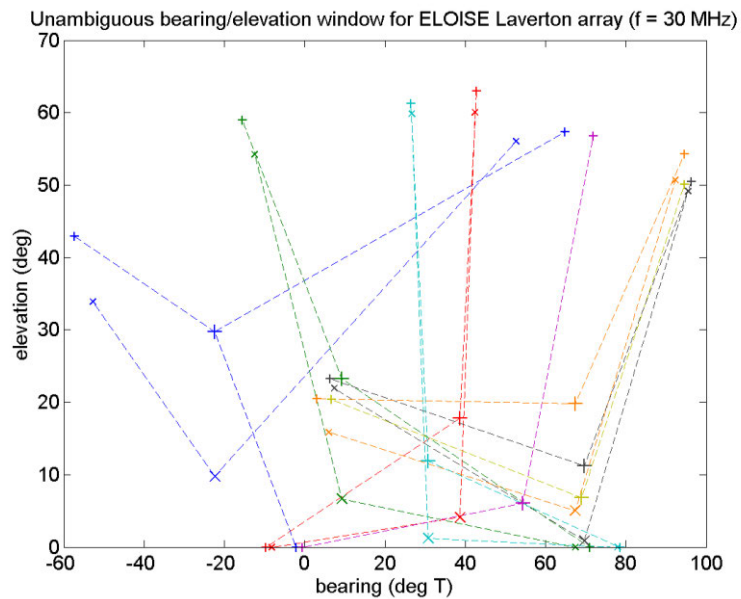
Figure 3.10. As above, but transformed to the bearing/elevation domain.
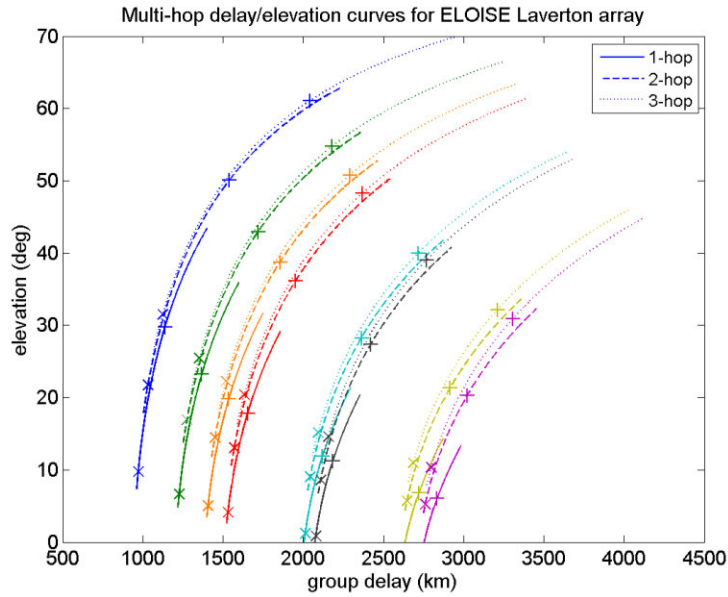
Figure 3.11. Elevation spread between different multi-hop geometries on each of the Laverton AoA paths (colour-coded the same as in **Figure 3.9**). The solid, dashed and dotted lines in each colour are for 1-, 2- and 3-hop mirror reflection, respectively, from virtual heights of 80–500 km. Nominal E and F mode elevations, assuming virtual heights of 100 km and 300 km, are represented by the "x" and "+" markers, respectively, where the lowest elevation is for 1-hop and highest elevation is for 3-hop.

Each antenna element was connected to its own direct-digital HF receiver, of the same design as those used previously in DST Group's Digital Oblique Receiving System (DORS) [Gardiner-Garden, Ayliffe, et al., 2008; Gardiner-Garden et al., 2011; Ayliffe et al., 2019; the latter is also reproduced in **Appendix F.2**]. These wideband receivers, developed specifically for HF radar applications [Frazer, 2007], support multiple independent down-converter channels, allowing many oblique propagation paths to be observed simultaneously. In this case, both 8- and 3-channel variants (NexGenR and MkIID models) were utilised at Laverton and Coondambo, respectively. Note that the front-end filters of the MkIID have a slightly more limited passband (5–35 MHz) than those in the more modern NexGenR (2–45 MHz), although it is still possible to make observations outside this range with reduced gain.

Existing DORS command and control software [Ayliffe et al., 2019] was used to schedule and collect raw sounding data on each receiver channel; that is, representing the in-phase and quadrature (I/Q) 32-bit floating-point samples on all 19 antennas after

50

digital deramping and down-conversion (also known as stretch processing, described in **Appendix A**). This software also managed the automatic gain control, which was synchronised between all receivers on the array and tasked to maintain ~10 dB headroom on the analogue-to-digital converters. The on-board processing, described further in **Section 4.2**, operated directly on the raw data, generating both CSF schedules for slaved sounder sites and processed data files that were sent back to a central server at DST Group Edinburgh, South Australia in near real time. Raw data were also recorded locally to support post-processing and further algorithm development, although for the most part, it is the results from the on-board processing that are presented in this thesis.

3.3　Array calibration

The ELOISE AoA system was designed to minimise the potential for calibration errors on the receiving array, including the use of electrically matched feed cables, a common GPS-disciplined time and frequency reference, and uniform element spacing along each arm. Nevertheless, small positional/timing errors, mutual coupling effects, and receiver mismatch can all pose accuracy limitations on the AoA estimates if left uncorrected. Due to their significant size and the complex medium through which their signals propagate, calibration of HF sky-wave arrays is a non-trivial problem, and requires many spatially-distributed sources, ideally in the far-field, to solve for the above factors [Solomon, 1998]. Among the main design limitations considered as part of the ELOISE array calibration problem were the following:

1　While the arrays were deployed by a highly skilled team of radio-frequency engineers from DST Group, making use of a theodolite and precision compass to place antenna elements in the field, there remains some uncertainty in the overall rotational alignments due to limitations in relating Magnetic North to True North. These uncertainties are expected to be on the scale of 1°.

2　For such a sparse and asymmetric 2D array, the influence of mutual coupling effects on the measured phase angle can vary considerably between elements. Furthermore, the challenge of operating over a wide frequency band (nominally 2–36 MHz) means that the array becomes quite over-sampled at the longer

wavelengths (element spacing ~10 m << wavelength ~100 m), making it even more prone to mutual coupling.

3 Operating at very low elevations makes measurements more susceptible to the finite ground mat properties (as specified in **Figure 3.6**), local multi-path from nearby reflecting surfaces, and potentially tropospheric refraction.

Each of these items will be explored further in **Chapter 5**, based on both modelling and large aggregated sets of ELOISE AoA estimates, analysed after completion of the experiment. However, for the purpose of on-board processing, a preliminary set of empirical correction factors were derived to partially address items (1) and (2) prior to the ELOISE experiment commencing. Distributions of the fitted phase residuals and absolute bearing estimates from the least-squares AoA fitting technique, described in **Section 4.2.2**, were used as the key measures for this preliminary calibration exercise. Gain differences between the elements/receivers were not explicitly addressed, as the interferometric AoA technique does not use signal amplitude, although these differences were found to be on the scale of ±1 dB.

Firstly, an assortment of sky-wave signals spanning various angles of arrival were observed on each of the Laverton and Coondambo arrays, and a fixed (frequency- and direction-independent) timing offset that minimised the overall spread of AoA fit residuals was determined for each array element. These offsets, typically of the order of ±2 ns, were adjusted manually by eye, one at a time, until the peaks in the residual distributions were roughly centred about zero. Secondly, known line-of-sight and surface-wave signals, with their more precisely defined "truth" directions, were used to provide a rotational (bearing) correction for the array. These two steps are illustrated for the Coondambo array in **Figure 3.12** and **Figure 3.13**, respectively. Phase residuals (in radians) have been converted to time residuals (in ns) for ease of comparing different frequencies, and bearing offsets are measured with respect to the true line-of-sight or great circle bearing.

Between the top and bottom panels of **Figure 3.12** there is a marginal reduction in the root-mean-square error (RMSE) of the fit (thick black line, calculated as per equation (4.7) in **Section 4.2.2.2**), indicating that the apparent wavefronts across the array are slightly more planar after the timing corrections. Nevertheless, a significant amount of
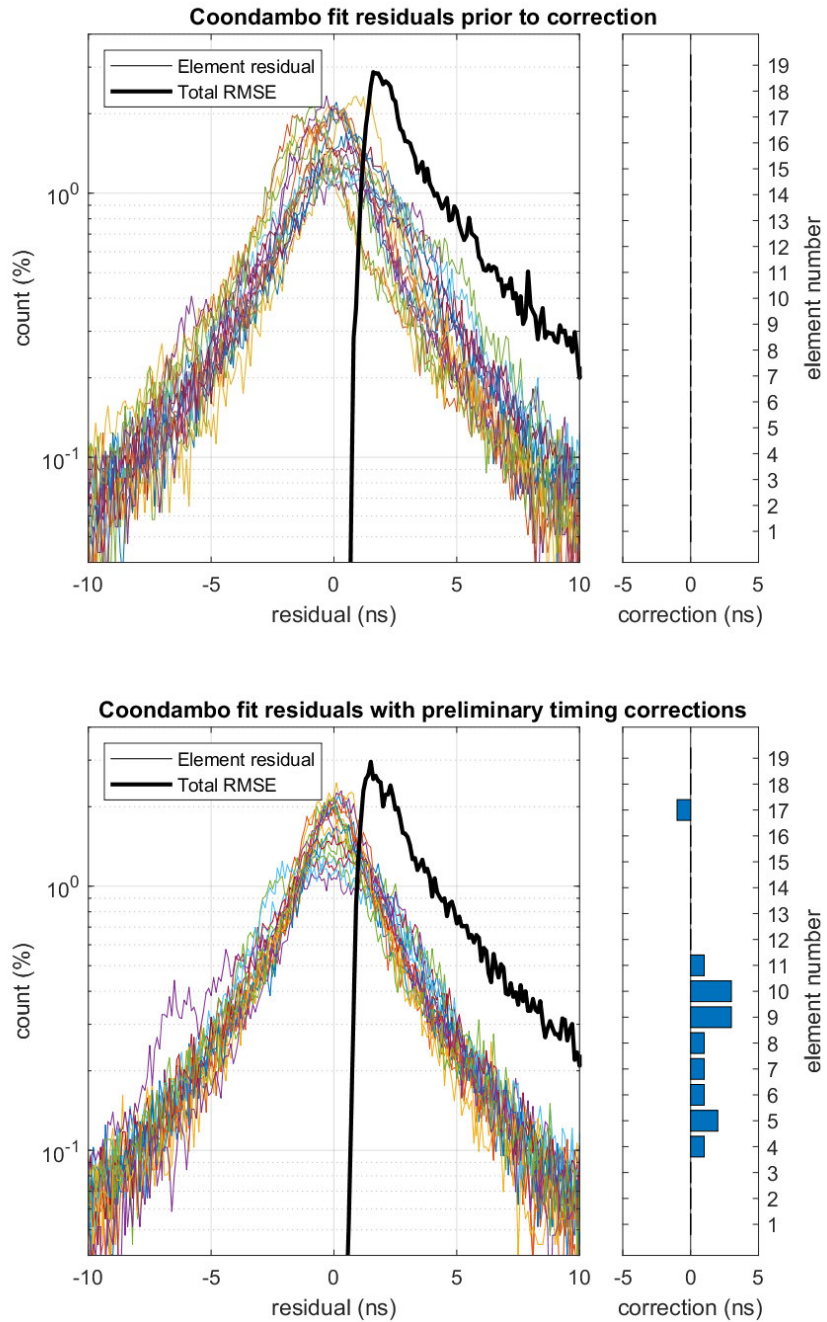
Figure 3.12. Histograms of AoA fit residuals before (top) and after (bottom) applying preliminary timing corrections, illustrated by the bar graphs on the right. Results are for a single thresholded ionogram, containing sky-wave propagation from Harts Range to Coondambo on 2015/08/26 05:30 UT. Only a modest improvement in the total RMSE distribution is attained with this simple representation. The remaining spread in the distributions is due to both calibration errors that are not adequately captured by this representation and physical wavefront distortions imparted by the ionosphere.

spread remains in these sky-wave observations. Closed-loop receiver testing, using a directly injected waveform, identified that the component due to receiver and cable mismatch (but excluding ionospheric distortions and positional errors) is less than ±1 ns. The reduced RMSE for line-of-sight observations in **Figure 3.13** suggests that much of the remaining non-planarity in **Figure 3.12** is indeed ionospheric in origin. A rather distinct positive bias of at least +1° is present in the bearing estimates of **Figure 3.13**, pointing to the need for a rotational correction, although it falls short of accounting for the frequency dependence in the top ("ionogram") panel.

Another example of estimating the rotational correction, this time for the Laverton array, is shown in **Figure 3.14**. In this case, a remote surface-wave signal was used; histogram counts are lower than **Figure 3.13** due to waveform discontinuities that prohibited overlapped processing. Once again, a positive bias is present in the bearing estimates. Given the above results, a preliminary rotational correction of -1° was adopted in the on-board processing of both Laverton and Coondambo arrays.

To provide further confidence in the calibrated output, an independent validation was set up using the nearby JORN R2E uniform linear array (ULA) at Laverton [Colegrove, 2000; Johnson et al., 2008; Sinnott, 2015]. Over a short test period, AoAs for 1-hop F2 mode propagation were measured simultaneously as CSF dwells on both ELOISE and JORN arrays; these are presented in **Table 3.3** for a single path (Cloncurry to Laverton). Results are specified as equivalent steer (cone) angles with respect to the JORN array boresight, each representing the azimuth angle of a conically ambiguous ULA beam at zero elevation. The ELOISE steer angles were derived from the joint AoA solutions on both arms (see **Section 4.2.2**), while the JORN steer angles (with inherent coning ambiguity and specified to 0.5° precision only) came from offline analysis and peak fitting of JORN baseband data using DST Group's radar signal processing test bed.

With the JORN radar being engineered and maintained to strict operational requirements, it is taken to be the "truth" in this comparison. At 3 km in length, it is also a substantially larger and more sensitive array. Despite some remaining discrepancies of up to 0.5°, which cannot be explained by the 1.6 km array separation alone, there is generally improved agreement between the ELOISE and JORN equivalent steer angles after applying the rotational correction to the ELOISE AoA estimates. It is noted that delay and Doppler estimates were also in very close agreement.
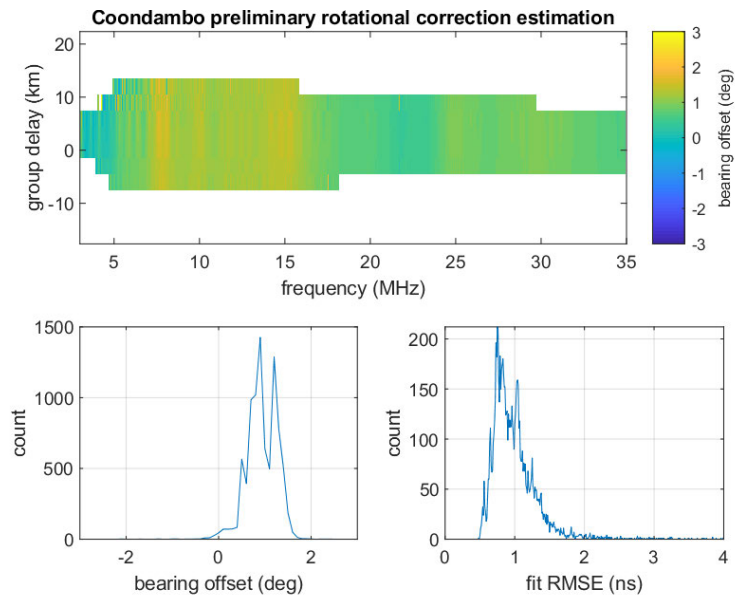
Figure 3.13. Line-of-sight (near-field) observations of a test transmitter located ~2 km north of the Coondambo array. The three panels show the thresholded "ionogram" (top) along with histograms of the bearing offset (bottom left) and fit RMSE (bottom right). Bearing offsets are measured with respect to the true bearing from the common element; a positive offset indicates a clockwise rotation in the apparent bearing, likely caused by an array boresight that is in fact less than the design value in **Figure 3.5**.
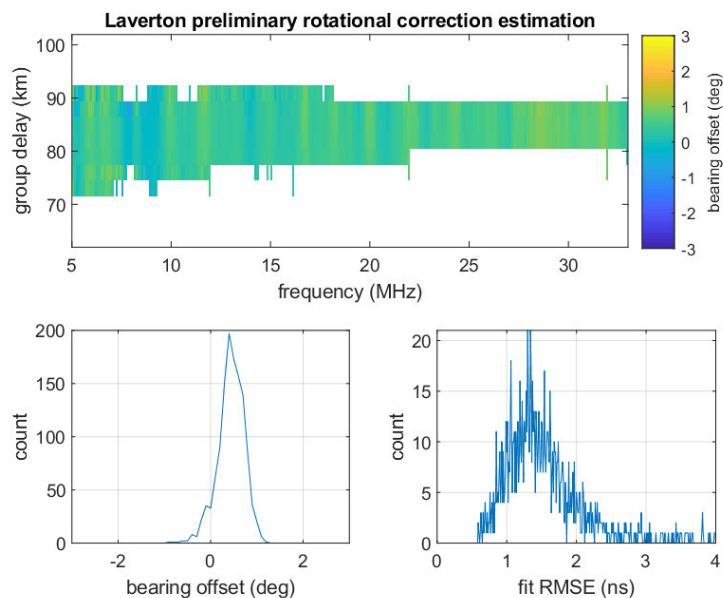
Figure 3.14. Surface-wave observations of a remote sounder transmitter located ~80 km east of the Laverton array. The panels are the same as in **Figure 3.13**. Bearing offsets are measured with respect to the great circle path bearing.

Table 3.3. Equivalent steer angles, measured with respect to the JORN R2E boresight of 35 °T, for simultaneous 1-hop F2 mode observations on the nearby ELOISE and JORN R2E arrays at Laverton. JORN R2E is a 480-element uniform linear array (ULA), spanning almost 3 km in length; it is separated from the ELOISE array by a distance of 1.6 km (centre to centre). The oblique path, from Cloncurry to Laverton (with great circle path bearing 70 °T), was observed over two CSF dwells at 12.183 MHz (22:03:39 UT) and 15.356 MHz (22:15:06 UT) on 9 August 2015. ELOISE results are presented both before and after the -1° rotational correction, and include the fitted bearing ($\hat{\theta}$) and elevation ($\hat{\beta}$) estimates used to derive the equivalent steer (cone) angles ($\hat{\alpha}$) according to the relationship $\sin(\hat{\alpha}) = \sin(\hat{\theta} - \theta_{\text{boresight}})\cos(\hat{\beta})$ [Fabrizio, 2013, eq. (4.73)]. JORN results are specified to 0.5° precision, comparable to the typical radar beamwidth at these frequencies.

	ELOISE array (uncorrected)	ELOISE array (corrected)	JORN R2E array
F2-low mode 22:03:39 UT	$35.2° \pm 0.5°$ ($\hat{\theta} = 70.2$ °T, $\hat{\beta} = 1.8°$)	$34.2° \pm 0.5°$ ($\hat{\theta} = 69.2$ °T, $\hat{\beta} = 1.8°$)	34.5°
F2-low mode 22:15:06 UT	$35.5° \pm 0.3°$ ($\hat{\theta} = 71.1$ °T, $\hat{\beta} = 9.2°$)	$34.5° \pm 0.3°$ ($\hat{\theta} = 70.1$ °T, $\hat{\beta} = 9.2°$)	35.0°
F2-high mode 22:03:39 UT	$35.0° \pm 0.6°$ ($\hat{\theta} = 71.1$ °T, $\hat{\beta} = 13.4°$)	$34.0° \pm 0.6°$ ($\hat{\theta} = 70.1$ °T, $\hat{\beta} = 13.4°$)	34.0°
F2-high mode 22:15:06 UT	$34.2° \pm 0.4°$ ($\hat{\theta} = 70.2$ °T, $\hat{\beta} = 13.2°$)	$33.2° \pm 0.4°$ ($\hat{\theta} = 69.2$ °T, $\hat{\beta} = 13.2°$)	33.0°

3.4 Chapter summary

- ELOISE was a unique experiment to observe ionospheric variability on a range of spatial and temporal scales, and with a variety of different instruments, set within Australia's dense network of HF ionosondes. A core collection period of about two weeks in September 2015 was surrounded by a broader set of synoptic sounder observations from neighbouring months. All up, ELOISE included a combined 68 VIS locations and QVIS/OIS midpoints dispersed across central and northern Australia, along with JORN backscatter radar and airglow observations.

- High-fidelity AoA measurements were made on a set of eleven intersecting oblique paths using two purpose-built HF receiving arrays. Nine of these paths were also equipped to make CSF measurements of the Doppler spectra. The array layout was designed to maximise the aperture size, subject to available land constraints, while managing the potential for ambiguities (grating lobes) due to spatial under-sampling. Path geometry predictions, based on mirror reflection, and prior experience of oblique off-angle propagation in the Australian region were each essential parts of the design methodology.

- Array calibration was carried out ahead of the core ELOISE period using a combination of known line-of-sight, surface-wave and aggregated sky-wave signals, covering different angles of arrival. Both the overall array alignment and inter-element phase delay offsets were corrected on scales of up to $\pm 1°$ and ± 2 ns, respectively, to compensate for small positional, feed delay, mutual coupling, and receiver mismatch errors. Side-by-side comparisons against the nearby JORN radar at Laverton confirm that the ELOISE array is able to deliver sufficiently accurate AoA estimates after corrections are applied, with steer angles agreeing to within $0.5°$, despite the much smaller dimensions of the ELOISE array with respect to the JORN array.

4 Signal processing of ELOISE angle-of-arrival data

4.1 Overview of on-board and offline signal processing

A significant amount of work in the lead-up to the ELOISE experiment was focused on developing automated real-time on-board processing to remotely monitor and evaluate observations collected by the two AoA arrays. The processing software, written by the author as part of this project, had to both reduce data volumes substantially, while also preserving sufficient detail and precision in the measured AoAs. Real-time clear channel evaluation and CSF scheduling algorithms were also implemented.

To support the offline (post-trial) analysis of ELOISE ionograms and CSF images, a peak detection, estimation and mode classification algorithm was written, allowing the 1-hop F2 mode to be isolated for subsequent processing. By applying simple mirror geometry for a spherical Earth and non-concentric spherical reflector, these classified peaks can then be transformed to reflection points in latitude and longitude, along with ionospheric tilt estimates at the path midpoint. This enables a more meaningful comparison of AoA estimates over different oblique path lengths and orientations.

This chapter describes the main signal/image processing capabilities of the ELOISE AoA system, using a small sample of the results collected during the 2015 campaign to illustrate each analysis step.

4.2 On-board signal processing and communications

A suite of on-board signal processing was deployed on the main server of each ELOISE array, designed to keep up with the ~4000 AoA ionograms and ~10000 CSF dwells per day at the peak of the experiment. The on-board processing included the following components, some of which were adapted from DST Group's existing DORS ionosonde [Gardiner-Garden, Ayliffe, et al., 2008; Gardiner-Garden et al., 2011; Ayliffe et al.,

2019], a modern oblique-incidence chirp-sounder receiver, where noted. All software components, other than the core command and control tools, were implemented in MATLAB®, with a parallel process running for each digital down-converter channel (oblique path).

- A fast and efficient implementation of ionogram and delay-Doppler map processing using basic chirp-sounder principles. This incorporated many components of already established DST Group software. A brief review of the fundamentals of chirp sounding is provided in **Appendix A**.

- A robust radio frequency interference (RFI) rejection scheme, adapted from the DORS ionosonde, to improve the clarity of ionogram trace features through auto-regressive interpolation of impulsive raw data samples [Turley et al., 2019; also reproduced in **Appendix F.4**]. This is run independently on each receiver channel prior to AoA estimation.

- An ionogram feature extraction and fitting algorithm, also adapted from the DORS ionosonde, to extract key image features such as the maximum observed frequency (MOF) of F2 propagation, and parameterise the midpoint electron density profile in terms of a multi-segment quasi-parabolic function with up to two embedded sporadic-E (Es) mirror reflectors [Heitmann & Gardiner-Garden, 2019; also reproduced in **Appendix F.3**].

- A clear channel evaluation (CCE) algorithm and adaptive scheduler for making channel scattering function (CSF) observations, to characterise the ionospheric Doppler spectra at key operating frequencies.

- A 2D angle-of-arrival (AoA) technique to extract bearing and elevation estimates for each range and frequency cell combination in the ionogram/CSF data, based on a least-squares planar wavefront fit to the elemental phases.

All transmitter and receiver sites communicated via a virtual private network, which enabled schedules to be distributed and processed results to be retrieved in near real time throughout the campaign. The standard files that were transferred to DST Group's central data server in Edinburgh, South Australia contained both thresholded and byte-scaled ionogram and CSF images, as well as the ionogram features and fitted parameters. The typical file size after data compression, for a single 3.75 min

observation interval on a given path, was 200–300 kB. Each AoA record was represented by a 3D data cube consisting of up to five pages of frequency versus group delay images; these contained amplitude, bearing and elevation estimates, along with bearing and elevation uncertainties.

The following sub-sections outline the specific properties, including scheduling and processing steps, relating to the ELOISE CSF collection and AoA estimation. Much of this content has been previously published in Heitmann et al. [2018] (also reproduced in **Appendix F.1**).

4.2.1 Channel scattering function scheduling and processing

Direct measurements of ionospherically-induced Doppler shifts on OIS paths provide information about the dynamics of the ionosphere, including TIDs. The ability to measure the full Doppler spectra, rather than just a single dominant Doppler component at each ionogram frequency and delay, was favoured for the ELOISE experiment, as it provided an extra dimension for separating propagation modes and characterising spread-Doppler conditions. The channel scattering function (CSF) mode, using a linear frequency-modulated continuous-wave (LFMCW) waveform centred on a fixed frequency, enabled delay-Doppler images to be collected alongside each OIS ionogram, without the need for any additional hardware. A similar technique has been demonstrated in previous experiments by DST Group [Harris & Frazer, 2005; Frazer et al., 2010; Fabrizio, 2013, ch. 6.2], although not with the same level of automation developed for ELOISE.

Six of the existing Digital Oblique Transmitting System (DOTS) OIS transmitter sites [Ayliffe et al., 2019] had their software upgraded ahead of the experiment to support an interleaved CSF operating mode, whereby each wideband (ionogram) sweep was preceded by three 6.6 s CSF dwells with a 10 Hz waveform repetition frequency (WRF). Effectively the CSF waveform consisted of 66 short 0.1 s sweeps (or "chirps", to distinguish from the wideband sweep), although only 64 of these were coherently processed, to allow for edge effects in the data series. Recognising the temporal decorrelation caused by dynamics in the ionosphere, this observation interval was chosen in line with typical values for HF sky-wave measurements, where coherent

integration times (CITs) of the order of ten seconds are the norm [Anderson & Lees, 1988]. The same output power (nominally 20 W) applied to both wideband and CSF waveforms. To accommodate the additional processing and CSF dwell time, the frequency limits of the wideband sweep were reduced from 2–45 MHz to 2–36 MHz, without any change to the 3.75 min revisit interval or 250 kHz/s sweep rate.

Adaptive scheduling of the CSF frequencies (within the ionogram sweep limits) and bandwidths (from 7.5 to 30 kHz) was controlled by a clear channel evaluation (CCE) process running on-board each of the Laverton and Coondambo receiving systems. The purpose of having real-time CCE was to ensure that, as a secondary user, ELOISE CSF transmissions were made on a non-interference basis with other active HF users and avoided known forbidden frequencies (e.g. those reserved for search and rescue). At any given time, each DOTS transmitter had its CSF scheduling assigned to either one of the Laverton or Coondambo sites, while the other site remained slaved. Real-time schedules were considered stale after 5 min and replaced by a fixed default schedule to prevent rogue transmissions when remote communications were unexpectedly lost.

The CCE algorithm is based on the approach used for frequency advice on the JORN radars [Earl & Ward, 1986; Earl & Ward 1987]. Each CCE system was fed by background narrowband noise data, extracted from the "guard band" of DORS ionograms (i.e. at group delays less than the great circle ground range), for just the single omni-directional antenna common to the two array arms. The entire 2–36 MHz band is processed into contiguous and overlapping 8 kHz channels. Any channel in which the median noise power spectral density is less than 7 dB above the lower decile over the surrounding 1 MHz interval is deemed to be *unoccupied*. Those 8 kHz channels that remain unoccupied for at least three of the last four observations (within 15 min) are deemed to be *clear*. Only contiguous blocks of clear and not forbidden channels spanning up to 30 kHz in bandwidth are candidates for CSF transmission.

Figure 4.1 and **Figure 4.2** show examples of the background narrowband noise data and CCE output, respectively, for 24 hours of Laverton observations. The F2 MOF for the Cloncurry to Laverton path is overlaid, exhibiting the typical pattern of diurnal variation embedded with smaller-period ionospheric disturbances. At a remote site such as this, most of the strong interferers (horizontal striations) and broad-band background noise in **Figure 4.1** are ionospherically propagated sources at very long distances,
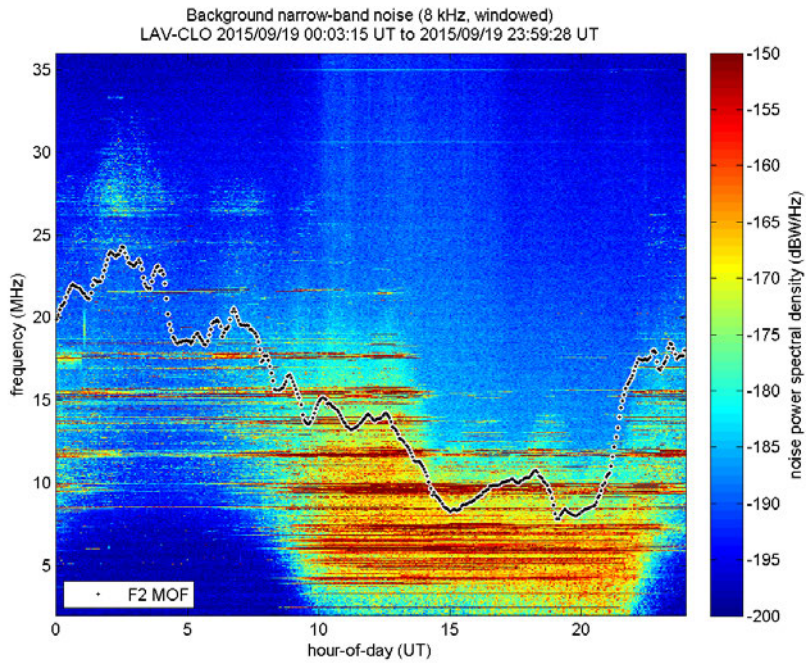
Figure 4.1. Background narrowband noise data from Laverton on 19 September 2015, with the F2 MOF for the Cloncurry to Laverton (2056 km) path overlaid. The local solar time at Laverton is approximately UT + 8 hr.
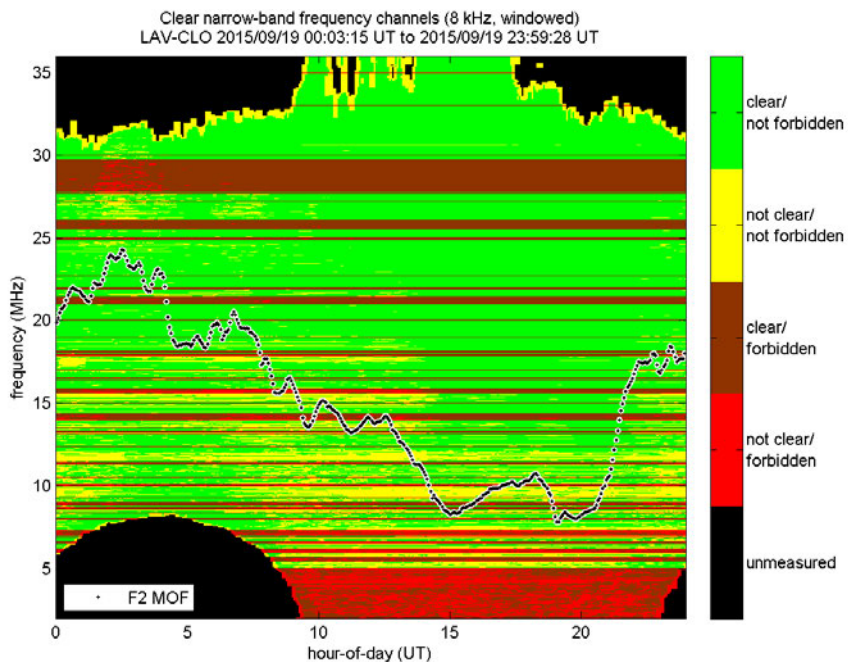
Figure 4.2. CCE output from Laverton on 19 September 2015, with the F2 MOF for the Cloncurry to Laverton path overlaid.

including terrestrial broadcast and communication signals, other HF radars/sounders, and natural phenomena such as lightning discharges. Radio sources above the ionosphere, including galactic emitters and satellites, also contribute to the background noise at higher frequencies. CCE identifies the green (clear/not forbidden) bands in **Figure 4.2** in which CSF dwells may be scheduled, and only those below the F2 MOF offer propagation support. Note that frequencies found to be internally noise limited on the DORS receiver are rejected and shown in black on **Figure 4.2**; however, these are not operationally useful frequencies, being contained to the lower HF band during the day and upper HF band during the night.

Target CSF frequencies are based on fixed logarithmic ratios of the F2 MOF, extracted automatically from the previous ionogram, and the frequency selection maintains temporal persistence whenever permitted by channel availability rules. Except at night, when the lower HF band becomes significantly congested, the preferred CSF bandwidth of 30 kHz was usually available in the CCE output. A time-staggered set of waveforms, with a fixed 10 Hz repetition frequency and a minimum 16 ms (~4800 km) time offset within the repetition interval, was implemented to allow the six CSF transmitters to potentially share some of the same frequency bands and further minimise the likelihood of interference. This waveform set is practically orthogonal on one-way oblique paths.

Figure 4.3 shows a sample of the CSF scheduler output for the same period as in **Figure 4.1** and **Figure 4.2**. Agile frequency and bandwidth allocations for each of the three CSF dwells scheduled every 3.75 min are indicated by the coloured markers in the top and bottom panels, respectively. At each time, the frequencies represent 72.9%, 81.0% and 90.0% of the median MOF over the last 15 min (subject to some tolerance for channel availability), which is thought to cover the band of most operational relevance to a JORN-like radar. During day-time observations, the CSF frequency allocations have a tendency to change more rapidly, as propagation support fluctuates with variations in solar irradiance and a greater prevalence of MSTIDs, although the use of the median MOF filter helps to maintain persistence.

The typical diurnal variation in the CSF bandwidth allocations across 13 days of schedules is presented in **Figure 4.4**. All paths exhibit a sharp drop in the percentage of 30 kHz dwells at night (notionally 0930–2100 UT on the Cloncurry to Laverton path), with a further drop in the pre-dawn period (1800–2100 UT). The day/night variation is

most severe on the shorter paths, for which propagation support is more limited in frequency. The penalty of operating with a narrower bandwidth is of course reduced group delay resolution.
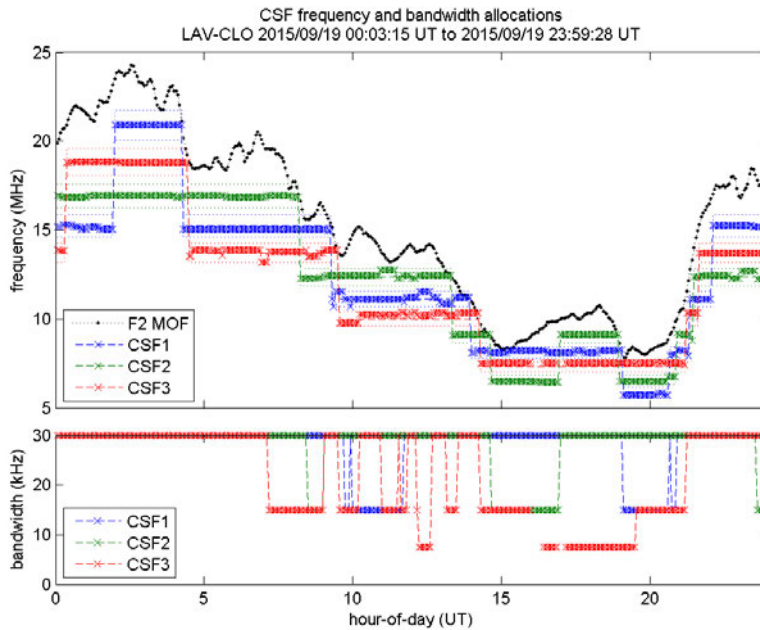
Figure 4.3. CSF scheduler output for the Cloncurry to Laverton path on 19 September 2015. The top and bottom panels show the agile frequency and bandwidth allocations, respectively, for each of the three CSF schedule entries (blue, green and red).
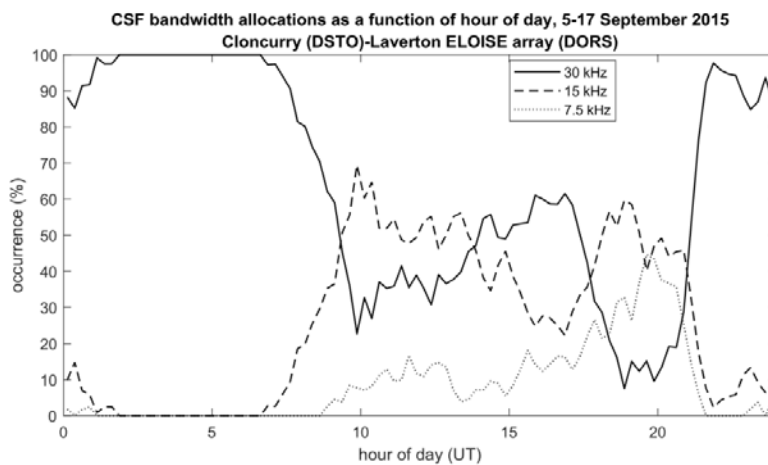
Figure 4.4. Percentage occurrence of CSF bandwidth allocations as a function of hour of day, from 5–17 September 2015 on the Cloncurry to Laverton path. For each 15 min time bin, the three lines representing 30, 15 and 7.5 kHz will sum to 100%.

Although the CSF frequencies were selected with a specific oblique path in mind, the ELOISE virtual private network facilitated schedule distribution to multiple OIS transmitter and receiver sites, enabling the collection of CSF dwells wherever ionospheric propagation existed. As such, the six transmitters were typically received at both Laverton and Coondambo simultaneously, giving up to twelve CSF paths. However, only nine of these paths (six into Laverton and three into Coondambo) were measured on the AoA arrays and, thus, have been listed in **Table 3.1** of **Section 3.1**.
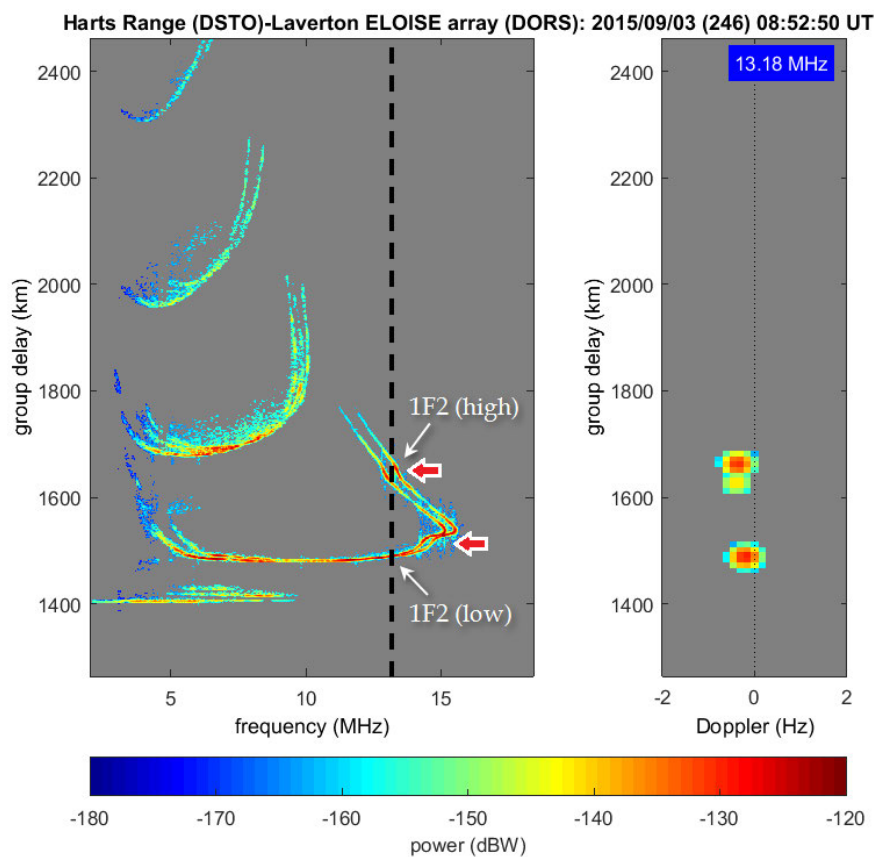
Figure 4.5. A sample ionogram (left) and corresponding CSF delay-Doppler image at 13.18 MHz (right), for the 1383 km OIS path from Harts Range to Laverton. The key 1-hop F2 low- and high-angle propagation modes at the CSF frequency (dashed black line) are annotated on the ionogram; see e.g. Davies [1990] or McNamara [1991] for a more detailed description of ionograms and their interpretation. The red arrows highlight the trace distortions ("kinks" or "wiggles") attributed to MSTIDs. Note that each pixel combines all elements of the array using a conventional phased array beam in the direction of the estimated AoA.

A sample CSF delay-Doppler image is shown in **Figure 4.5**, alongside an ionogram that probed the same frequency 69 s later. The CSF has an unambiguous Doppler window of ±5 Hz, although in practice ionospheric Doppler shifts rarely exceed ±1 Hz [Davies, 1990, p. 239], so the image here has been cropped for clarity. The kink-like distortions on the 1-hop F2 trace (indicated by the red arrows) are the characteristic signatures of MSTIDs (e.g. Lobb & Titheridge [1977]; Harris et al. [2012]). The same group delay structure is clearly seen in both images, albeit at lower resolution in the CSF owing to its lower processing bandwidth (i.e. 30 kHz versus 200 kHz). Under more complicated propagation conditions, it is possible for the CSF measurement to detect weaker modes than in the ionogram, owing to its longer CIT (6.4 s versus 0.8 s) and more judicious choice of frequency. This particular case is representative of mid-latitude ionospheric conditions in the early evening, when the F2 layer height is increasing, as evident from the slightly negative Doppler shifts in the CSF image.

While high-resolution parametric and subspace-based spectral analysis techniques were considered, the on-board delay-Doppler processing for CSF image generation followed a conventional matched-filtering approach [Fabrizio, 2013, ch. 4], using a windowed fast Fourier transform without any impulsive noise suppression. This was due to a combination of the tight deployment timeframe, lesser computational load, and a priority on robustness. As an example of the potential for future improvements, **Figure 4.6** shows an enlarged version of the same CSF dwell as in **Figure 4.5** with an auto-regressive data extrapolation algorithm [Swingler & Walker, 1989; Turley & Voigt, 1992; Turley, 2008] applied prior to windowing. This improves the Doppler resolution and reduces windowing losses in the bottom panel, compared to the standard processing in the top panel, and thus provides a better chance of resolving fine multi-mode structure (although it is insufficient alone to separate O/X components on the F2-low mode). Similar extrapolation techniques may also be applied in the frequency domain, to increase the effective signal bandwidth and range resolution [Nguyen & Turley, 2015], although this has not been tested on ELOISE CSF data.

Note that as with any swept waveform, including the wideband DORS sounder sweep, there is an inherent delay-Doppler ambiguity in the ELOISE CSF observations. With a 10 Hz WRF, the CSF signal is effectively sweeping at 75–300 kHz/s (for 7.5–30 kHz bandwidth), and delay-Doppler coupling [Fitzgerald, 1974] leads to a group delay error of between 4 km and 1 km for a representative 1 Hz ionospheric Doppler shift. This is

equivalent to 10% of a delay cell, and as such, the error can usually be neglected. For the ionogram sweep rate of 250 kHz/s, this increases to a 1.2 km error, or 80% of a delay cell. A positive/negative Doppler frequency shift will correspond to a negative/positive apparent delay offset for an up-chirp, causing the peak to appear slightly closer/further in delay than reality.
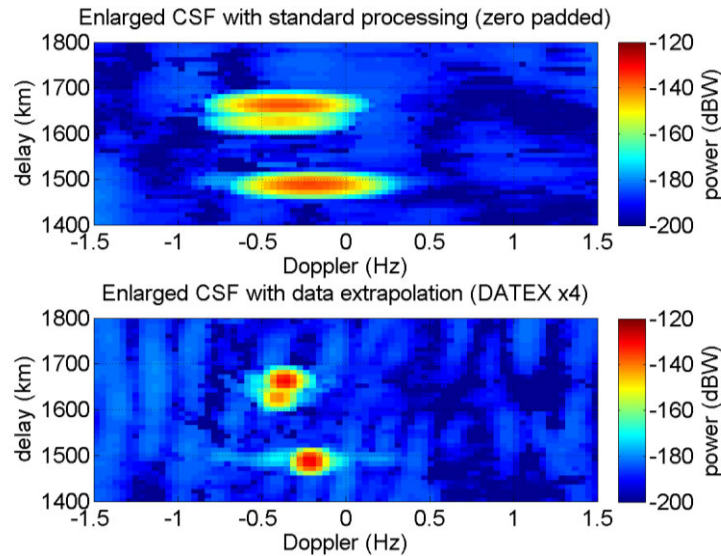
Figure 4.6. Enlarged versions of the same CSF dwell from **Figure 4.5**, analysed using both standard processing (top panel) and data extrapolation ("DATEX") by a factor of 4 (bottom panel). A Blackman-Harris 92 dB window has been applied to each, and zero padding has also been included in the top panel to match the Doppler step size of the bottom panel.

4.2.2 Angle-of-arrival estimation

Both ionogram and CSF observations received on the AoA arrays were subject to on-board AoA estimation, which applied an interferometric technique to fit (in a least squares sense) a planar wavefront jointly across the phase angles on all 19 elements of both array arms. Although many more sophisticated algorithms exist for estimating the direction of radio wave propagation with phased arrays, including subspace-based methods such as MUSIC and ESPRIT, and parametric methods such as auto-regressive model fitting [Gething, 1991], these are largely considered overkill (i.e. slower) and/or less robust for a high signal-to-noise-ratio (SNR), discrete-mode problem such as one-

67

way oblique HF reception. For a single propagation mode, testing of the ELOISE implementation has confirmed that the AoA result, using only phase information, is indeed consistent with the direction of peak output power from a conventional beamformer (see **Appendix B**).

The AoA estimation technique was carried out independently for each range and frequency cell combination in the ionogram/CSF data, using the predicted great circle bearing and 1-hop equivalent mirror elevation (based on group delay only) to resolve grating lobe ambiguities. The principal assumption is that each range/frequency cell contains a single dominant mode; separation of modes must therefore make use of information from the other data dimensions (i.e. frequency, delay and Doppler). Given point-to-point propagation via discrete modes, this is believed to be a reasonable assumption for a significant proportion of the ELOISE data set. Where there are unresolved modes (e.g. O/X components), either one mode will tend to dominate or else the resultant non-planarity will be flagged by a corresponding increase in the fit uncertainty.

In short, the AoA estimation algorithm consists of the following steps:

1   Resolve 2π phase ambiguities on each array arm using predicted geometry.

2   Fit planar wavefronts jointly to the unwrapped phase angles across the two arms.

3   Transform the least-squares fits in wave vector space, including uncertainties, to bearing and elevation.

4   Construct AoA data cube, including a beamformed estimate of power in each pixel of the ionogram/CSF image.

The following sub-sections address each of these steps in turn.

4.2.2.1 Ambiguity resolution

For a general array element n, at position $\boldsymbol{r}_n \equiv [x_n, y_n, z_n]^T$ in a local East-North-Up (ENU) Cartesian coordinate system, the basic planar wave equation that relates the AoA wave vector $\boldsymbol{k} \equiv [k_x, k_y, k_z]^T$ to the elemental phase angle φ_n , relative to the phase at the origin φ_0, takes the following form:

$$\boldsymbol{k} \cdot \boldsymbol{r}_n = \varphi_n - \varphi_0, \tag{4.1}$$

where $\boldsymbol{k} = \frac{2\pi}{\lambda} \begin{bmatrix} \sin\theta\cos\beta \\ \cos\theta\cos\beta \\ \sin\beta \end{bmatrix}$, for a given wavelength λ, bearing (azimuth) θ, and elevation β. Note that this wave vector, which points away from the receiving array, normal to the planar wavefront, is directed in the opposite sense to the physical wave vector of the signal for notational convenience. It is assumed to be constant across the array aperture, as expected for a planar wavefront from a far-field source.

In practice, what is measured is a wrapped phase angle $\psi_n \in [-\pi, \pi)$, which differs from φ_n by an integer M_n number of cycles:

$$\varphi_n = \psi_n + 2\pi M_n, \tag{4.2}$$

where by convention $\varphi_0 = \psi_0 \in [-\pi, \pi)$ at the origin.

With the ELOISE array being under-sampled at frequencies above 7.5 MHz (i.e. antenna spacing $d > \lambda/2$ on the long arm for $\lambda < 40$ m), the integer factors M_n (for $n = 1, \dots, 19$) are potentially ambiguous. In other words, they may differ by more than one between adjacent elements and, thus, the phases cannot be trivially unwrapped by working outwards from the common element (here taken to be at the origin). Therefore, it is first necessary to resolve 2π ambiguities in the measured phase angles (with calibration corrections applied). Note that the implemented algorithm included the ability to easily remove degraded receiving elements from the AoA solution, meaning that uniform spacing could not be assumed, although this functionality was rarely used in practice.

The approach taken for ambiguity resolution is to use the known great circle bearing and 1-hop equivalent mirror elevation at the measured delay (i.e. the predicted geometry) and find the "closest" valid solution in wave vector space. This is achieved by forming the phase angle offset

$$\delta\varphi_n \equiv \varphi_n - \varphi_{n,p} = \psi_n + 2\pi M_n - \varphi_{n,p}, \tag{4.3}$$

where $\varphi_{n,p} = \boldsymbol{k}_p \cdot \boldsymbol{r}_n$ is the relative phase for predicted wave vector $\boldsymbol{k}_p$. For a single uniform linear arm, the closest solution for φ_n (defined by integer M_n) is that for which the difference in phase offsets between adjacent elements $\delta\varphi_{n+1} - \delta\varphi_n$ lies on the interval $[-\pi, \pi)$. With non-uniform spacing (e.g. if degraded elements have been removed), extra care is required to ensure that a compatible set of integer ambiguity factors are assigned across various multiples of the fundamental spacing; the root-mean-square error (RMSE) of the fit will quickly reveal if an invalid solution is found.

The end result of this step is a vector of phase angles that are effectively unwrapped along each array arm, as illustrated by the "x" markers in **Figure 4.7**. While there remain some small variations in the inter-element phase differences, due to both ionospheric distortions and uncorrected calibration errors, the phase angles defining the wavefront are very close to linear on each arm. Noting that this is a fairly typical example, the assumption of a planar wavefront is felt to be justified. In this case, the slopes are related to the steer angles on each arm; the frequency dependence may be removed by converting phase angle to phase delay.
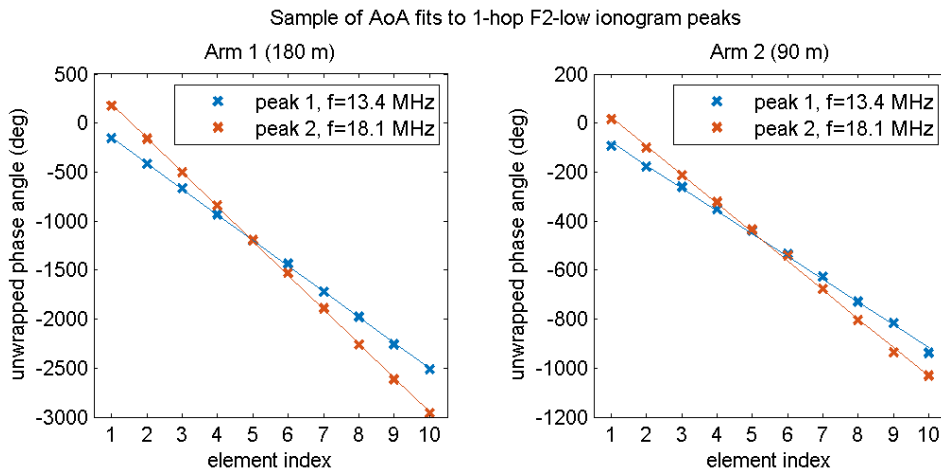
Figure 4.7. Measurements of unwrapped phase angles for two ionospheric modes at different frequencies on the 1-hop F2-low mode from Cloncurry to Laverton. The left and right panels contain observations made on the long (180 m) and short (90 m) arms, respectively. Lines of best fit are overlaid to provide a sense of linearity, although it is noted that the actual AoA estimation algorithm fits both arms jointly.

4.2.2.2 Planar wavefront fit

From equation (4.1), with $z_n = 0$ for a horizontal planar array (strictly speaking, not Earth-conformal), the linear equations that relate the phase angles on each of the 19 elements of the ELOISE array can be written in matrix form as follows:

$$X\boldsymbol{p} = \boldsymbol{b}, \tag{4.4}$$

where design matrix $X = \begin{bmatrix} x_1 & y_1 & 1 \\ \vdots & \vdots & \vdots \\ x_{19} & y_{19} & 1 \end{bmatrix}$ contains the element positions in the local coordinate system; parameter vector $\boldsymbol{p} = [k_x, k_y, \varphi_0]^T$ contains the corresponding horizontal components of the wave vector and the phase angle at the origin; and measurement vector $\boldsymbol{b} = [\varphi_1, \cdots, \varphi_{19}]^T$ contains the unwrapped phase angles on each of the elements. Note that for this step, the axes of the (x, y)-coordinate system are rotated to be aligned with the semi-major and semi-minor axes of the least-squares uncertainty ellipse for the ELOISE array geometry (the semi-major axis being 74.6° anti-clockwise/clockwise from boresight, towards the short arm, for Laverton/Coondambo). While this does not affect the fit, it does simplify the subsequent transformation of uncertainties to bearing/elevation space.

The least-squares solution to the normal equations above can then be expressed by the standard formula

$$\hat{\boldsymbol{p}} = (X^T X)^{-1} X^T \boldsymbol{b}, \tag{4.5}$$

where the hat notation is used to denote an estimated variable. In practice, the Cholesky decomposition of $X^T X$ is used for numerical stability, rather than actually computing the pseudo-inverse as written above.

Goodness of fit is measured in terms of the sum of squared phase residuals or its square-root, the RMSE. The AoA fit residuals (in radians) are defined as

$$\boldsymbol{e} \equiv [e_1, \cdots, e_{19}]^T = \boldsymbol{b} - X\hat{\boldsymbol{p}}, \tag{4.6}$$

while the RMSE (an unbiased estimator of the standard deviation) is given by

$$E = \sqrt{\frac{1}{N-P} \sum_{n=1}^{N} e_n^2},$$ (4.7)

where the mean is summed over all $N = 19$ array elements and normalised by N minus the number of parameters $P = 3$. Note that both the fit residuals and RMSE are scaled in units of time throughout this thesis; that is, by dividing by the angular wave frequency.

The resultant components of $\hat{\boldsymbol{p}}$ define the planar surface that jointly fits the phases on both arms according to equation (4.1); $\hat{k}_x$ and $\hat{k}_y$ being the partial derivatives along the horizontal (x, y) axes, and $\hat{\varphi}_0$ being the intercept at the origin. As $\hat{\varphi}_0$ is fitted across the whole array, in general it will not exactly match the observed phase on the common element φ_8. Note that for direct line-of-sight propagation from near-field sources (e.g. in **Figure 3.13**, **Section 3.3**), the wave vector varies over the array aperture, and a planar fit is no longer appropriate without modification. In this case, the wavefront can be approximated by a spherical surface to first order, allowing the phase measurements to be "flattened" prior to calculating the fit; this is described further in **Appendix C**.

As an alternative to the joint fitting approach, it is also possible to solve for the phase slope and intercept along each arm independently (i.e. as a 1D least-squares problem) then combine. In this case, the fitted slope represents the horizontal wave vector component along the axis of the arm, or equivalently, the sine of the steer (cone) angle scaled by the wave number $(2\pi/\lambda)$. Such decoupled fits have been overlaid as lines in **Figure 4.7**, for sake of illustration, although the joint solution is preferred for routine use because it does not duplicate the contribution from the common element.

4.2.2.3 Bearing/elevation transformation

Estimates of bearing (azimuth) $\hat{\theta}$, measured clockwise from True North, and elevation $\hat{\beta}$ are related to the fitted wave vector components by the simple trigonometrical equations

$$\hat{\theta} = \text{atan2}\big(\hat{k}_x, \hat{k}_y\big) + \theta_y, \text{ and}$$ (4.8)

$$\hat{\beta} = \text{acos}\left(\frac{\lambda}{2\pi}\sqrt{\hat{k}_x^2 + \hat{k}_y^2}\right), \tag{4.9}$$

where θ_y defines the bearing of the y-axis (after rotation).

Note that for a physical solution,

$$\sqrt{\hat{k}_x^2 + \hat{k}_y^2} \leq \frac{2\pi}{\lambda}. \tag{4.10}$$

Assuming no uncertainties in the element positions, the standard deviations for each component of the parameter vector estimate $\hat{\boldsymbol{p}}$ are directly related to the sum of squared residuals in the fit. This allows uncertainties in the corresponding bearing and elevation estimates to be characterised. Higher uncertainties may be indicative of multi-mode structure unresolved in frequency, delay and Doppler, for example. Measures of non-planarity and spatial mode complexity are investigated further in **Section 6.3**. Besides unresolved multi-mode, calibration errors, receiver noise, residual RFI, and spatial incoherency across the array aperture can also contribute to AoA uncertainties. Array misalignment and antenna gain patterns are further sources of potential error, but manifest as fixed biases rather than contributing to the variance.

In order to better understand how the measured phase noise on each receiver maps to signal bearing and elevation errors under this transform, simulated distributions of AoAs estimated using the above three steps are shown in **Figure 4.8**. These are based on spherical mirror model propagation across the eight oblique paths into Laverton, and for two representative virtual heights: 100 km (E mode) and 300 km (F2 mode). The effects of receiver noise, caused by instabilities in the local oscillator and a finite external SNR, were modelled by introducing a random Gaussian phase error to the theoretical wrapped phase angle ψ_n on each element. A standard deviation of 5° was selected for this phase error, based on closed-loop receiver tests at around 20 dB SNR, and results in the bearing and elevation spread pictured. This plot highlights both the general distribution of oblique paths in AoA-space (like **Figure 3.10**, **Section 3.2**), as well as some of the challenges faced by an array of this size operating at low elevations, with representative E mode returns of 0–10° and F2 mode returns of 5–30° in elevation.
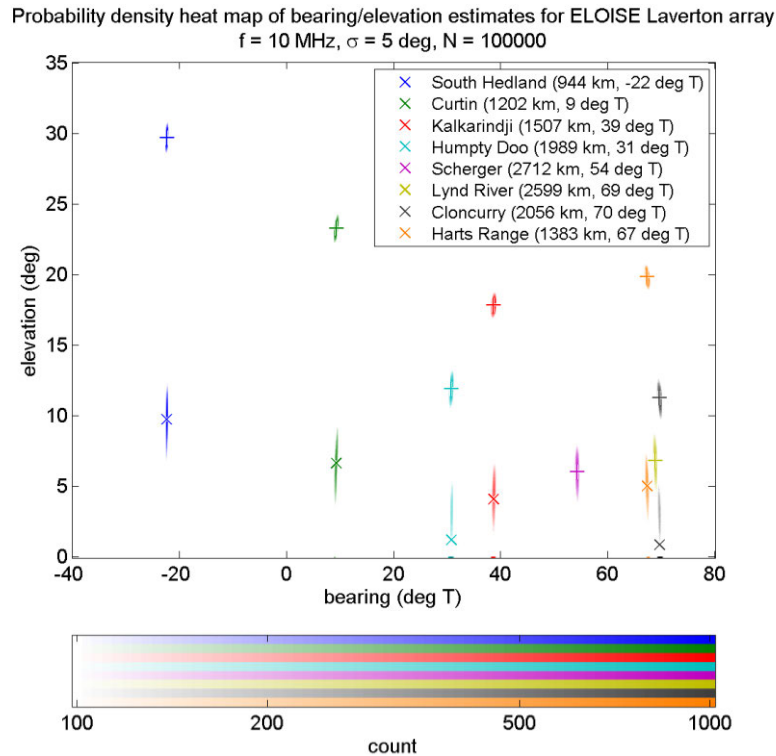
Figure 4.8. Simulated distributions of AoA estimates in bearing/elevation space at a frequency of 10 MHz, assuming white Gaussian phase noise with a standard deviation of 5°. For each oblique path, returns from both 1-hop E and F2 modes (at 100 km and 300 km virtual heights, respectively) are evaluated, without regard for the antenna gain patterns, using 100,000 randomised samples. The true AoA positions are indicated by the "x" (E mode) and "+" (F mode) markers, as in **Figure 3.10** of **Section 3.2**.

Below 10° in particular, the synthetic AoA distributions are very broad in elevation, and potentially become biased (overestimated) as symmetrically distributed errors in wave vector space map to a skewed distribution in bearing/elevation space, and unphysical solutions are clipped at an elevation of zero. The null in the antenna gain pattern at 0° elevation is believed to further exacerbate this problem in real observations, as the actual elevation distribution of the signal is modulated by the sharp roll-off in reception. Whereas the standard deviation in bearing is typically around 0.2°, for elevation this increases to 0.4–0.7° above 10°, and 1.0–1.5° below 10°. It will be seen later, in **Table 4.1**, that this bears a close resemblance to the standard deviations in real AoA estimates derived from observations. Overall performance will generally be better at higher SNR, although the geometrical limitations remain.

4.2.2.4 Angle-of-arrival data cube construction

The final product of the AoA processing is a 3D data cube representing the AoA images in each dimension (i.e. amplitude, bearing/elevation estimates, and bearing/elevation uncertainties). By default, this is quantised and thresholded to keep data volumes manageable. Byte-scaling supports 0.5 dB quantisation in amplitude (power) and 0.2° in AoA, the latter being on the same scale as the typical standard deviations quoted above.

It is advantageous to make use of the bearing/elevation information to combine array elements and form classical beams in the direction of the AoA estimate for each pixel; in theory, this offers array gains of up to ~13 dB in the amplitude (power) image. This extra sensitivity can reveal weaker mode structure due to complex scattering effects that may not be detected on a single element. Such weakly scattered returns usually have a much broader range of AoA estimates than the regularly refracted returns.

A sample of the on-board AoA output is shown in **Figure 4.9** and **Figure 4.10**, for the same sounding as in **Figure 4.5**. Bearing and elevation offsets are defined as the differences between the measured angles and the predictions, based on the great circle and 1-hop mirror geometry, while bearing and elevation uncertainties are the transformed standard deviations in $\hat{k}_x$ and $\hat{k}_y$. The uncertainties are very much dependent on SNR and tend to be larger where O/X components are unresolved, or elevation angles are low (particularly along the 1-hop E trace).

An alternative presentation of the same AoA ionogram is given in **Figure 4.11**, this time as a collection of scatter plots showing the fitted AoA estimates and their offsets as a function of frequency. The frequency axis may also be interpreted as representing (linear) time over the duration of the sounding. Many of the frequency-dependent fluctuations in AoA on the scale of 0.5–1 MHz (equivalent to 2–4 s at a sweep rate of 250 kHz/s) exceed the AoA uncertainties from **Figure 4.10**, hinting at real physical variability, though it is conceded that some of these are likely to be mutual coupling effects. A slight lag in the AoA fluctuations between O and X high-ray components also appears to be present, indicative of height-dependent perturbations in the electron density which are sampled at different locations due to O/X ray divergence [Munro & Heisler, 1956b].
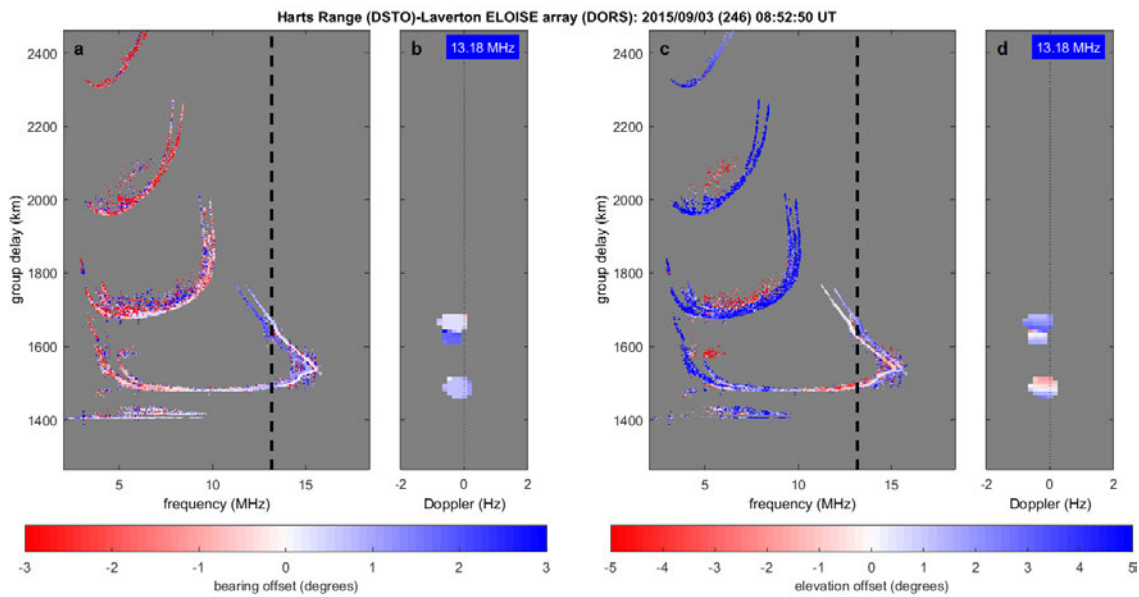
Figure 4.9. Estimates of bearing offset (panels a and b, measured clockwise with respect to the great circle bearing) and elevation offset (panels c and d, measured with respect to the 1-hop equivalent mirror elevation) for the same ionogram and CSF images as in **Figure 4.5**.
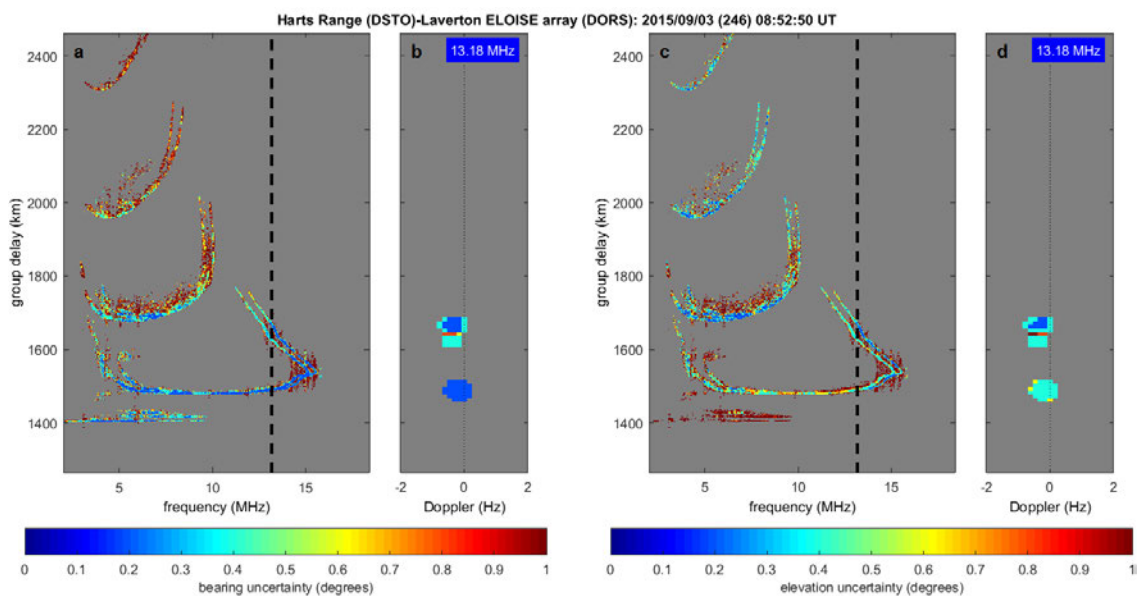
Figure 4.10. Estimates of bearing uncertainty (panels a and b) and elevation uncertainty (panels c and d) for the same ionogram and CSF images as in **Figure 4.5**.
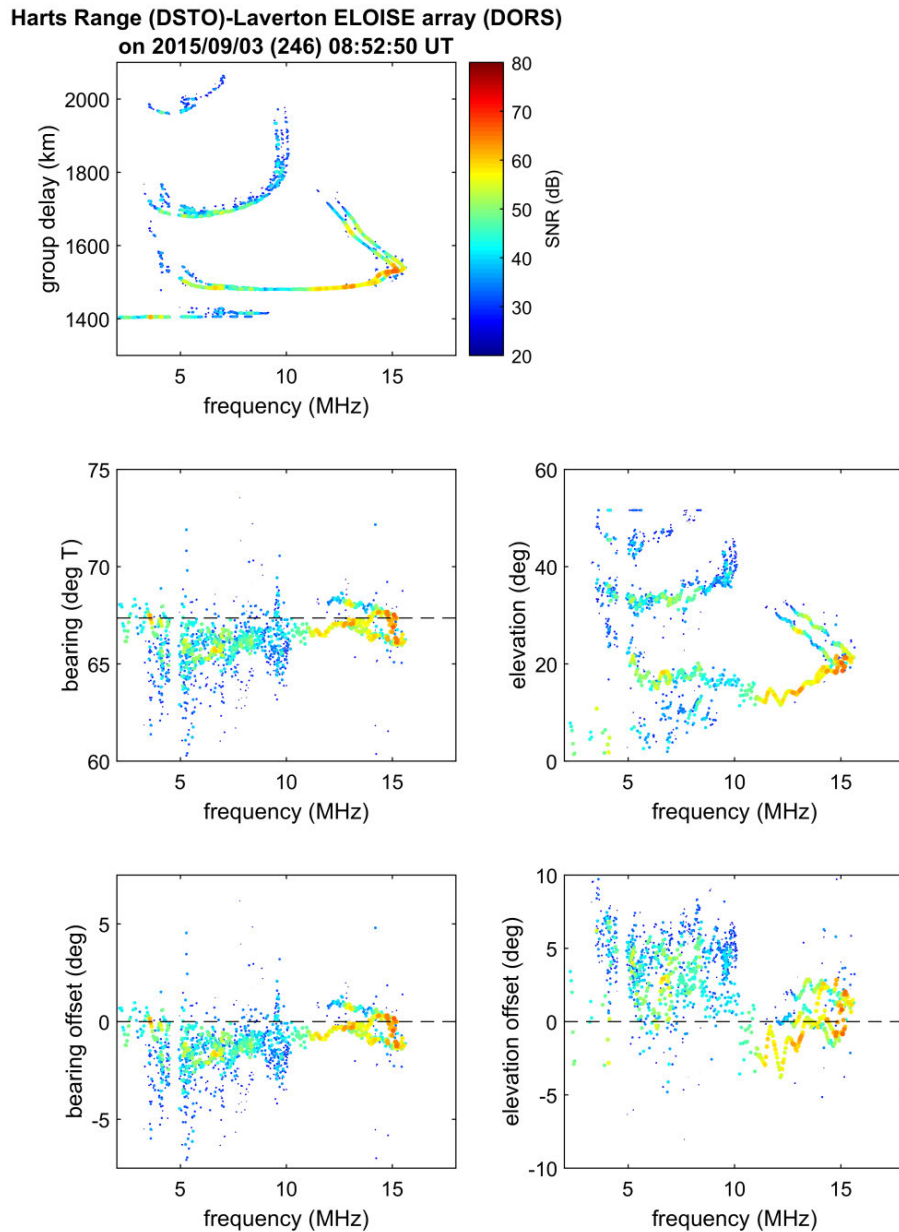
Figure 4.11. An alternative presentation of the ionogram from **Figure 4.5** and **Figure 4.9**, showing AoA estimates as a function of frequency. From top to bottom and left to right, the vertical axes represent group delay, bearing, elevation, bearing offset and elevation offset. The marker colours/sizes are fixed across all panels based on SNR.

For the 1-hop F2 low and high ray modes observed in the CSF image (as annotated in **Figure 4.5**), the magnitudes of the bearing offsets are mostly less than 1° in **Figure 4.9**; that is, they are received in line with the great circle bearing of 67 °T (measured clockwise from True North), as is typical under fairly benign ionospheric conditions.

However, note that the ordinary (O-mode) and extraordinary (X-mode) polarisation components of the 1-hop F2-high mode are subtly distinct, with the O-mode being more perturbed by the passage of the MSTID at this particular time (although the reverse is true at other times). While geomagnetic effects alone can influence the AoA, simulations by Dao et al. [2016] have indicated that these O/X deviations are likely to be insignificant compared to the potential effects of horizontal gradients and disturbances on longer oblique paths such as this.

The corresponding elevation offsets in **Figure 4.9** are slightly larger, and again there is an O/X-mode difference, but the perturbations are still relatively small on the whole (less than 3°). As the predicted geometry varies with delay, an artificial delay-dependent gradient is found in the elevation offsets of the modes in the lower-resolution CSF image (panel d); this is purely an artefact introduced by subtracting the equivalent elevation and does not relate to unresolved mode components. The multi-hop returns in the ionogram (panel c) also tend to have positive offsets, as the 1-hop predicted geometry is invalid and underestimates elevation for these multiple hops.

4.3   Ionogram peak detection, estimation and mode classification

In further offline processing, peaks are detected and fitted at each frequency in the thresholded ionograms and CSF dwells, giving estimates of power/SNR, frequency, group delay, Doppler, bearing and elevation, along with their uncertainties. A heuristic algorithm was developed and tuned specifically for the thresholded ELOISE data, which detects and fits peaks in the amplitude image (i.e. as a function of group delay and, for CSF, Doppler), then estimates peak AoA parameters across the other pages of the data cube. It consists of the following steps, illustrated in the four panels of **Figure 4.12** and **Figure 4.13** for ionogram and CSF images, respectively:

1   Impose a raised threshold on the SNR-scaled image, no less than 30 dB below the peak pixel at each frequency, to eliminate residual noise, RFI and side-lobe artefacts. The ionogram threshold is higher where RFI suppression has occurred.

2   Apply a small amount of convolutional smoothing over a 3×3 neighbourhood (panel a), to reduce the number of local maxima due to random amplitude fluctuations.

3   Perform morphological reconstruction on the filtered image (panel b) and use this as a cut-off for detecting prominent peaks that are at least 3 dB above all other local maxima in their connected region. Grey-scale dilation (over ±1 pixels in each fitted dimension) is initially applied to the image from step 2, to merge any peak regions that may have become detached due to thresholding, then this image is scaled down by 3 dB and further dilated, repeatedly, while being constrained by the same image prior to scaling. Local maxima in the smoothed image from step 2 that are greater than or equal to the output of the morphological reconstruction are declared to be candidate peaks. In the rare case that there are more than 16 separate candidates at any single frequency, those with the lowest peak amplitude values are discarded to bring the number down to 16; this enables the peak data to be stored in regular arrays.

4   For each candidate peak, identify the centroid of the connected local maximum region (panel c), noting that due to quantisation there may be a cluster of equal-valued pixels. A bounding box surrounding the local maximum is then expanded until it encapsulates at least a 3 dB roll-off on all sides; this condition may be relaxed for closely-spaced peaks or in cases where the effects of thresholding are prohibitive. Several checks are also carried out to flag poorly defined peak regions, including those that are excessively flat, anomalously shaped, inseparable from adjacent peaks, or truncated by the edges of the image.

5   Fit a quadratic curve or surface (i.e. a parabola or paraboloid, for ionogram and CSF, respectively), to dB-scaled pixels within the bounding box of each candidate peak (panel d). This is equivalent to a 1D or 2D Gaussian in linear amplitude units, being a function of group delay (ionogram) or Doppler and group delay (CSF). Note that while the peak detection, described above, is performed using morphological operations, the fitting itself is done on the smoothed-only image. A linear least-squares formulation is used here, returning the peak amplitude, position and half-widths at half-maximum (-3 dB). Highly asymmetric fits, for which the widths in each dimension differ by more than a

factor of two, in contrast to the shape of the point-spread function, are rejected at this point.

6   Calculate peak bearing and elevation estimates using a weighted mean of pixel values. A Gaussian weighting function is adopted, with half-width at half-maximum (0.5) equal to the half-width of the fitted peak from step 5. Standard deviations over the peak cluster are also calculated using an unbiased estimator of the weighted sample variance; this should not be confused with the transformed uncertainties in the AoA fits for each pixel (i.e. characterising the non-planarity across the array).
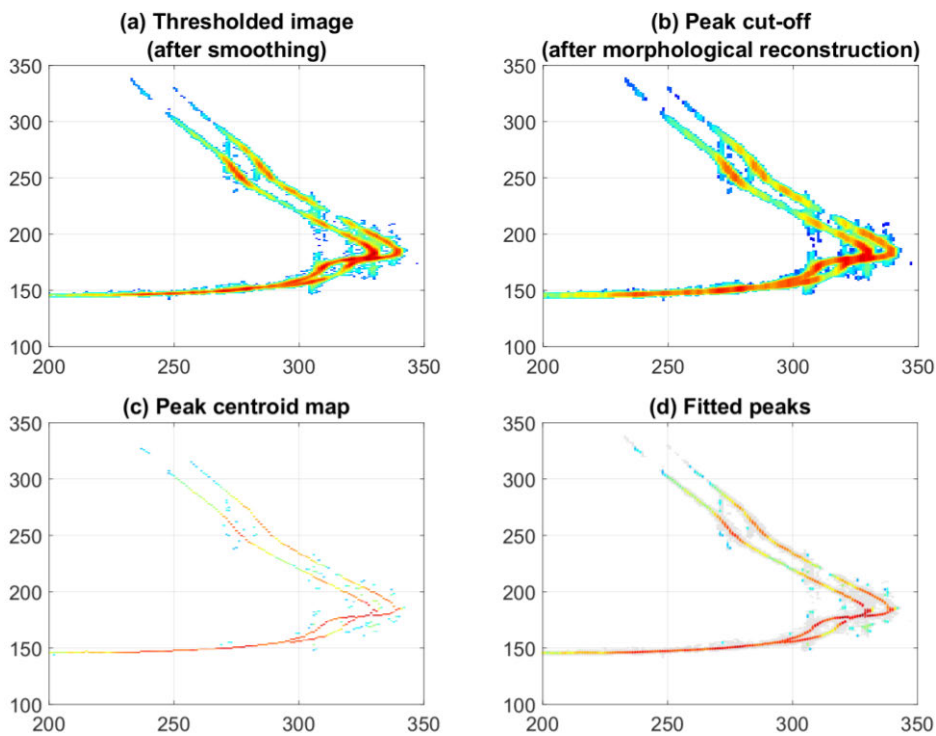
Figure 4.12. Ionogram peak detection and estimation process, illustrated using an enlarged portion of the left panel from **Figure 4.5**. The panels show (a) the thresholded image with smoothing applied (step 2); (b) the cut-off above which candidate peaks are detected (step 3); (c) the centroids of the candidate peaks (step 4); and (d) the fitted peaks (step 5) with small bars representing the 3 dB width and shaded areas indicating the original pixel positions from panel a. The axes labels represent pixel indices.
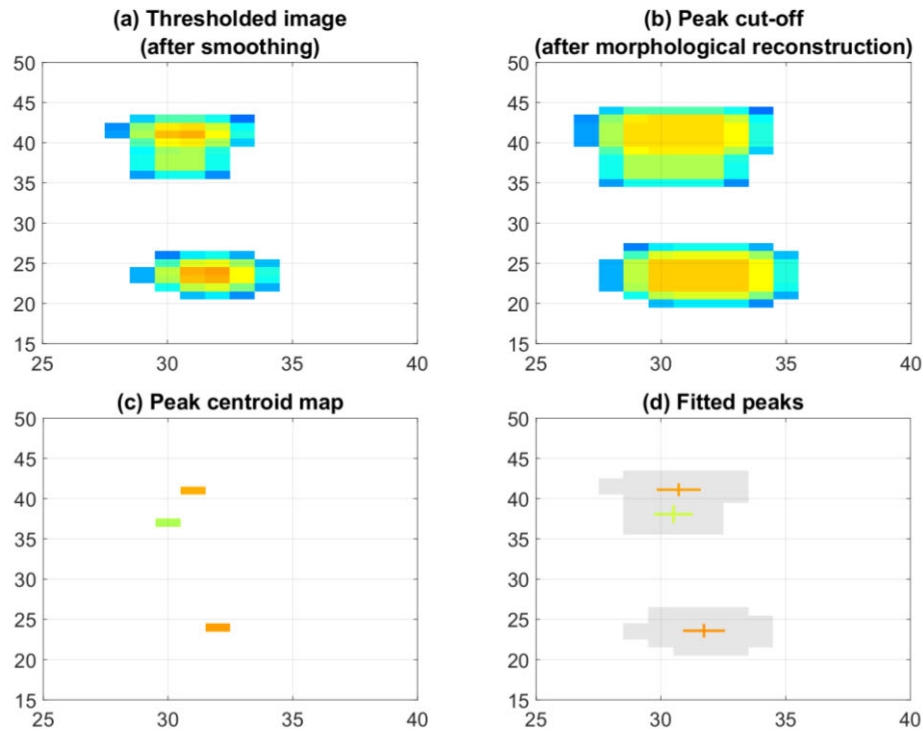
Figure 4.13. CSF peak detection and estimation process, illustrated using an enlarged portion of the right panel from **Figure 4.5**. The panels show the same processing steps as in **Figure 4.12**, although in this case the peaks are fitted in 2D.

Note that the predicted geometry is re-evaluated at the peak centre position to produce bearing and elevation offset estimates. This is important, as spectral leakage means the signal power is not confined to just one delay/Doppler cell, and taking the difference before peak fitting leads to the aforementioned gradient artefact in the elevation offsets of **Figure 4.9**. Nevertheless, the phase relationship between array elements and, hence, the AoA does not fundamentally change with spectral leakage into neighbouring cells, provided there remains sufficient SNR and separation from adjacent peaks.

Figure 4.14 overlays the peaks from **Figure 4.12** and **Figure 4.13** on the original ionogram and CSF images, demonstrating that they offer a good representation of the key propagation information, in a compact form well suited to subsequent manipulation. The three resolvable F2 modes in the CSF image, namely 1F2-low, 1F2-high(O) and 1F2-high(X), are all easily identified in this example.

The final stage in the peak processing is to automatically assign a mode classification, based on a parameterised model of the midpoint electron density profile. The ten-

parameter JORN ionospheric profile specification [Gardiner-Garden et al., 2018], which is routinely fitted to ionograms as part of DORS on-board processing [Heitmann & Gardiner-Garden, 2019], is used for this purpose; refer to **Appendix D** for a brief description of the JORN profile parameters and construction. The quasi-parabolic layer parameters are individually filtered over time, to help eliminate outliers, and synthetic traces are generated by analytic ray tracing through a spherically-symmetric ionosphere, as described in Heitmann & Gardiner-Garden [2019]. Doppler shifts can be calculated from a simple finite difference of the phase path. As the synthetic traces are a smoothed approximation (±7.5 min) of the propagation mode characteristics at any instant, limited to three basic layers (E, F1 and F2), and without any valleys or gradients, they only capture the primary mode structure. Nevertheless, as will be seen in **Chapter 7**, this is often sufficient to account for many larger-scale disturbances. A fixed group delay tolerance about the synthetic traces (±30 km for F2-low and ±60 km for F2-high), along with an additional frequency tolerance for the ionogram only (±1 MHz), is designed to draw in surrounding peaks into a single broad mode classification, without regard to micro-multipath or off-angle behaviour; this is illustrated in **Figure 4.15**, where 1-hop F2-low has been identified.
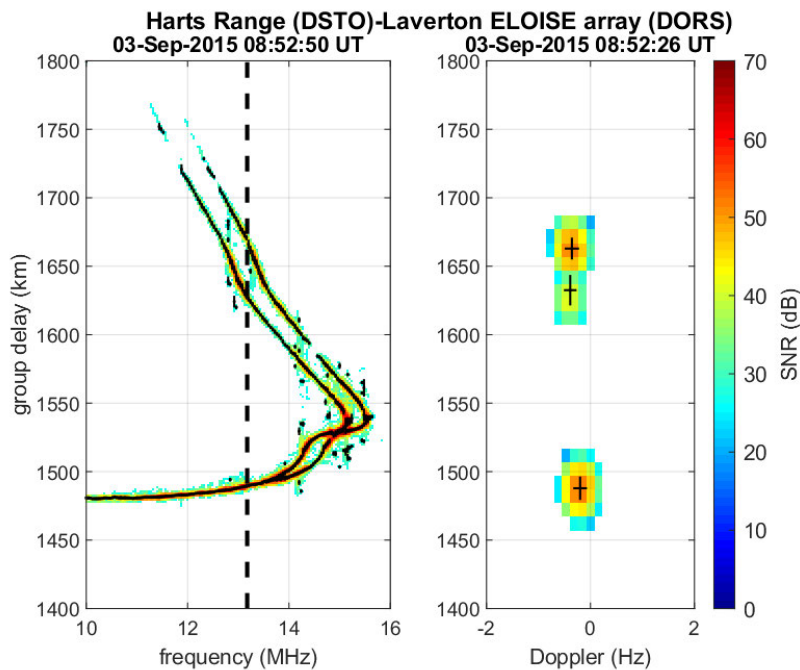
Figure 4.14. Fitted peaks from **Figure 4.12** and **Figure 4.13**, overlaid as black markers on the original ionogram (left) and CSF (right) images, scaled in dB SNR.
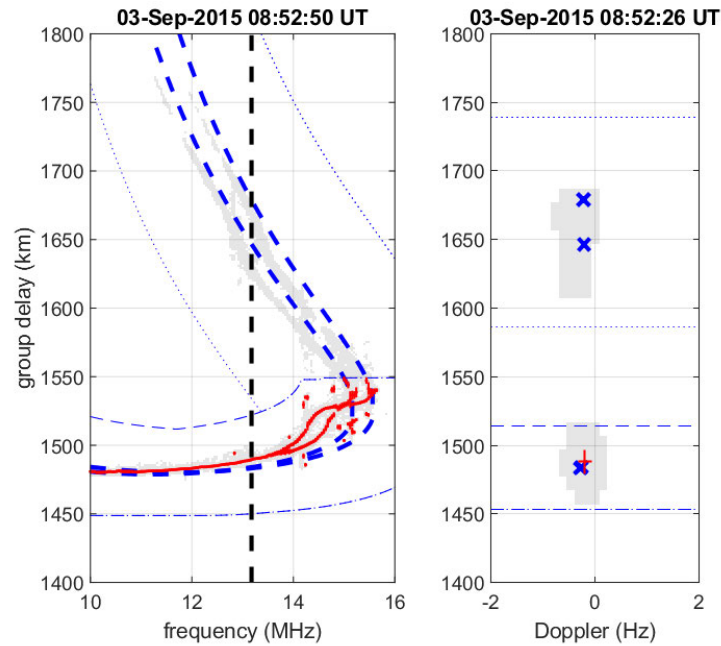
Figure 4.15. Fitted peaks from **Figure 4.14**, after 1-hop F2-low mode classification. Only those peaks that fall within a fixed delay tolerance of the synthetic F2 traces (thick blue dashed lines and crosses) are deemed to be F2 mode, and those below the nose of the trace are further classified as F2-low (red markers). For the ionogram panel (left), a frequency tolerance is also included. The effective classification region is bounded by the thin blue dashed lines, with the shaded areas indicating the original pixel positions.

It is important to note that the F2-low O/X components remain unseparated in **Figure 4.15**; this applies to all ionogram and CSF peak data regardless of their broad (no-field) mode classification. Without polarisation measurements, O/X separation is not a trivial task, and while some efforts were made to automate this process on amplitude-only ionograms, they were ultimately not reliable enough for wider use. All O/X classifications presented in this thesis are therefore manually determined by eye.

With regard to multi-hop modes, these image peaks are currently fitted, but do not progress any further in the analysis. Mode classification itself would not necessarily be difficult using the parameterised profile; however, interpreting the multi-hop HF observables in terms of localised perturbations in the (multiple) ionospheric reflection points is a non-trivial problem, given that AoA is only measured on the final hop. To overcome such ambiguities, a more generalised framework for relating and combining measurements across the ELOISE AoA network would be required, or AoA on both

transmit and receive could potentially be considered in future (e.g. with a multiple-input multiple-output or MIMO-like architecture [Frazer et al., 2010; Frazer et al., 2014]).

Although the analysis and interpretation of peak data will be reserved for subsequent chapters, at this point it is worth inspecting the vast quantities of peaks to gain a sense of the typical bearing and elevation uncertainties achieved by the ELOISE AoA system. Median values are plotted in **Figure 4.16** for each path as a function of SNR, using only the 1-hop F2-low returns (thus eliminating less stable propagation modes). The left
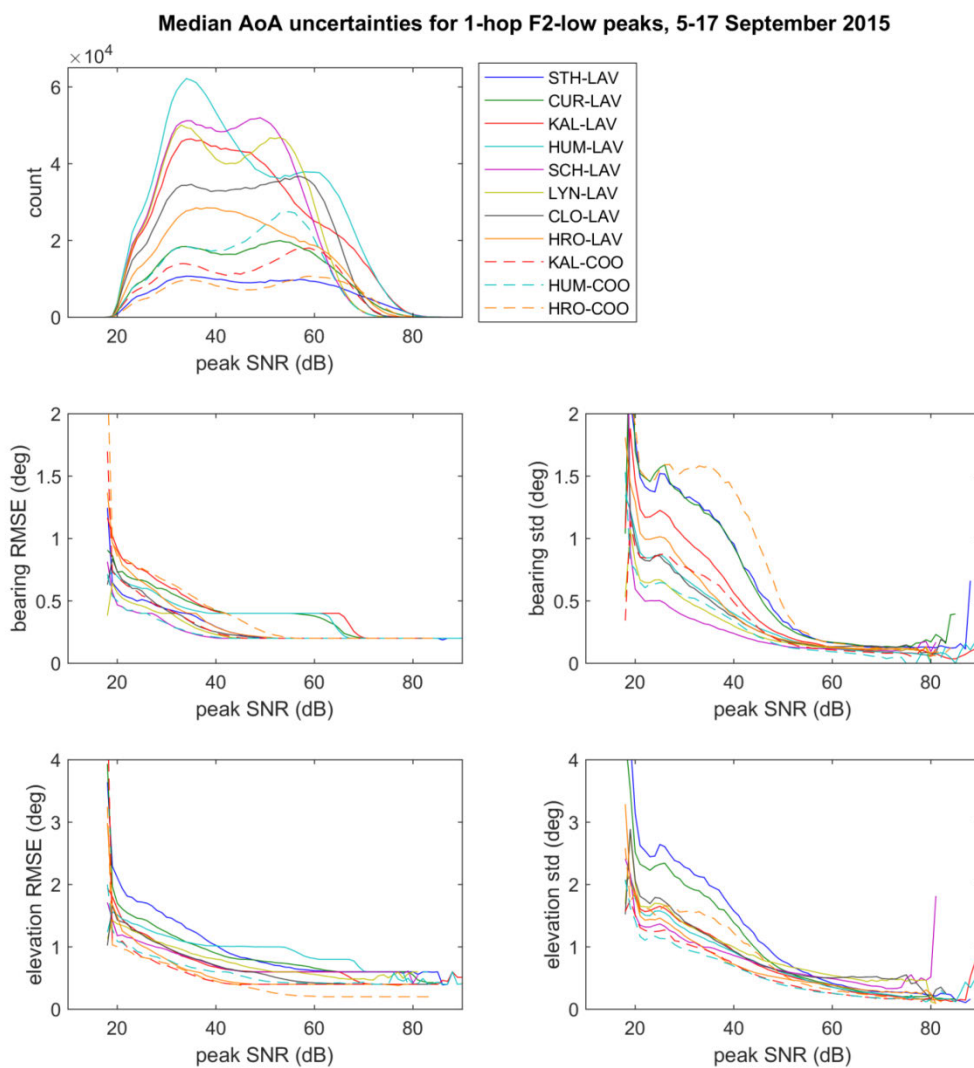
Figure 4.16. Median bearing and elevation uncertainties aggregated over all 1-hop F2-low ionogram peaks from 5–17 September 2015. The top panel shows a histogram of the number of contributing peaks, while the second and third rows represent bearing and elevation, respectively, in terms of both the RMSE of the planar AoA fit (left) and standard deviation over the image peak (right).

84

column shows the transformed RMSEs from the least-squares planar wavefront fits (i.e. a measure of spatial uncertainty across the array, taking a weighted average over the width of the peak), while the right column shows the weighted sample standard deviations of the AoA estimates themselves (i.e. a measure of how much the AoAs vary from pixel to pixel in the ionogram/CSF image).

Not surprisingly, the two uncertainty measures both exhibit the same diminishing relationship with increasing SNR. However, the standard deviations across image pixels (right column) tend to be larger at lower SNR, perhaps due to the merging of multiple marginally resolved modes in the delay dimension during the smoothing and morphological reconstruction steps. As the on-board AoA estimates are quantised in $0.2°$ steps, the uncertainties do contain some quantisation noise which prevents asymptotic decay to zero; other than this, the remainder at high SNR represents just calibration errors and intrinsic ionospheric effects. There is clearly some dependence on path geometry, yet it is not necessarily in the manner one might expect: at lower SNR in particular, the shorter paths (e.g. HRO–COO, STH–LAV and CUR–LAV) are among those with the highest uncertainties, even in elevation. It is believed this is because the SNR distributions (top panel) are shifted up slightly for short (less lossy) paths, in effect exposing more disturbed propagation at the lower SNRs.

4.4 Reflection point mapping

While variations in peak AoA parameters on a single path alone can reveal interesting patterns and periodicities, the results are inherently coupled to the path geometry and therefore of limited use in comparing different oblique paths across the ELOISE AoA network. An intuitive mapping scheme for transforming the observables (i.e. oblique frequency, delay, bearing and elevation) to reflecting layer parameters (i.e. equivalent vertical frequency, latitude, longitude and virtual height) was thus implemented to aid the interpretation of results, and is used widely in **Chapters 7** and **8**. It is founded on the equivalent geometry theorems of Breit & Tuve [1926] and Martyn [1935], along with the modified secant law [Davies, 1990; McNamara, 1991], extended to a tilted mirror surface. As the Doppler frequency shift is largely invariant to oblique path length for the same equivalent vertical frequency [Pickering, 1975; Davies, 1990, ch. 7.8.6], CSF Doppler measurements are not transformed as part of this mapping.

The reflecting mirror geometry is sketched in **Figure 4.17**, for a receiver-centred Cartesian coordinate system with the y-axis directed along the great circle path bearing and the z-axis being the local vertical. Simple 1-hop propagation geometry over a spherical Earth is assumed, using radius $R_E = 6376$ km, the best fit to the WGS-84 geoid over northern Australia [Gardiner-Garden, 2002]. Other known constants are the great circle path length R_0, receiver location $\mathbf{r}_R \equiv [x_R, y_R, z_R]^T = [0, 0, 0]^T$, and transmitter location $\mathbf{r}_T \equiv [x_T, y_T, z_T]^T = [0, \ R_E \sin(R_0/R_E), \ R_E \cos(R_0/R_E) - R_E]^T$.
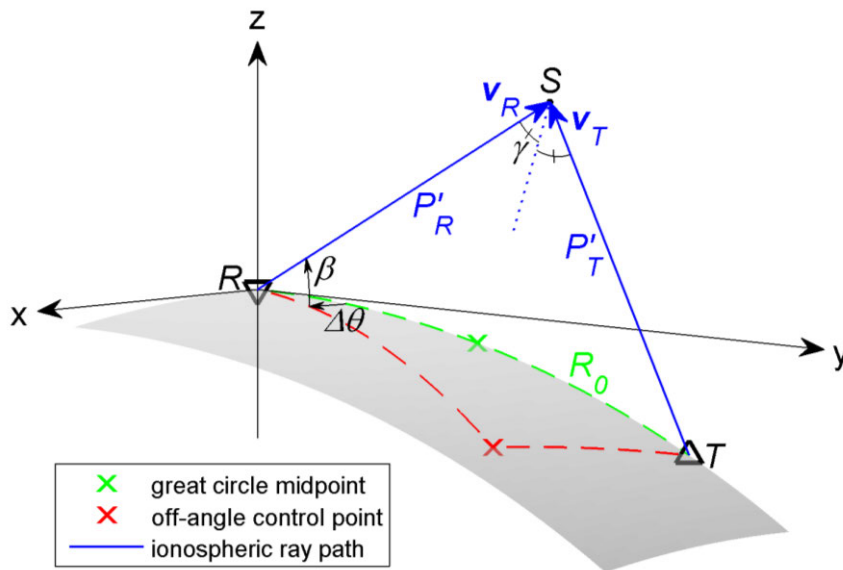
Figure 4.17. Diagram of receiver-centred Cartesian coordinate system used for reflection point mapping (not to scale). A perturbed (off-angle) path from transmitter "T" to receiver "R" via the ionospheric reflection point "S" is shown in blue; this is completely defined by the bearing offset $\Delta\theta$ and elevation β at the receiver, along with the total path length $P' = P'_R + P'_T$. A tilt angle for the reflecting surface is implied by the angle of incidence γ and the condition for specular reflection. The great circle path (dashed green line) and perturbed path projected on the surface of the Earth (dashed red line) are also shown.

The input (measured) parameters to the mapping process are the total path length (group delay) P', bearing offset $\Delta\theta \equiv \theta - \theta_0$ (for great circle path bearing θ_0), and elevation β. The output parameters are the latitude, longitude, virtual height of reflection, and

equivalent "vertical" frequency f_v, generalised to represent propagation normal to the reflecting surface (yet not necessarily at vertical incidence), and thus, approximate the plasma frequency. The unknown location of the reflection point, which can be related to the receiver and transmitter segment lengths P_R' and P_T', is given by

$$\boldsymbol{r}_S \equiv \begin{bmatrix} x_S \\ y_S \\ z_S \end{bmatrix} = P_R' \begin{bmatrix} \sin \Delta\theta \cos \beta \\ \cos \Delta\theta \cos \beta \\ \sin \beta \end{bmatrix}, \tag{4.11}$$

where

$$P' = P_R' + P_T', \tag{4.12}$$

$$P_R'^2 = x_S^2 + y_S^2 + z_S^2, \text{ and} \tag{4.13}$$

$$P_T'^2 = x_S^2 + (y_S - y_T)^2 + (z_S - z_T)^2. \tag{4.14}$$

Combining equations (4.12), (4.13) and (4.14) gives

$$P_T'^2 = P'^2 - 2P'P_R' + P_R'^2 = P_R'^2 - 2y_T y_S - 2z_T z_S + y_T^2 + z_T^2. \tag{4.15}$$

After substituting in y_S and z_S from equation (4.11) and rearranging, the following result is obtained for the receiver ($\overline{RS}$) segment length:

$$P_R' = \frac{P'^2 - y_T^2 - z_T^2}{2P' - 2y_T \cos \Delta\theta \cos \beta - 2z_T \sin \beta}. \tag{4.16}$$

The reflection point $\boldsymbol{r}_S$ and, if required, the transmitter ($\overline{TS}$) segment length P_T' can then be easily determined using equations (4.11) and (4.12), respectively. An additional transformation from $\boldsymbol{r}_S$ in Cartesian coordinates to latitude, longitude and virtual height yields a more natural representation for further analysis and interpretation. Note that geocentric latitude is converted to geodetic latitude using the WGS-84 eccentricity factor.

The final reflection parameter to calculate is f_v, which can be found by applying the modified secant law [Davies, 1990, ch. 6.6.2]:

$$f_o = \frac{k f_v}{\cos \gamma},$$
(4.17)

where f_o is the oblique frequency, γ is the angle of incidence to the reflecting surface, and

$$k = \begin{cases} 1, & R_0 < 500 \text{ km} \\ 1 + 0.000032 R_0, & R_0 \geq 500 \text{ km} \end{cases}$$
(4.18)

is an empirically derived range-dependent modification factor [Lynn, 2000]. Although k was strictly fitted assuming midpoint geometry, it is thought to be an acceptable first-order approximation for the more general case of off-angle reflection considered here. For specular reflection, γ can be calculated as half the angle between the two vectors $\boldsymbol{v}_R$ ($\overrightarrow{RS}$) and $\boldsymbol{v}_T$ ($\overrightarrow{TS}$):

$$\gamma = \tfrac{1}{2} \text{acos} \left(\frac{\boldsymbol{v}_R \cdot \boldsymbol{v}_T}{\|\boldsymbol{v}_R\| \|\boldsymbol{v}_T\|} \right).$$
(4.19)

An aggregated heat map of 1-hop F2-low ionospheric reflection points is shown in **Figure 4.18** for all Laverton and Coondambo AoA ionograms from 5–17 September 2015. Each fitted peak has been mapped using the transforms described above and combined to give a probability density distribution about each path midpoint. On the shortest four paths (900–1400 km), the distributions are fairly symmetric and unbiased about the great circle midpoints, as should be expected for observations that span many days and a multitude of ionospheric disturbances. Indeed, the only persistent gradient over these time scales should be associated with the equatorial anomaly, but this passes well to the north of all ELOISE AoA midpoints (at geomagnetic latitudes between 30 and 36 °S). On the longer paths (1500–2700 km), there is increasing evidence of biases in the down-range distances, caused by the overestimation of elevation noted earlier. For example, on Lynd River to Laverton, the centre of the distribution is displaced by approximately 250 km in **Figure 4.18**, well beyond the outer 120 km radius of the
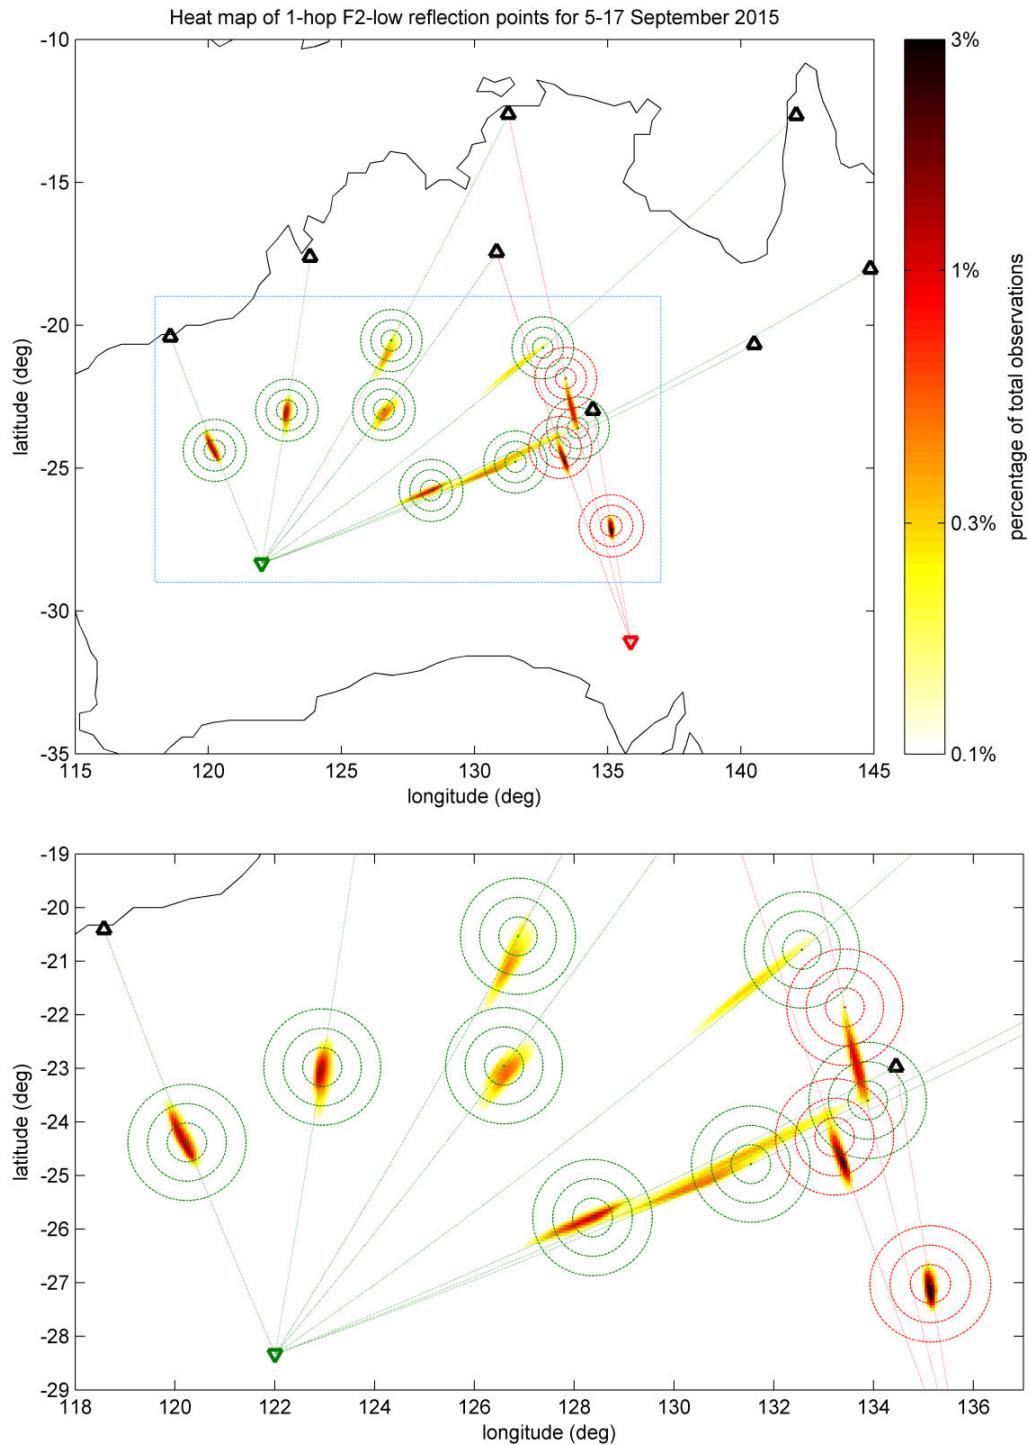
Figure 4.18. Heat maps showing the probability density distribution of 1-hop F2-low reflection points on all eleven ELOISE AoA paths from 5–17 September 2015. The bottom panel is an enlarged view of the central region marked by the light blue box in the top panel. Each oblique path has been independently normalised by the total number of fitted ionogram peaks across all frequencies. The bullseyes represent the great circle midpoints and 40, 80 and 120 km ground radii surrounding them.

relevant bullseye, and in fact closer to the Cloncurry to Laverton midpoint. Such a result is not surprising, given the synthesised uncertainty distributions in **Figure 4.8**, but the effect is much more pronounced in this geographic domain than in bearing/elevation space.

Perhaps more importantly, what the elongated shapes of the distributions in **Figure 4.18** do reveal is a strong sensitivity to cross-range reflection point perturbations, driven by variations in bearing rather than elevation. While this leaves potential blind spots in the ability to measure the full directional spectrum of ionospheric disturbances on a single oblique path, the orientation diversity of the ELOISE AoA network as a whole promises to provide a much more complete picture of the spatial structure of disturbances in the region. The scatter and motion of reflection points from one time epoch to the next in a geographic display such as this is also a valuable visual tool for detecting wave-like TID behaviour, moving horizontally over large distances. However, it is noted that reflection points do not actually track a fixed phase position along the disturbance, as the geometry for specular reflection changes as the ray is perturbed.

Finally, in addition to mapping the reflection points themselves, it is insightful to transform the AoA uncertainties (in degrees) to typical down-range and cross-range uncertainties (in km) for each path. This quantifies the relationship between path length and range uncertainty, including the unfavourable impact of elevation uncertainty on longer paths. The results, presented in **Table 4.1** for all 1-hop F2-low ionogram peaks, have been ordered from the shortest (top row) to the longest (bottom row) path, and list both the median values and median uncertainties (standard deviations) in group delay, bearing offset and elevation estimates, along with the transformed range uncertainties. Note that these are the same uncertainties that were previously plotted as a function of SNR in the right column of **Figure 4.16**.

Median uncertainties in group delay and bearing are reasonably consistent across all paths, at 1.2–1.6 km and 0.2–0.5°, respectively. These transform to down-range and cross-range uncertainties of 2–5 km and 3–5 km, respectively, in the reflection points. While the uncertainty in elevation shows a similar amount of variation, from 0.4–0.9° (roughly double the bearing uncertainty), this transforms to a much larger down-range uncertainty on the longest path (52 km) compared to the shortest path (9 km). Note that the low-elevation bias is an additional error on top of this.

Table 4.1. Median group delay, bearing offset and elevation estimates for 1-hop F2-low ionogram peaks observed on each of the eleven ELOISE AoA paths from 5–17 September 2015. The paths are ordered from shortest to longest. In each three-column set is the median value itself, the median of the sample standard deviations (calculated for each fitted peak, with the 3 dB widths in group delay scaled assuming a normal distribution), and this same result mapped via tilted mirror geometry to a down-range uncertainty (for group delay and elevation) or cross-range uncertainty (for bearing offset).

	Group delay ($\widehat{P}'$)			Bearing offset ($\Delta\widehat{\theta}$)			Elevation ($\widehat{\beta}$)		
	med.	std.	range unc.	med.	std.	range unc.	med.	std.	range unc.
HRO–COO (908 km)	1110 km	1.6 km	2 km	-0.2°	0.5°	4 km	31°	0.6°	9 km
SHD–LAV (944 km)	1125 km	1.4 km	2 km	-0.5°	0.4°	4 km	28°	0.9°	15 km
CUR–LAV (1202 km)	1365 km	1.4 km	3 km	0.1°	0.5°	5 km	23°	0.8°	19 km
HRO–LAV (1383 km)	1539 km	1.3 km	3 km	-0.2°	0.3°	4 km	21°	0.6°	17 km
KAL–LAV (1507 km)	1652 km	1.3 km	3 km	0.8°	0.4°	5 km	18°	0.7°	22 km
KAL–COO (1593 km)	1739 km	1.3 km	3 km	-0.1°	0.2°	3 km	18°	0.4°	15 km
HUM–LAV (1989 km)	2123 km	1.3 km	4 km	0.7°	0.3°	5 km	13°	0.7°	32 km
CLO–LAV (2055 km)	2188 km	1.3 km	5 km	-0.4°	0.3°	4 km	14°	0.8°	36 km
HUM–COO (2098 km)	2232 km	1.2 km	4 km	-0.1°	0.2°	3 km	13°	0.5°	22 km
LYN–LAV (2598 km)	2731 km	1.3 km	5 km	-0.5°	0.3°	5 km	11°	0.9°	51 km
SCH–LAV (2712 km)	2846 km	1.3 km	5 km	-0.4°	0.2°	5 km	8°	0.8°	52 km

As a guide, the bearing and elevation uncertainties for the CSF peaks are roughly half that of the ionogram peaks, in a median sense, but the group delay uncertainties are considerably larger, being inversely proportional to the processing bandwidth. Across all nine CSF paths, the standard deviations in delay, bearing and elevation were in the range of 8–15 km, 0.1–0.2° and 0.2–0.4°, respectively. The lower AoA uncertainties are likely to result from the longer CIT and higher SNR of the CSF peaks. The standard deviation in Doppler was a consistent 0.12 Hz (approximately three-quarters of a Doppler cell) on all CSF paths.

4.5 Ionospheric tilt estimation

Embedded in each solution for the off-angle path geometry, under the condition for specular reflection, is information about tilts in the reflecting surface. If the ionosphere is modelled as a spherical mirror with radius R_i that is non-concentric with the spherical Earth [Folkestad, 1968; Dyson & Bennett, 1992], then any reflection path defines a point $\boldsymbol{r}_S$ on the spherical mirror and a unit normal vector to the surface at this point (directed outwards):

$$\widehat{\boldsymbol{v}}_{norm} = \frac{\boldsymbol{v}_R + \boldsymbol{v}_T}{\|\boldsymbol{v}_R + \boldsymbol{v}_T\|}. \tag{4.20}$$

Given only a single point and its normal vector, the unknown parameters that define the sphere, being the centre position $\boldsymbol{r}_0 \equiv [x_0, y_0, z_0]^T$ and radius R_i, are underdetermined. To address this, the coordinate system from **Figure 4.17** is initially rotated about the Earth centre $[0, 0, -R_E]^T$ in the y-z ($x = 0$) plane, to reposition the great circle midpoint at the origin (i.e. such that the new z^*-axis is now the local vertical at the midpoint). Displacements to the centre of the spherical mirror are then constrained to the local horizontal plane, with $z_0^* = -R_E$ (fixed); in other words, restricting the offset to the cross-range (x^*) and down-range (y^*) directions only, with a varying radius controlling the reflection height. Note that an asterisk denotes rotated coordinates in the midpoint-centred system.

Starting with the unperturbed case (i.e. $x_0^*, y_0^* = 0$), the solution for the rotated but constrained centre position $\boldsymbol{r}_0^* = [x_0^*, y_0^*, -R_E]^T$ and its corresponding mirror radius is

found in an iterative fashion. At each iteration, the radius to each reflection point $\boldsymbol{r}_S^* = [x_S^*, y_S^*, z_S^*]^T$ is first updated using the spherical formula:

$$R_i = \|\boldsymbol{r}_S^* - \boldsymbol{r}_0^*\| = \sqrt{(x_S^* - x_0^*)^2 + (y_S^* - y_0^*)^2 + (z_S^* + R_E)^2}. \tag{4.21}$$

This is followed by an update to the centre position of the mirror, based on knowledge of the unit normal:

$$x_0^* = x_S^* - R_i(\hat{\boldsymbol{v}}_{norm}^*)_x, \text{ and} \tag{4.22}$$

$$y_0^* = y_S^* - R_i(\hat{\boldsymbol{v}}_{norm}^*)_y. \tag{4.23}$$

Trimmed sample means and standard deviations for x_0^* and y_0^* are computed from all available 1-hop F2 reflection points in a given sounding, weighted by SNR. This draws out the large-scale spatial gradients from the noise and smaller-scale (frequency- and height-dependent) components, and becomes the starting point for the next iteration. The trimmed statistics exclude the lower and upper deciles of the centre estimates in each dimension to provide robustness against outliers. Upon convergence, the midpoint height and horizontal gradients can be calculated directly from the mean x_0^*, y_0^* and R_i estimates using the spherical formula and its partial derivatives:

$$H = z_S^* \Big|_{x_S^*, y_S^* = 0} = \sqrt{R_i^2 - x_0^{*2} - y_0^{*2}} - R_E, \tag{4.24}$$

$$H_{x^*} = \frac{\partial z_S^*}{\partial x^*}\Big|_{x_S^*, y_S^* = 0} = \frac{x_0^*}{\sqrt{R_i^2 - x_0^{*2} - y_0^{*2}}}, \text{ and} \tag{4.25}$$

$$H_{y^*} = \frac{\partial z_S^*}{\partial y^*}\Big|_{x_S^*, y_S^* = 0} = \frac{y_0^*}{\sqrt{R_i^2 - x_0^{*2} - y_0^{*2}}}. \tag{4.26}$$

A further coordinate rotation in the horizontal plane aligns the x and y axes with the zonal (eastward) and meridional (northward) directions at the midpoint:

$$H_{x(E)} = \cos(\theta_m)\, H_{x^*} + \sin(\theta_m)\, H_{y^*}, \text{ and} \qquad (4.27)$$

$$H_{y(N)} = -\sin(\theta_m)\, H_{x^*} + \cos(\theta_m)\, H_{y^*}, \qquad (4.28)$$

where θ_m is the great circle path bearing at the midpoint. Zonal and meridional tilt angles are given by the arctangent of these horizontal gradients.

Despite the simplistic propagation assumptions involved, the resultant tilt estimates are found to sensibly capture the characteristic negative and positive zonal gradients present around the dawn and dusk terminators, respectively, as well as significant day-to-day variations over periods of an hour or more (see **Chapter 7**). Thus, with appropriate smoothing, these large-scale background gradients provide a means to remove part of the AoA variability in a physically meaningful way, to better reveal the medium-scale (<1 hour) components that are typically the subject of MSTID studies. There remains an outstanding question over the influence of geomagnetic effects; for example, Dao et al. [2016] demonstrate how field effects can masquerade as tilts. However, given the scale of variations in the geomagnetic field, this is expected to present as a small bias rather than additional variance in the tilt estimates.

4.6   Chapter summary

- On-board signal processing for the ELOISE AoA system produces high quality ionogram and delay-Doppler (CSF) images, fits the midpoint profile parameters, evaluates clear channel availability, and performs robust AoA estimation for the dominant propagation mode in each frequency/Doppler and group delay cell. These steps were designed to run in real time to support frequency-agile CSF scheduling and remote monitoring of the data.

- The AoA estimation algorithm jointly solves for the bearing and elevation angles over both arms of the 2D array, using an interferometric least-squares technique to deliver both estimates and their uncertainties. Ambiguities due to grating lobes are resolved using predicted path geometries.

- Some of the challenges in estimating elevation on very long oblique paths have been identified, using both sample observations and simulations. Resolution limits of conventional processing, and the impacts of unresolved multimode have also been discussed, with high-resolution spectral analysis proposed as one possible solution and demonstrated with an auto-regressive data extrapolation algorithm.

- In further offline analysis, peaks are detected and fitted in the AoA images, and 1-hop F2 returns are classified using the midpoint profile parameters. The AoA peak data can then be transformed via tilted mirror geometry to give reflection points and a characterisation of large-scale ionospheric gradients, which are used to compare AoA estimates across different oblique paths in subsequent chapters. Relationships between the path length and reflection point uncertainties have been explored.

5 Understanding preliminary results and instrumental limitations

5.1 Overview of preliminary results and analysis

The ELOISE experiment represented the first extensive testing for the angle-of-arrival (AoA) receiving system, having been designed specifically for this campaign. Between late July 2015 and mid-January 2016, approximately 300,000 AoA ionograms and 550,000 channel scattering function (CSF) dwells were collected across the two arrays, cataloguing a diverse set of ionospheric conditions and disturbances. This included a moderate number of geomagnetically disturbed days, with Kp indices frequently in the 4–7 (Active to Storm) range on a scale of 0–9. Solar activity was at a moderate level throughout the collection period, with a smoothed sunspot index R_{12} (version 2) of 65.9 for the core month of September 2015 [Space Weather Services, 2018].

Following completion of the experiment, one of the first activities was to document, day by day, the quality of the data collected and ionospheric conditions (variability) observed. This initiated a considerable effort to better understand the instrumental effects that were manifested in the data and correct or account for them where possible. While on-board calibration focused on empirical corrections only, using very limited data, the hope was that a more physical basis could be established for the AoA discrepancies, particularly at low elevation angles, although ultimately the results were largely inconclusive.

This chapter covers several distinct subjects, each representing preliminary steps that were taken to interpret the ELOISE AoA data in terms of ionospheric disturbance signatures and help separate the geophysical effects of interest from the instrumental limitations. The material presented here therefore acts to bridge the gap between the pre-trial development (in **Chapters 3** and **4**) and the post-trial analysis (in **Chapters 6**, **7** and **8**). The first section presents examples of typical ionospheric variability identified

in the preliminary AoA data analysis, to broadly establish the applicability of the ELOISE observations to the F2 variability studies in later chapters. Having identified residual AoA biases arising with just the preliminary (on-board) calibration, primarily the overestimation of elevation on long paths, this chapter then models and quantifies two lower-order effects that have the potential to impact the results; namely, tropospheric refraction and array mutual coupling. As the modelling alone is not found to satisfactorily account for the biases, empirical (path-dependent) bearing and elevation corrections are instead estimated from aggregate ionospheric returns and applied to all subsequent results.

5.2   Examples of disturbances in angle-of-arrival data

Ionospheric variability manifests in many different forms in ground-based ionosonde observations. Among the common components of variability identified in the ELOISE AoA data set were:

- TID-like signatures, including kinks in the trace, "breathing" modes, and satellite traces, depending on the amplitude and horizontal wavelength of the disturbance. These signatures often appear in the 2-hop trace first, then subsequently the 1-hop trace, indicative of horizontally propagating structure. A more detailed description of such effects can be found later in **Section 8.1**.

- Night-time mid-latitude spread-F conditions, including the characteristic curvature in the lower edge of the CSF delay-Doppler response, related to the horizontal velocity of the reflecting surface [Lynn, 2008].

- Geomagnetic storm conditions, with continuous observations over multiple days tracking the ionospheric response across the commencement, main and recovery phases.

- The presence of sporadic-E layers, including instances of discrete traces caused by off-angle returns from patchy structure. Although the equinox period of ELOISE represents a seasonal low in sporadic-E occurrence [Whitehead, 1989], Es layers were nevertheless still observed, with blanketing Es on some nights.

- The pre-dawn collapse [Rishbeth, 1988] and midday bite-out [Kohl & King, 1967; Saryo et al., 1989; Lynn et al., 2014; 2016] of the F2 critical frequency, associated with layer compression/rarefaction and field-aligned plasma transport caused by changes in thermospheric winds.

This section will present typical examples from the first four of these categories, which most benefit from having access to AoA information.

A sample ELOISE ionogram from the Laverton array is shown in **Figure 5.1**, with colours scaled according to received power (left), bearing offset (centre), and elevation offset (right). The path from Cloncurry to Laverton spans 2055 km with a great circle bearing of 70 °T. The contorted shape ("kink") identified on the 1-hop F2-low trace corresponds to multiple simultaneous propagation paths through a localised electron density perturbation, probably associated with a medium-scale TID; such disturbances
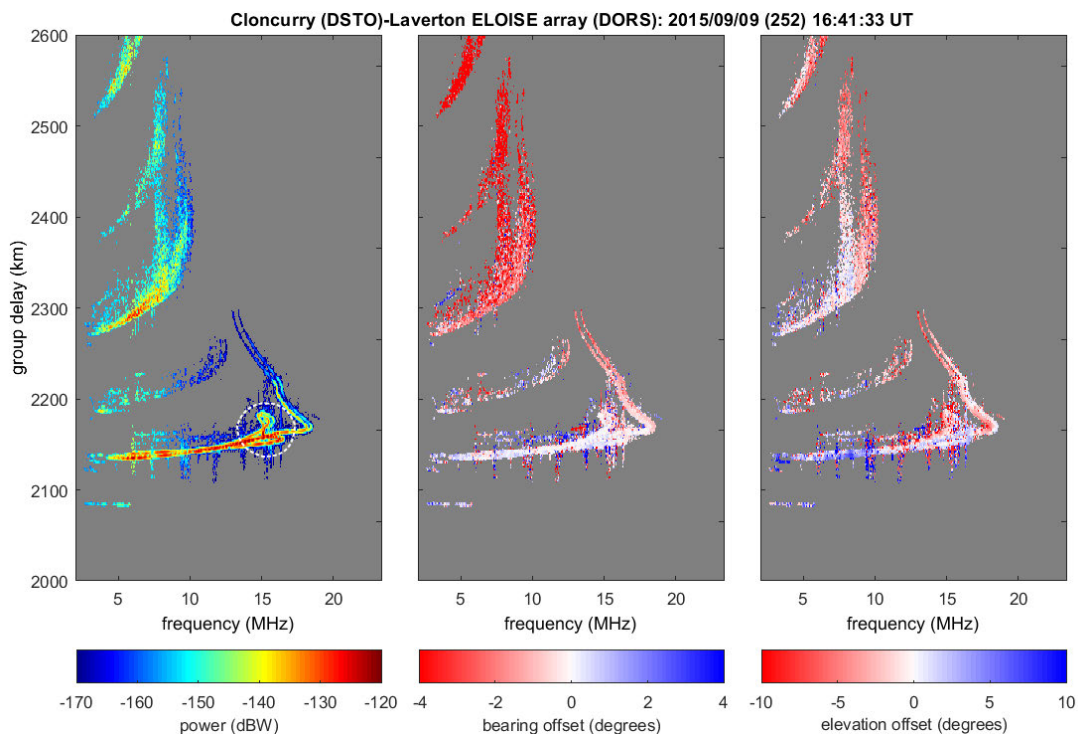
Figure 5.1. Sample of an ELOISE ionogram scaled in terms of power (left panel), bearing offset (centre panel), and elevation offset (right panel). The local time at Laverton was 12:41am (UT+8 hours). The dashed white line in the left panel identifies a particularly prominent trace "kink" on the 1-hop F2-low trace, thought to be caused by the passage of a medium-scale TID.

are discussed further in **Chapter 8**. This particular case occurred in the middle of the night, and over the course of about an hour, the kink feature progressed all the way down the high-ray and along the low-ray F2 trace. Despite the clear multiplicity of propagation paths in group delay at around 15 MHz, the AoA offsets barely register the feature, due to it having a much smaller spatial scale than the length of the oblique path.

Another example of a disturbed ionogram is shown in **Figure 5.2**, representing relatively degraded night-time propagation conditions (approximately 3 hours after sunset at F2 heights) on the oblique path from Humpty Doo (near Darwin) to Laverton, with ground range 1989 km and great circle bearing 31 °T. The off-angle ("satellite") trace that appears across most frequencies at group delays between 2100 and 2200 km, with a bearing offset of almost 20° in parts, is a signature of large-scale TIDs frequently seen in the post-sunset ionosphere at mid-latitudes. Elevation offsets are mostly negative along the satellite trace, as the equivalent great-circle path with the same delay predicts a larger elevation than the true off-angle geometry implies.
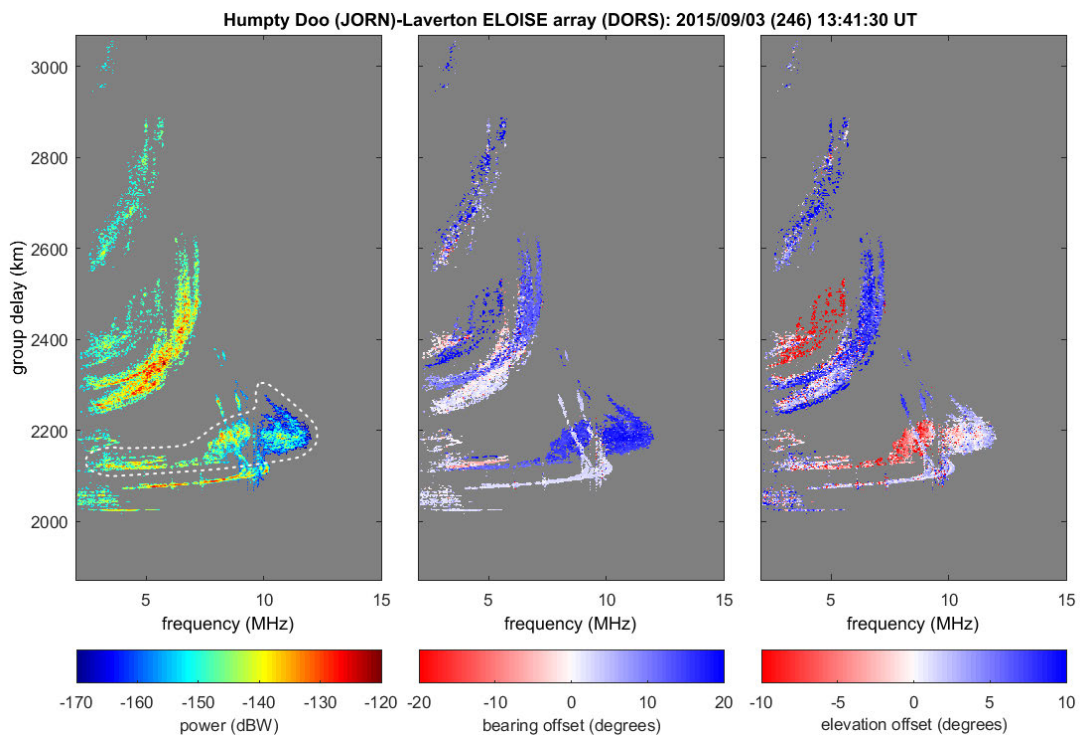
Figure 5.2. Another disturbed ELOISE ionogram, presented in the same format as **Figure 5.1**, showing the off-angle signatures of larger-scale ionospheric structure. The local time at Laverton was 9:41pm. The dashed white line in the left panel identifies an additional ("satellite") 1-hop F2 trace, characterised by extreme deviations in bearing.

Large-scale structures such as in this example are often found to be a precursor to spread-F irregularities [Bowman, 1981; Bowman & Monro, 1988; Bowman, 1991b; Lynn et al., 2013], which pervade the post-midnight ionosphere at mid-latitudes, particularly during low-to-moderate solar activity [Bowman, 1960]. Over the surrounding 45 minute period, shown in **Figure 5.3**, the additional off-angle traces move outwards and inwards in group delay as their bearing offsets shift from negative (red) to positive (blue) values. At times, these traces partially overlap in the ionograms (e.g. 1337 UT), and the bearing offset represents only the stronger of the two returns.
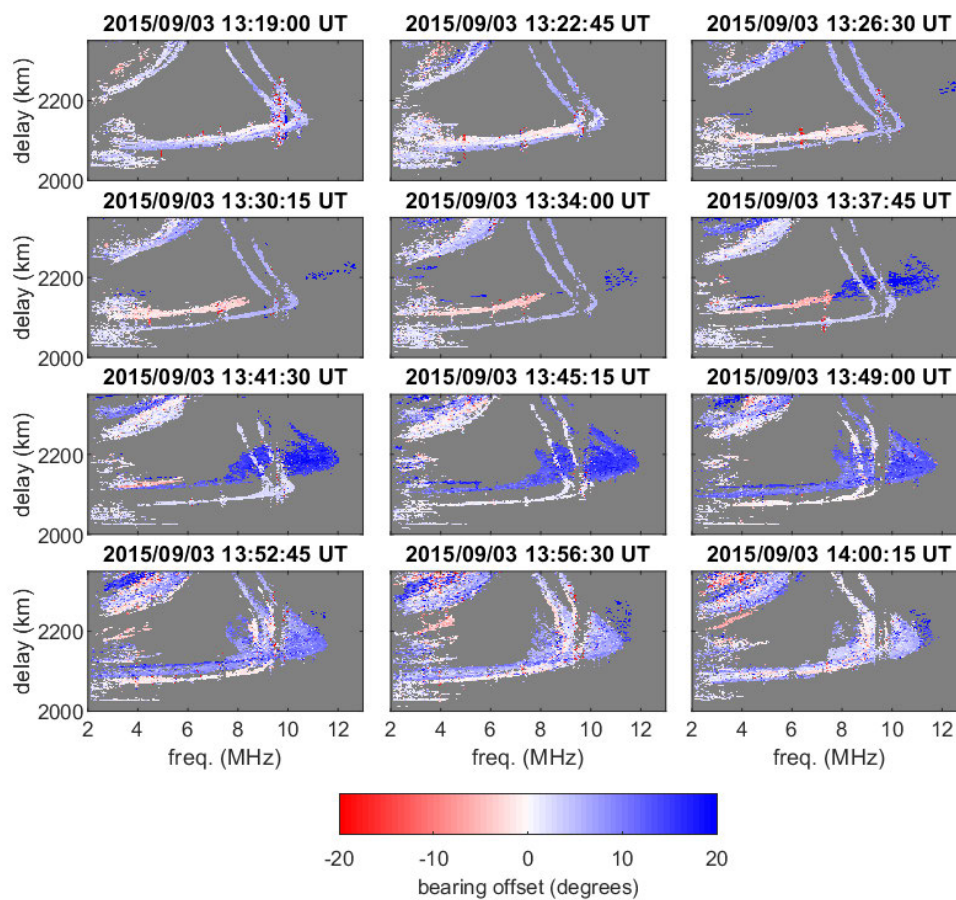
Figure 5.3. A time sequence of bearing offset images, measured before and after the ionogram in **Figure 5.2**. From left to right and top to bottom, the panels show one "satellite" trace breaking away from the unperturbed (direct) path and moving outwards in group delay (1319–1341 UT), followed by another trace, from the opposite direction, moving inwards in delay and merging with the unperturbed path (1337–1400 UT).

An enlarged portion of the same ionogram from **Figure 5.2** is shown in **Figure 5.4** alongside a neighbouring CSF dwell at 5.7 MHz. This frequency corresponds to the dashed vertical line on the ionogram, which captures the lower frequency tail of the off-angle structure identified in **Figure 5.2**. The automatic peak detection and fitting algorithm, operating across each page of the AoA data cube, has identified three 1-hop F2 returns in the CSF image, including the unperturbed shorter path (at 0.0 Hz Doppler and 2076 km group delay), and two off-angle longer paths: one with its reflection point advancing (at +0.8 Hz and 2123 km) and one receding (at -0.7 Hz and 2134 km). Their great circle bearing offsets were +10° and -7°, respectively, indicative of a TID structure propagating transversely (roughly north-westward) to the direction of the HF reception (north-eastward). The time evolution of ionograms in **Figure 5.3** supports this interpretation. Furthermore, it is consistent with the systematic westward and equatorward propagation directions reported elsewhere for night-time TIDs at mid-latitudes (e.g. Shiokawa et al. [2009]). The 2-hop F2 mode structure, above 2300 km group delay, is considerably more complicated, although also contains both advancing and receding components in the CSF dwell.
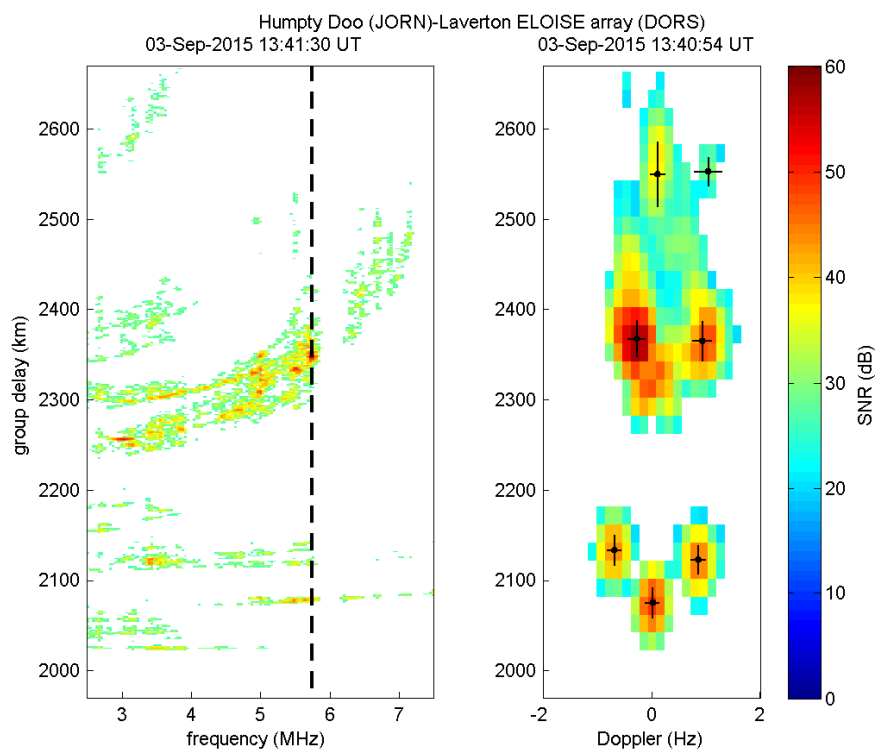
Figure 5.4. Enlarged portion of an ELOISE ionogram (left panel, rescaled from **Figure 5.2**) alongside the matching CSF dwell at 5.7 MHz (right panel, with fitted peaks).

A more degraded ionogram from later on the same night, but a different path, is shown in **Figure 5.5**. In this case, rather than there being a discrete off-angle trace, the returns are spread in delay, and contain both advancing and receding spread Doppler components of increasing magnitude at greater delays. Nevertheless, as in **Figure 5.4**, the advancing (positive Doppler) components tend to have a positive bearing offset, while the receding (negative Doppler) components have a negative bearing offset, indicative of a similar horizontal drift motion.
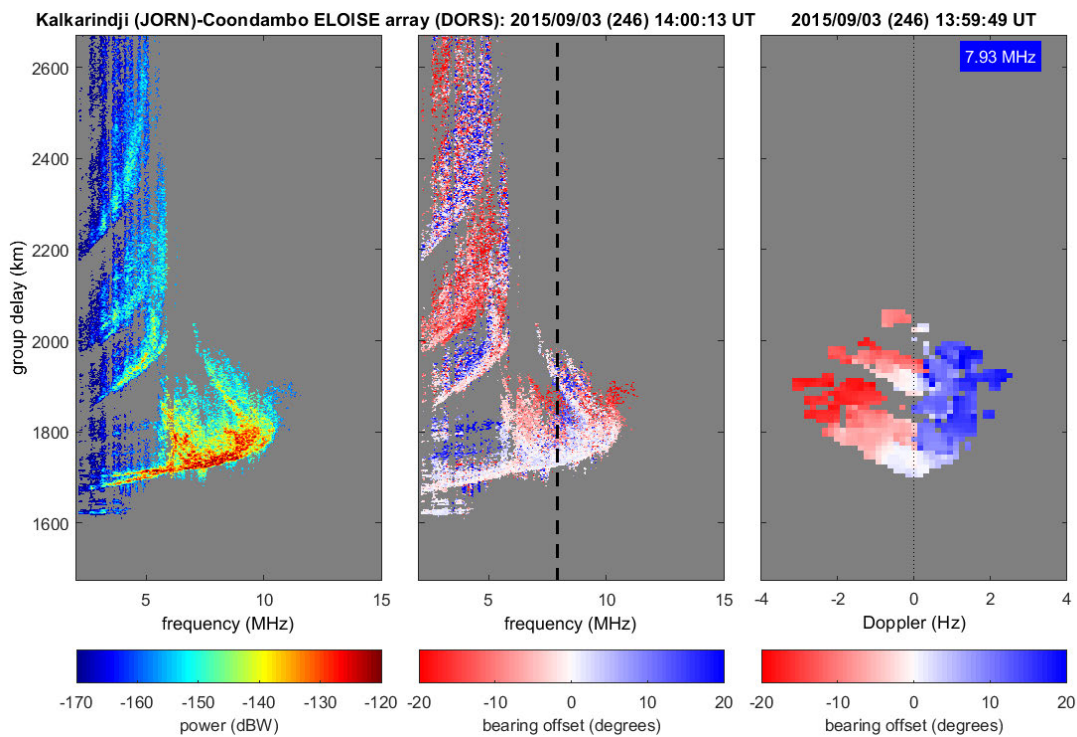
Figure 5.5. A night-time spread-F example, showing the ionogram (scaled in terms of power, left panel, and bearing offset, centre panel) alongside its corresponding CSF dwell (scaled in terms of bearing offset, right panel).

It is worth noting that the disturbances described above are not directly attributable to any solar or geomagnetic events, as the observations followed several days of quiet conditions. However, enhanced geomagnetic activity was encountered later in the month, caused by a high-speed solar wind stream from a coronal hole. Significant depressions in F2 propagation support were observed on 10–11 September 2015, in response to the recovery phase of the storm [Gardiner-Garden et al., 2019]. An example

of the impact on ELOISE AoA observations is shown in **Figure 5.6**. The F2 returns, starting at ~1700 km delay, are severely constrained in frequency, and unusually diffuse for this time of day (early afternoon). There is a distinct negative-to-positive (red-to-blue) transition in the elevation offset on either side of the F2 high-ray traces, pointing to a down-range spatial gradient in the F2 critical frequency (foF2); this is consistent with the interpretation of mid-latitude frequency spread presented by Clarke [1972], Bowman et al. [1988], and Bowman [1991a].
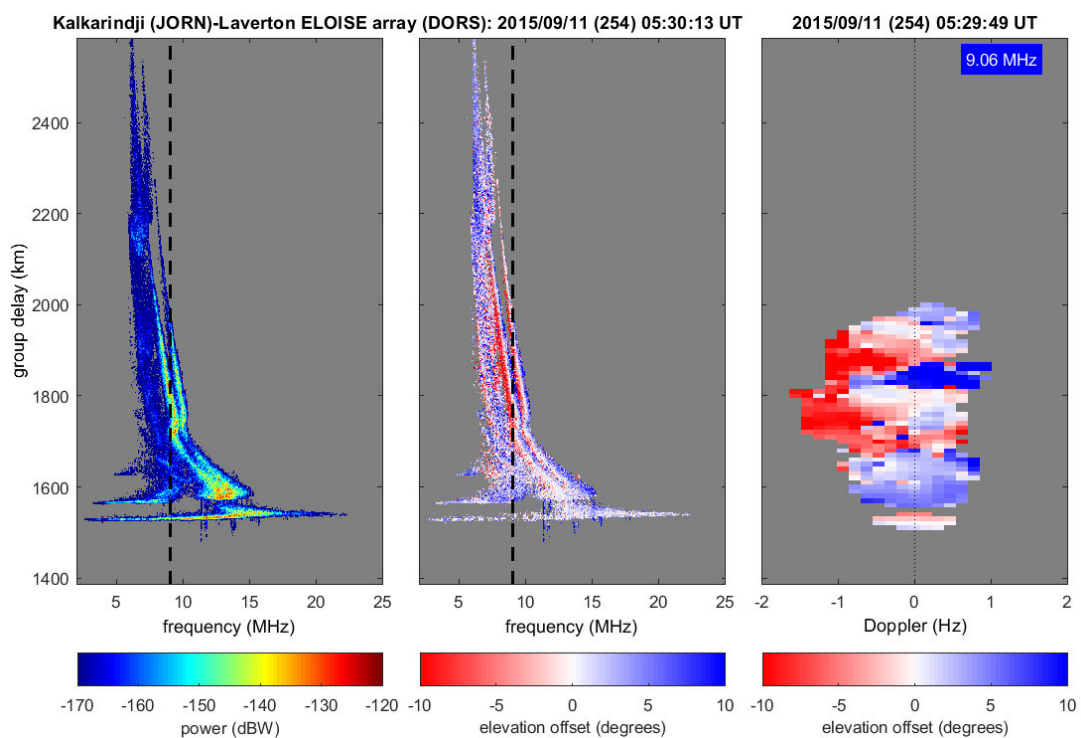
Figure 5.6. A daytime example with severely depressed F2 ionisation and degraded propagation characteristics, associated with increased levels of geomagnetic activity. The three panels show the ionogram (scaled in terms of power, left panel, and elevation offset, centre panel) alongside its corresponding CSF dwell (scaled in terms of elevation offset, right panel).

A final ionogram example in **Figure 5.7** presents the off-angle characteristics of patchy night-time sporadic-E (Es) structure. At the time, the Es layer was blanketing the F2 layer, which was experiencing a very depressed critical frequency (foF2). At least three near-horizontal traces, indicative of returns from a thin layer at multiple reflection

points, are marginally separated in group delay, yet have strikingly different bearing and elevation angles. The longest propagation path is the most perturbed in AoA, meaning its reflection point is furthest from the great circle midpoint. Note that this example is quite similar to the ionogram in Figure 10 of Ayliffe et al. [2019], without AoA, and highlights the difficulties in interpreting such data using delay measurements only.
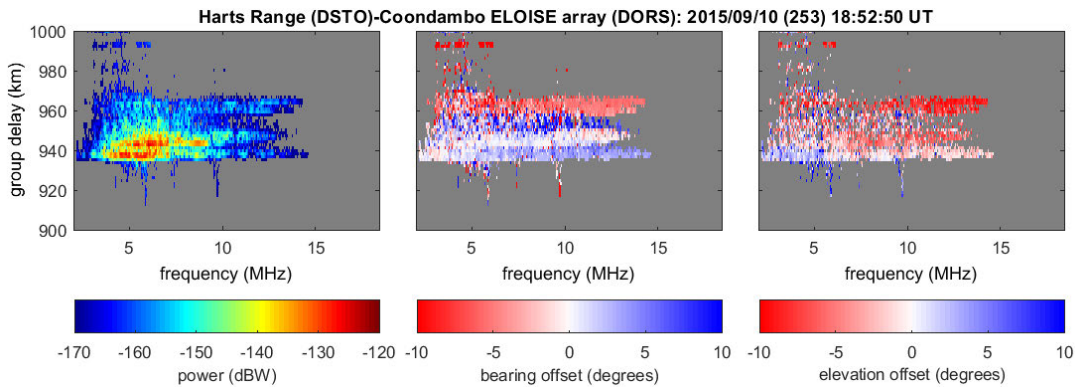
Figure 5.7. A night-time example containing multiple discrete sporadic-E (Es) traces, each with their own distinct off-angle direction. The three panels show the enlarged E-region of the ionogram scaled in terms of power (left panel), bearing offset (middle panel), and elevation offset (right panel).

5.3   Tropospheric refraction

The focus for the remainder of this chapter now shifts to understanding instrumental effects and limitations. Although far less significant than in the ionosphere, refraction in the troposphere can influence HF propagation in a similar manner, whereby a negative refractivity gradient as a function of height causes the ray trajectory to be bent back towards the Earth [Altshuler, 1998; Doerry, 2013]. This has the potential to increase the apparent elevation angle at low elevations, albeit by only a small amount. This section seeks to identify whether the impact should be measurable in the ELOISE AoA data.

While the effects of tropospheric refraction have been well studied at VHF, UHF and microwave frequencies, most recently in the context of satellite communications, the research activity has been comparatively less active in the HF band. In theory, however, the HF implications should be the same; it is the ability to measure them that is more

difficult. A good summary of tropospheric refraction in the HF context can be found in Millman [1977].

In the first instance, three relatively simple methods based on geometrical optics were used to evaluate the effects of tropospheric refraction on ELOISE AoA paths:

1 The 4/3 effective-Earth approximation [Schelleng et al., 1933; Miller, 1951; Bean & Thayer, 1959a]. By raising the Earth radius, the problem is transformed to a coordinate system in which the effective ray trajectory through the troposphere is a straight line (i.e. equivalent to free-space propagation); this greatly simplifies the computations. It derives from the assumption of a constant vertical refractivity gradient, of the order of -40 km^{-1}, where refractivity N is related to the refractive index n by the formula $N \equiv (n - 1) \times 10^6$. Strictly speaking, the optimal radius-scaling factor ("K") varies with the initial elevation angle [Miller, 1951], as well as the meteorological conditions affecting the refractivity profile, but to a first-order approximation it is assumed to be a constant value of 4/3. Typical errors in this approach are explored by Doerry [2013].

2 Numerical ray tracing through a realistic model refractivity profile of the lower atmosphere, such as the Central Radio Propagation Laboratory (CRPL) Exponential Reference Atmosphere [Bean & Thayer, 1959b]. The CRPL profile, presented in equation (5.1) as a function of height h above mean sea level, is based on the segmented linear/exponential model developed by Bean & Thayer [1959a]. It depends only on the surface refractivity $N_0 = N(h_0)$ derived from meteorological data. While the average global surface refractivity is about 325 [Altshuler, 1998], a value of 300 is adopted here to better match a sample of Alice Springs radiosonde data collected at 0 UT on 4 September 2015.

$$
N(h) = \begin{cases} N_0 + (h - h_0)\Delta N, & h_0 \leq h < h_0 + 1\,\mathrm{km} \\ N_1 \exp\left(\frac{-\ln(N_1/105)}{8 - h_0}(h - h_0 - 1)\right), & h_0 + 1 \leq h < 9\,\mathrm{km} \\ 105 \exp\left(-0.1424(h - 9)\right), & h \geq 9\,\mathrm{km} \end{cases} \quad (5.1)
$$

where $\Delta N = -7.32 \exp(0.005577 N_0)$ is the surface refractivity gradient, and $N_1 = N_0 + \Delta N$ is the refractivity at 1 km above the surface.

3 Numerical ray tracing through a tropospheric refractivity profile derived from radiosonde temperature (T), pressure (p), and dew-point (T_d) measurements. The standard two-term empirical formula for the refractivity of dry air [Smith & Weintraub, 1953] was used to calculate the profile:

$$N = \frac{77.6}{T}\left(p + \frac{4810e}{T}\right),\tag{5.2}$$

where vapour pressure $e = 6.112 \exp\left(\frac{17.62T_d}{243.12+T_d}\right)$ hPa for T_d in °C [Magnus-Tetens formula, World Meteorological Organization, 2012, Annex 4.B]. The example in **Figure 5.8** shows the Alice Springs radiosonde data from 0 UT on 4 September 2015 and its resultant refractivity profile. This is compared against the CRPL profile in **Figure 5.9**.

For all three approaches, rays are fired off at a range of elevations through the troposphere, and ray homing is used to converge on a specified oblique ground range. The ionosphere is treated as a simple mirror reflector. For method (1), an upper height of 8 km was adopted for the refracting region of the troposphere, to provide satisfactory agreement with the CRPL model results; the region above, up to the ionospheric virtual height, is treated as free space. For methods (2) and (3), the ray elevation at each spherically-stratified layer of the troposphere is calculated by the piecewise application of Bouguer's formula [Schulkin, 1952]. It is assumed that the refractivity profile is smoothly varying, to avoid anomalous effects like ducting and caustics.

Figure 5.10 shows the elevation differences calculated for equivalent ionospheric paths with and without tropospheric refraction. Three representative ground ranges (1000, 2000 and 3000 km) were considered for each of the three methods, while allowing the virtual height of reflection to vary from 80 to 600 km. There is a slight dependence on the choice of ground range due to the spherical Earth geometry. Longer paths suffer slightly more refraction for the same uncorrected elevation angle, causing the small spread (sometimes almost indistinguishable) between lines of the same colour. Across all elevation angles, this spread remains less than about 10% of the absolute elevation difference, which itself tends to be reasonably consistent between all three methods, particularly (2) and (3). Note that the effect of the troposphere on the modelled group delay and its apparent virtual height is negligible (<1 km), and is thus not shown here.
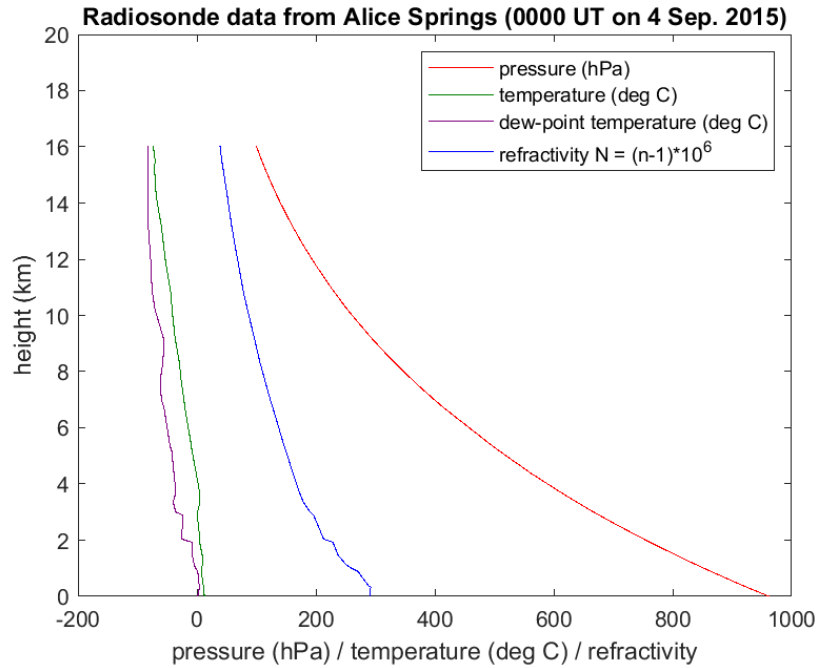
Figure 5.8. Sample radiosonde data (red, green and purple lines) from Alice Springs, and the resultant refractivity profile calculated using equation (5.2) (blue line). The observations were made by the Australian Bureau of Meteorology [Newton, 2018].
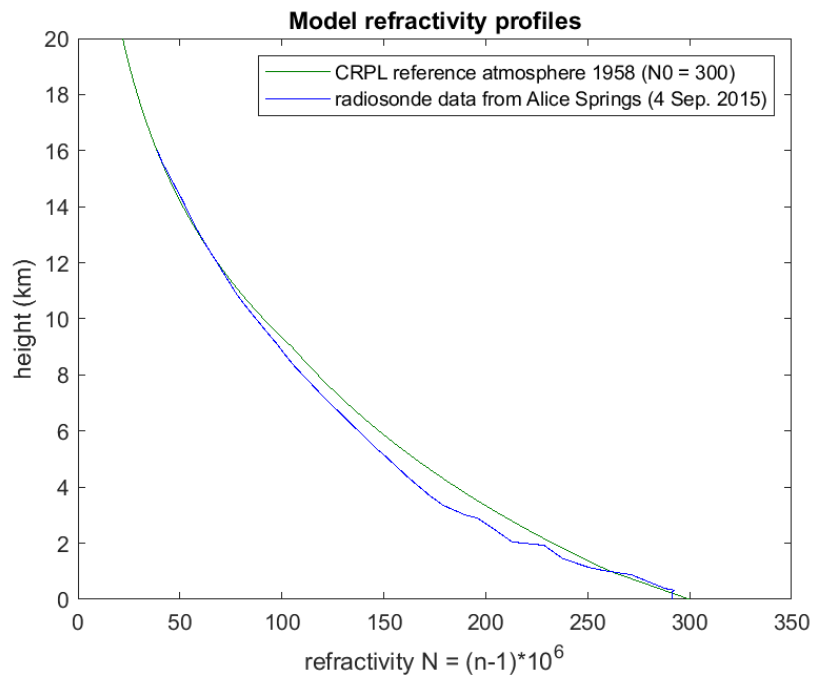
Figure 5.9. Comparison of refractivity profiles calculated using the CRPL reference model and Alice Springs radiosonde data.
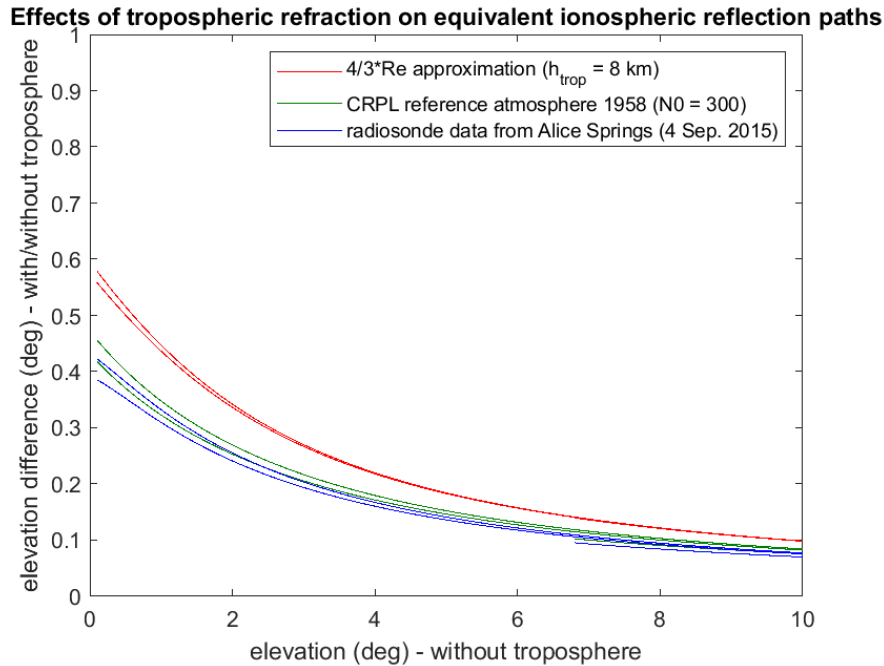
Figure 5.10. Equivalent elevation differences with and without tropospheric refraction, for the three alternative methods described: (1) 4/3 effective-Earth approximation (red), (2) ray tracing through the CRPL refractivity profile (green), and (3) ray tracing through a radiosonde-derived refractivity profile (blue). For each method, there are three lines of the same colour, representing 1000, 2000 and 3000 km ground range, which overlap in elevation where different virtual heights support the same launch angle. The shorter path (1000 km) only appears above 7°, corresponding to the minimum virtual height of 80 km.

Given that 60% of ray bending occurs in the first 1 km [Bean & Thayer, 1959a], and less than 10% occurs in the region above 10 km [Shulkin, 1952], it is not unreasonable to assume a constant refractivity gradient, as in method (1). Nevertheless, the more realistic profiles of methods (2) and (3) are believed to provide a greater level of accuracy. The fact that the elevation differences calculated from these methods never exceed a few tenths of a degree places any tropospheric correction below the measurement uncertainties of the ELOISE elevation estimates presented in **Section 4.4** (**Table 4.1**). Indeed, at elevations above 10°, the tropospheric effect is reduced to less than 0.1°, which is expected to be undetectable by the ELOISE AoA system. As such, tropospheric refraction is believed to be of only minor significance, and certainly cannot account for the magnitude of the low-elevation biases observed.

5.4   Array mutual coupling

The on-board ELOISE AoA processing, described in **Section 4.2**, adopted empirical calibration corrections to compensate for unknown positional/timing errors, mutual coupling effects, and receiver mismatch. However, some proportion of this correction has the potential to be predicted by simple electromagnetic modelling of the array and the interactions between antenna elements. The amplitude and phase distortions caused by mutual coupling can be expressed in terms of the array manifold, which is effectively a set of complex steering vectors for "phasing up" the array over the full hemisphere, and knowledge of the true manifold enables the effects to be counteracted during AoA estimation [Friedlander & Weiss, 1991; Singh et al., 2013]. This section considers two such mutual coupling models in terms of their ability to capture frequency and AoA dependent deviations in the inter-element phase differences from the planar wavefront assumption (i.e. the theoretical array manifold).

A first-order representation of array mutual coupling is via a mutual impedance matrix, which considers the induced currents between isolated pairs of vertically polarised dipoles (equivalent to monopoles on a perfectly conductive, infinite ground plane). It is based on the method of moments and assumes one piecewise sinusoidal mode on each centre-fed virtual dipole [Kraus, 1988, ch. 9.17]. Self-impedance is calculated as a special case of mutual impedance, with separation distance replaced by the antenna radius.

The resistive (R_{mn}) and reactive (X_{mn}) components of the mutual impedance ($Z_{mn} = R_{mn} + iX_{mn}$) between monopoles m and n of equal length l and separation distance d are given by equations (30) and (31), respectively, from Cox [1947]. These are reproduced as equations (5.3) and (5.4) below.

$$
\begin{aligned}
R_{mn} = \frac{15}{\sin^2(kl)} \{ & 4Ci(u_1) - 2Ci(u_0) - 2Ci(v_0) \\
& + \cos(2kl)\left[Ci(w_1) - 2Ci(v_0) + Ci(x_1) - 2Ci(u_0) + 2Ci(u_1)\right] \quad (5.3) \\
& + \sin(2kl)\left[Si(w_1) - 2Si(v_0) - Si(x_1) + 2Si(u_0)\right] \},
\end{aligned}
$$

and

$$X_{mn} = \frac{15}{\sin^2(kl)} \{-4Si(u_1) + 2Si(u_0) + 2Si(v_0)$$
$$+ \cos(2kl)\left[-Si(w_1) + 2Si(v_0) - Si(x_1) + 2Si(u_0) - 2Si(u_1)\right] \quad (5.4)$$
$$+ \sin(2kl)\left[Ci(w_1) - 2Ci(v_0) - Ci(x_1) + 2Ci(u_0)\right]\},$$

where Si and Ci are the standard sine and cosine integrals,

$$Si(x) \equiv \int_0^x \frac{\sin t}{t} dt, \quad (5.5)$$

$$Ci(x) \equiv -\int_x^\infty \frac{\cos t}{t} dt, \quad (5.6)$$

k is the wave number, and $u_0 = k(\sqrt{d^2 + l^2} - l)$, $u_1 = kd$, $v_0 = k(\sqrt{d^2 + l^2} + l)$, $w_1 = k(\sqrt{d^2 + 4l^2} + 2l)$, and $x_1 = k(\sqrt{d^2 + 4l^2} - 2l)$ are the limits of integration. In the case of the ELOISE AoA array, the antenna length is 6.5 m, with outer diameter 127 mm. Strictly the antennas were not fed at the base, as this calculation assumes, but approximately one-third of the length above the ground.

The mutual impedance matrix Z, normalised to the load impedance $Z_L = 450 \, \Omega$, is then given by [Gupta & Ksienski, 1983; Hui, 2007]

$$Z = I + \frac{1}{Z_L}[Z_{mn}], \quad (5.7)$$

where $[Z_{mn}]$ is a 19×19 matrix in the case of the ELOISE arrays, containing the mutual impedances for each antenna pair.

Under this model, the phase of the measured voltage ψ_n is offset from the phase of the theoretical manifold ψ_{n0} according to the matrix inverse of Z (the so-called mutual coupling matrix). The array manifold phase correction associated with mutual coupling is therefore given by

$$\Delta\psi_n \equiv \psi_n - \psi_{n0} = \text{angle}(Z^{-1}). \qquad (5.8)$$

A more sophisticated approach is to use the method of moments to numerically estimate (in an iterative manner) the currents and, hence, the array manifold for a wire-segment representation of the array. This computation was carried out by Mr Charlie Williams for the ELOISE array configuration, using the Numerical Electromagnetics Code (NEC) 4.2 [Burke, 2011], sampled in 0.5 MHz and 0.5° steps. Realistic dimensions of the monopoles were included in the model, including the feed points, but a perfect ground plane was assumed for simplicity. Reradiated signals and array shading effects are captured by this approach; in other words, the field impinging on the array is non-uniform.

Figure 5.11 compares the array manifold phase corrections calculated using both the NEC model (top panel) and the Z-matrix approximation (bottom panel), as a function of bearing and elevation for a fixed frequency of 20 MHz. The common element of the ELOISE array (i.e. at the intersection between the two arms) has been chosen here, as it represents the antenna with the greatest number of neighbouring antennas, and therefore the most likely to be subject to mutual coupling effects. It is encouraging to see that the two models exhibit many of the same characteristics as a function of AoA, but the simpler Z-model tends to overestimate the magnitude of the correction. At lower frequencies, the correction varies much more slowly with AoA.

Regrettably, there was not a framework in place to directly measure mutual coupling effects during the ELOISE experiment itself. However, it is possible to compare the two models against phase correction factors estimated from entire days of ionospheric observations on the paths of interest; an example of this is shown in **Figure 5.12**. The estimated corrections are simply the means of the AoA fit residuals, for just the strongest 1-hop F2-low peak at each frequency of each ionogram, binned in terms of frequency and AoA. While all paths can potentially contribute to each bin in the 3D frequency/bearing/elevation data cube, the distribution is naturally dominated by a single path in the direction of its great-circle bearing. As such, only a handful of discrete bearings are well characterised, and it is more meaningful to focus on a single bearing slice; in the case of **Figure 5.12**, this is the Lynd River to Laverton bearing (after adjusting for the rotational array correction).
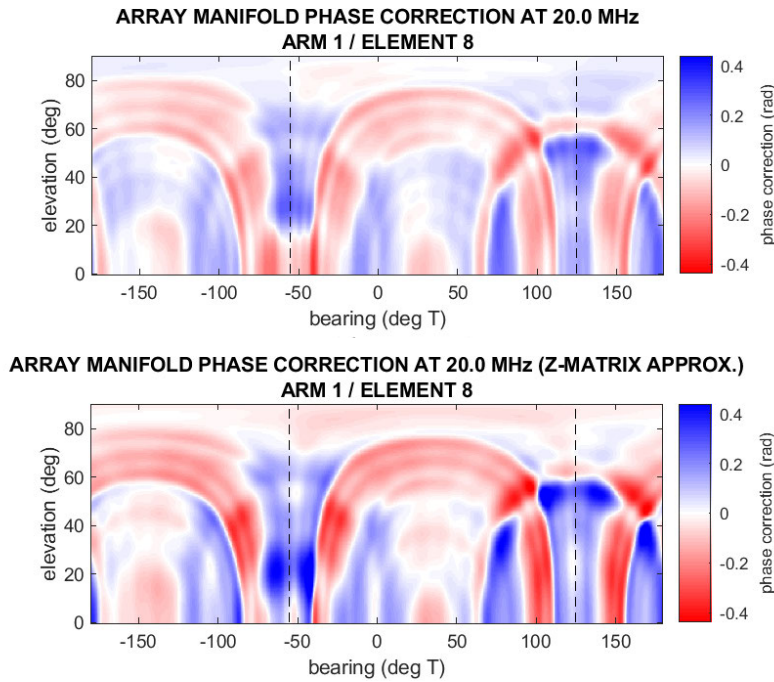
Figure 5.11. Array manifold phase corrections for the common ELOISE array element (#8) at 20 MHz, calculated using both the NEC model (top panel) and the Z-matrix approximation (bottom panel). The dashed vertical lines indicate ±90° from boresight for the Laverton array, within which the oblique paths of interest reside.

The top panel of **Figure 5.12**, showing the number of samples, resembles a full day of ionograms superposed, and gives a sense of where in the frequency/elevation space the estimated corrections are likely to be most reliable. Where the number of samples is zero, the estimated correction is undefined. The subsequent panels show the estimated and modelled array manifold phase corrections, rescaled in terms of metres to remove most of the frequency dependence. Although the general sizes of the corrections are comparable, the pattern estimated from observations (panel b) bears little similarity with the modelled patterns (panels c and d). In particular, the low-frequency positive correction in the observations (from 2–4 MHz) is not reproduced in the models. It is worth noting that this feature and others were found to be consistently observed from day to day. A possible explanation is that the finite dimensions of the ground mat may have a more pronounced influence at these longer wavelengths; due to space constraints, the mat extends less than half a wavelength in some directions below 15 MHz. Of course, the rescaling to metres does tend to amplify the apparent magnitude of this low-frequency feature, and it is not such a strong outlier in terms of degrees.

Another comparison of phase corrections, but for a different array element, is presented in **Figure 5.13**. Again, there is a banded pattern with frequency, and some features do look faintly similar, but the agreement is frankly insufficient for the model to be of value for calibration purposes. On the other hand, the estimated corrections do not reliably cover the whole 3D frequency/bearing/elevation space and are computationally expensive to generate over the full data set (requiring reprocessing of raw I/Q data). While not shown here, the modelled array manifolds were also used to evaluate the frequency-dependent effects on bearing estimation, but the results were sufficiently inconsistent with the patterns in **Figure 3.13** and **Figure 3.14** that this was not pursued further.
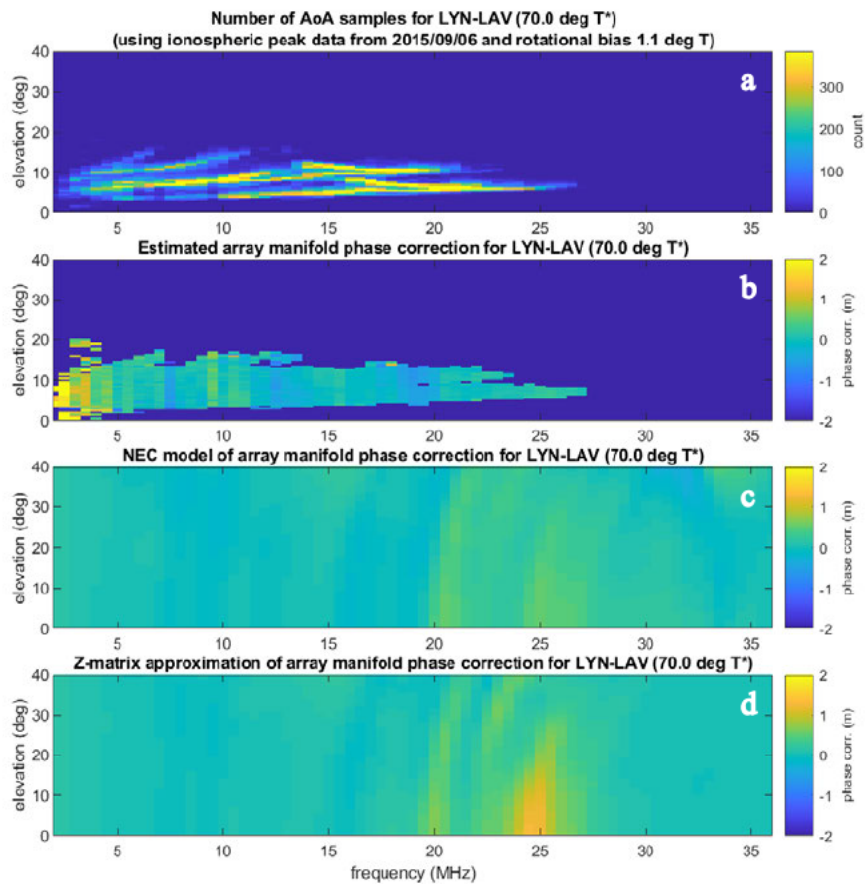
Figure 5.12. Comparison of array manifold phase corrections estimated from ionospheric observations at Laverton (panel b), alongside model predictions (panels c and d, for the NEC model and Z-matrix approximation, respectively). The results apply to the common ELOISE array element (#8) and the Lynd River path bearing. The number of samples contributing to the estimate in each frequency and elevation bin is also shown (panel a).
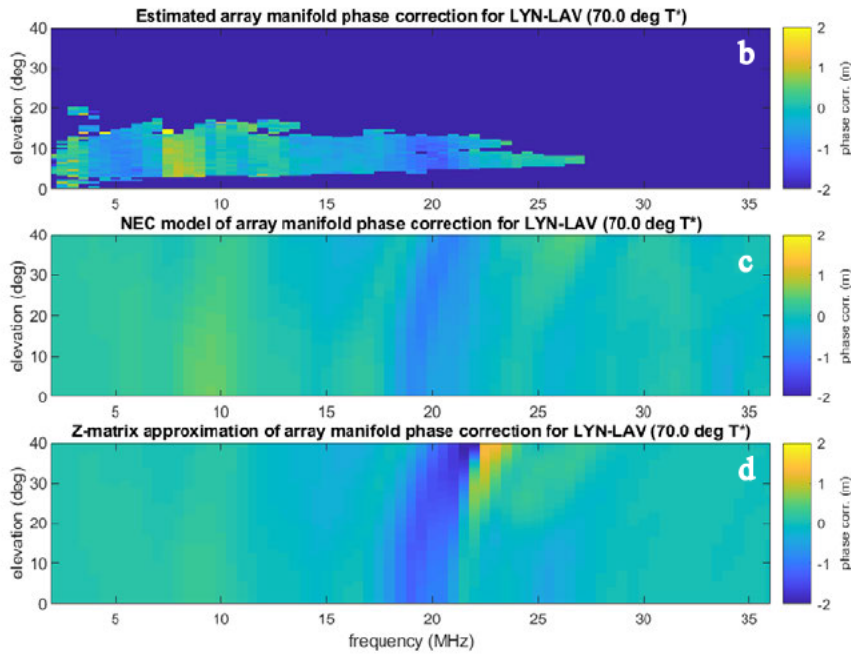
Figure 5.13. As in **Figure 5.12**, but evaluated for a different array element (#11, being the end element of the shorter arm). The number of contributing AoA samples (panel a in **Figure 5.12**) does not depend on the array element and is therefore omitted here.

5.5 Bearing and elevation bias corrections

Since the preceding analysis was unable to quantitatively account for AoA biases and phase residuals on the basis of physical models, an empirical approach to the calibration problem was ultimately pursued in post-processing. As in the previous section, aggregated ionospheric data from the strongest 1-hop F2-low peak at each frequency of each ionogram was used; in this case, spanning all observations from 5–17 September 2015. Path-dependent bias corrections were estimated for the post-fit AoA estimates only, as there was insufficient evidence to suggest that revisiting the inter-element phase differences (or timing offsets) from the on-board processing would result in any marked improvement in results (i.e. in terms of minimising the RMSE of the fit). Two implicit assumptions are that the long-term distribution of ionospheric gradients has zero mean, so gradient effects can be eliminated by aggregating over a sufficiently large data set, and that the biases do not vary significantly from day to day (at least, over the core ELOISE period). In practice, both magnetic deflection and persistent meridional

gradients around the shoulder of the equatorial anomaly will prevent the true AoA offsets from averaging to exactly zero.

The approach taken for estimating the path-dependent bearing biases is illustrated in **Figure 5.14** and **Figure 5.15**, for the Laverton and Coondambo arrays, respectively. Histograms of the measured offsets from the great-circle bearing are constructed for each path separately, and a parabola is fitted about the maximum to estimate the peak offset. The preliminary rotational bias of +1°, identified in **Section 3.3**, has been left uncorrected here, and therefore it is not surprising to see that the peak position across all paths (black dashed line) is in fact very similar to this preliminary estimate for both Laverton and Coondambo (cf. **Figure 3.13** and **Figure 3.14** in **Section 3.3**). Bearing biases on individual paths range from +0.2° (South Hedland to Laverton) to +2.1° (Kalkarindji to Laverton), indicating a distinct variation across the different directions. The widths of the distributions are believed to be dominated by ionospheric variability rather than measurement errors. Evidence of asymmetry in the distributions reveals a slight limitation of the simple aggregation approach for estimating bearing biases, with at least some bimodal behaviour (e.g. on the Kalkarindji to Laverton path in **Figure 3.13**) caused by the combination of a narrower daytime distribution and a broader night-time distribution of bearing offsets.

The elevation bias is a little more troublesome to estimate, as the predicted reference based on equivalent geometry (i.e. group delay) is itself uncertain. **Figure 5.16** shows the probability density functions of the observed elevation estimates as a function of the excess delay (i.e. the difference between the group delay and ground range in km), with the equivalent elevation predictions overlaid (dashed lines). As well as the influence of the ionospheric propagation support itself, the omni-directional gain patterns of both transmitting and receiving antennas, which notionally roll off at low and high elevation angles, may affect the number of peaks detected at each elevation. The standard Breit-Tuve relation [Reilly, 1985] tends to underestimate elevation, by an amount that increases with group delay [George, 1970]; thus, the apparent overestimation of elevation may be accentuated. Although empirical modifications to the Breit-Tuve relation have been proposed [George, 1970; Reilly, 1985], these have not been tested over the full range of ELOISE path lengths or require assumptions about the true height of reflection. As such, the standard relation was retained as the reference.
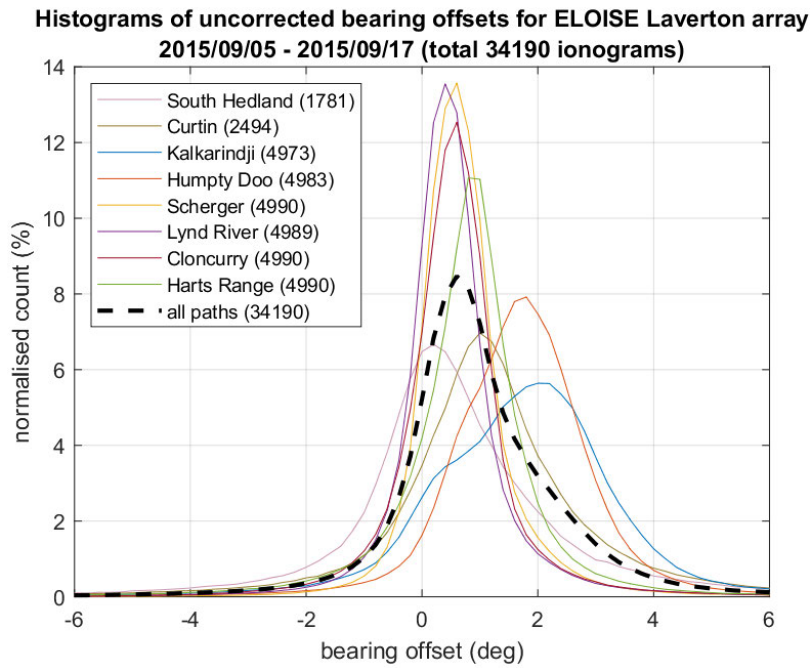
Figure 5.14. Histograms of uncorrected bearing offsets aggregated over 13 days of 1-hop F2-low ionogram observations from all Laverton paths. The numbers alongside each path indicate the number of contributing ionograms, and the black dashed line shows the total distribution over all data.
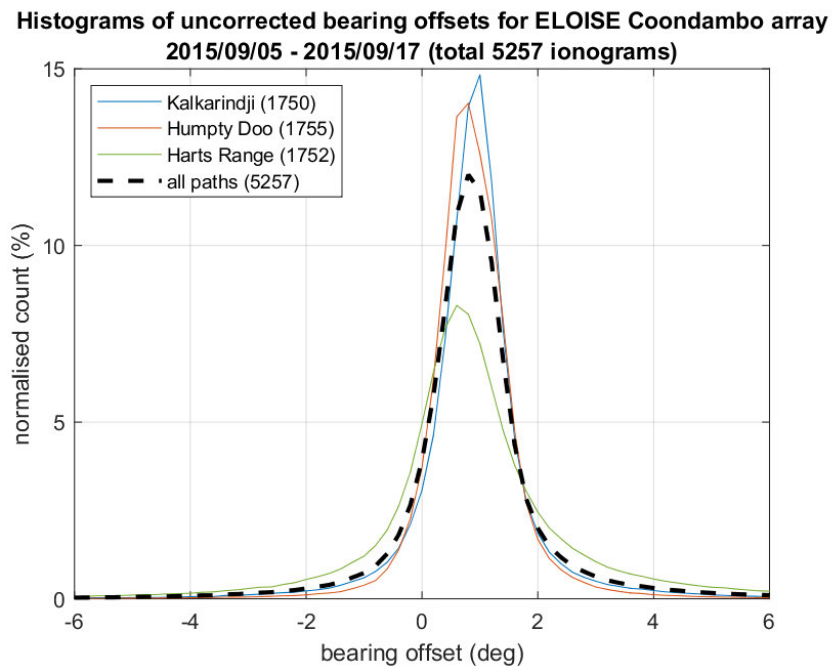
Figure 5.15. As above, but for the Coondambo array. Note that far fewer ionograms were collected from Coondambo over the course of the ELOISE experiment.

Once the equivalent elevation is subtracted off, the result is the offset plots in **Figure 5.17**. Given that each ELOISE path covers only a limited range of elevations, over which there is minimal variation in the elevation offset, a single scalar elevation bias was defined for each path based on the trimmed mean offset; this is shown by the magenta dashed lines, ranging from -1.2° (South Hedland to Laverton) to +3.2° (Lynd River to Laverton). The bias consistently exceeds +1° on the longer paths (>2000 km), although the proportion of this that is true overestimation of elevation versus deficiencies in the equivalent geometry reference remains unclear.
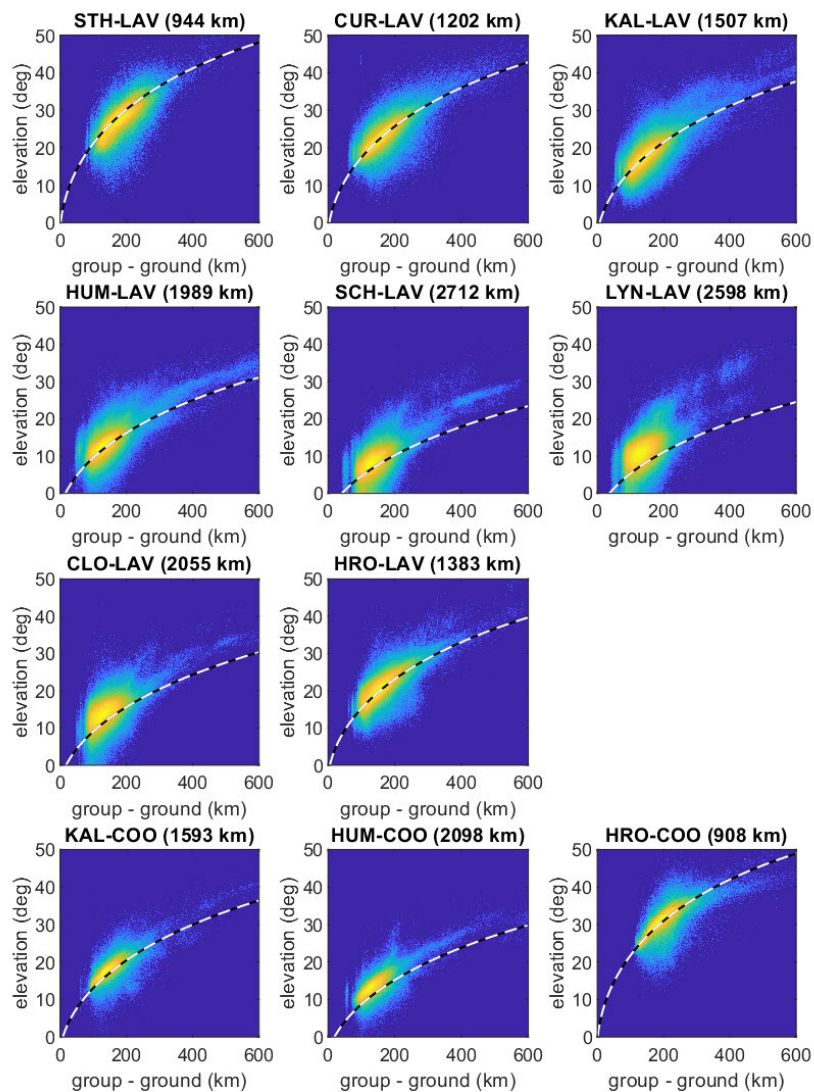
Figure 5.16. Probability density functions of uncorrected elevation versus excess delay (group – ground) for each of the ELOISE AoA paths, aggregated over 13 days of 1-hop F2-low ionogram observations. The standard Breit-Tuve equivalent geometry reference is overlaid as a dashed line on each panel.

117
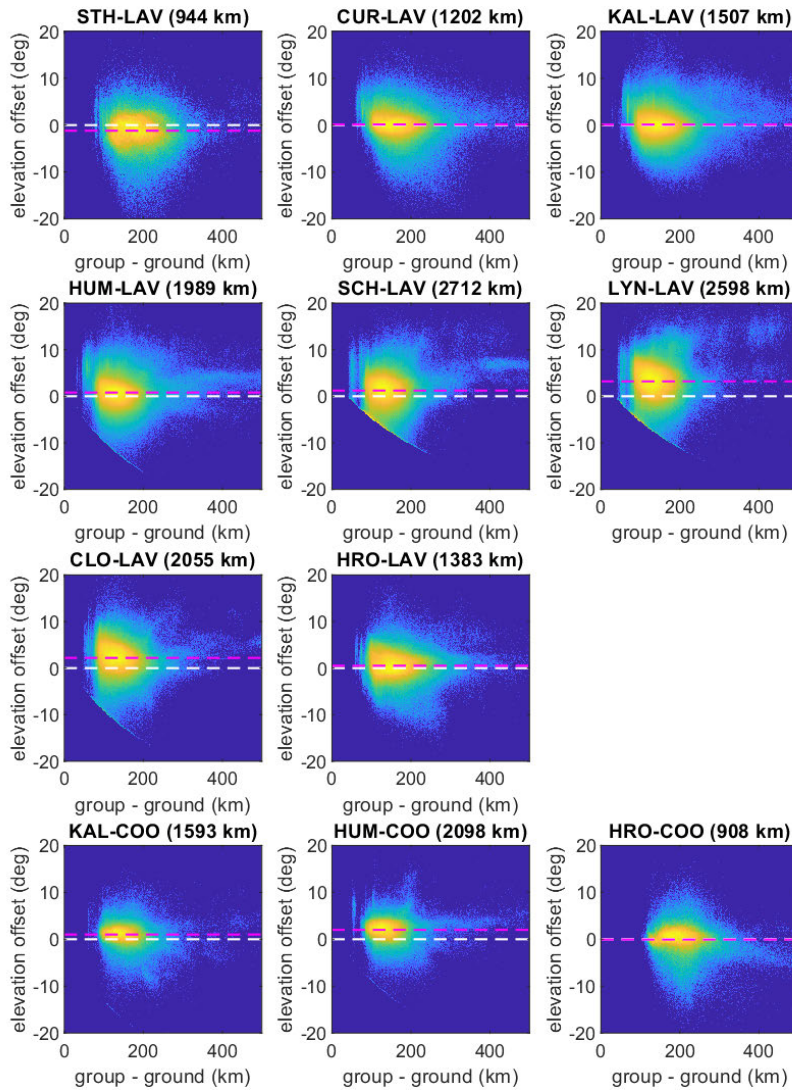
Figure 5.17. As in **Figure 5.16**, but rescaled in terms of the elevation offset (i.e. with respect to the predicted reference). The trimmed mean offset for each path (magenta dashed line) is used to define the path-dependent elevation bias.

Following bearing and elevation bias corrections, the resultant geographic heat map of 1-hop F2-low reflection points is shown in **Figure 5.18** (cf. **Figure 4.18** in **Section 4.4**, prior to corrections). The centres of the distributions are much more closely aligned with the great-circle midpoints (bullseyes) than before. While there are still some outstanding questions about the sources of the biases and errors in the equivalent geometry, which cast doubt on the absolute elevation accuracy, the power of this approach is that relative AoA displacements are now better characterised. This is most important for the perturbation analysis in subsequent chapters.

118
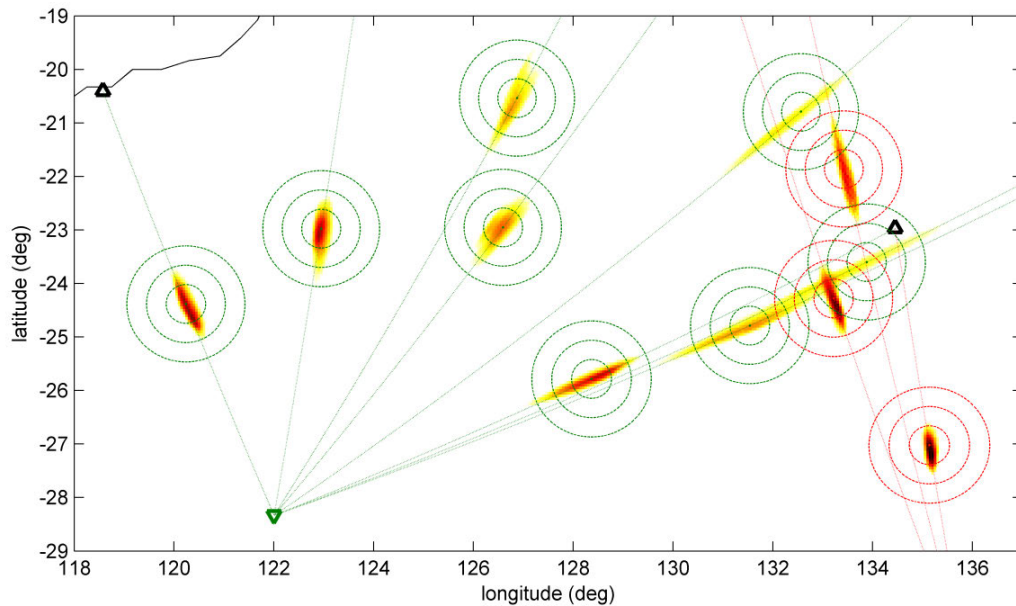
Figure 5.18. Heat maps showing the probability density distribution of 1-hop F2-low reflection points on all eleven ELOISE AoA paths from 5–17 September 2015, after path-dependent bearing and elevation bias corrections are applied.

5.6 Chapter summary

- Preliminary AoA ionograms and CSF dwells from the ELOISE experiment have been used to illustrate just a sample of the diverse set of mid-latitude ionospheric disturbances encountered, including TIDs, spread-F conditions, and the response to geomagnetic storms.

- Tropospheric refraction cannot account for the majority of the observed low-elevation bias, and the effect is predicted to be less than the measurement uncertainty of ELOISE elevation estimates. As such, it will be neglected in all the subsequent analysis.

- Mutual coupling between antenna array elements may introduce AoA errors on the scale of those observed. However, the models tested did not provide a sufficiently good match to the observations to offer a feasible correction strategy. Improvements to the model and/or a framework for directly measuring the correction factors more accurately are areas identified for future work.

119

- Path-dependent bearing and elevation bias corrections were ultimately used to account for offsets in aggregated ionospheric returns that could not otherwise be explained. These are empirically derived and serve to shift the distribution of mapped reflection points to be centred on the great-circle path midpoints.

6 Characterisation of small-scale (intra-dwell) ionospheric variability

6.1 Overview of intra-dwell channel scattering function data

The extensive collection of channel scattering function (CSF) dwells from the ELOISE experiment contains an interesting sample of ionospheric variability and mode interactions on a very small temporal scale. Although the standard ELOISE angle-of-arrival (AoA) processing adopts a coherent integration time (CIT) of 0.8 s for the wideband sweep (ionogram) and 6.4 s for the CSF, the ionosphere is not always stationary on these timescales, for example due to small turbulent structures associated with plasma instabilities and geomagnetic activity/pulsations. Unresolved mode mixing in particular can lead to rapid signal fading and apparent AoA fluctuations. Most commonly, where O/X components are unresolved, such effects are believed to be due to Faraday fading [Davies, 1990, ch. 7.7.3], caused by the polarisation axis of the combined O/X wave rotating with respect to the vertically polarised receiving array. Deep fading can severely degrade both HF communications and radar detection/tracking performance. It is therefore pertinent to investigate the scale of the small-scale ionospheric variability that exists within a single CIT, both from a geophysical standpoint and to identify implications on instrument fidelity.

By processing the CSF data on a sweep-by-sweep (chirp-by-chirp) basis, the 6.4 s CIT can be reduced to the inter-pulse period of 0.1 s, albeit without Doppler information, and the variation over the 64 contiguous sweeps gives a sense of what is being averaged over and thus sub-sampled by the full-length estimates. These so-called intra-dwell estimates were computed for both the strongest 1-hop F2-low peak and the strongest 1-hop F2-high peak in each CSF dwell. Typically O and X modes are not separated on low-ray returns at the CSF delay resolution, so greater signal instability (i.e. variation in power and AoA) is expected [Harris & Frazer, 2005].

A daytime example with minimal AoA variability is presented in **Figure 6.1**. The solid lines are the single-sweep (delay-only) peak estimates as a function of time, while the horizontal dashed lines are the standard (delay-Doppler) peak estimates from the full dwell. Only a single array arm, in this case the long arm, has been included to relax the assumption of 2D planarity and simplify the plot; on the Kalkarindji to Laverton path shown, variations in the steer angle for this arm are best interpreted as variations in apparent elevation. Steer angle uncertainties, for both single sweeps and the full dwell, are spatial standard deviations derived from the fit RMSE. The rate of change in the relative phase path P from one sweep to the next can be interpreted as a measure of the ionospheric Doppler shift f_{id}; that is, based on the formula $f_{id} = -\frac{f}{c}\frac{dP}{dt}$ [Davies, 1990, ch. 1.5.5].
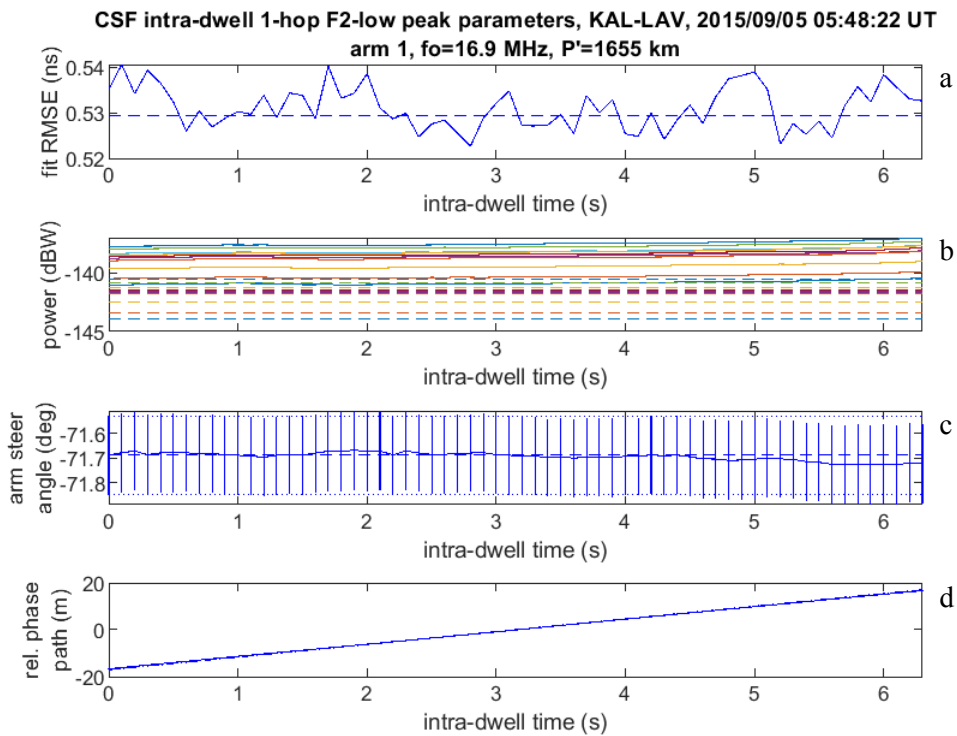
Figure 6.1. Sample intra-dwell 1-hop F2-low peak parameter estimates from Kalkarindji to Laverton, using only the long arm of the array. The panels, from top to bottom, show (a) the root-mean-square error (RMSE) of the AoA fit, (b) the peak power measured on each of the ten elements, (c) the estimated steer angle (with uncertainty bars representing one standard deviation), and (d) the relative phase path (integrated Doppler), each plotted as a function of intra-dwell time. The solid lines are the single-sweep estimates, while the dashed lines are the full-dwell estimates.

Without Doppler discrimination, the single-sweep estimates include additional power contributions, from both noise/interference sources and potentially extra signal modes; as such, it is not surprising to see higher power levels for the solid lines in the second panel. The spread between the ten elements (in different colours) is 3–4 dB, of which <1 dB was observed across the digital receivers using a directly injected waveform. The remainder is thought to be due to the combined mismatch in the cables, antenna transformers, and mutual coupling effects. Nevertheless, in terms of the key results, being steer angle and phase path, there is very little difference between the solid and dashed lines.

A contrasting example, with clear evidence of mode mixing and a fading period of about 2 s, is shown in **Figure 6.2**. This is similar to the CSF fading characteristics reported by Harris & Frazer [2005], who attributed amplitude and AoA fluctuations on the scale of tens of seconds to unresolved O/X components. In this case, with each periodic 20–30 dB dip in the peak power, the fit RMSE increases significantly and the estimated steer angle rapidly diverges from the mean, flipping through up to 20° at the null point. Such variations are non-physical, and point to incoherencies in the mode mixing pattern (including signal polarisation) across the 180 m aperture of the arm; this is equivalent to 3.5 wavelengths at 5.9 MHz. The fact that the null position is offset slightly between arm elements lends support to this explanation.

The variation in phase path is also intriguing, as it points to a change in sign of the apparent Doppler shift about the nulls. The full-dwell phase path estimate (i.e. the dashed line in the bottom panel) is derived directly from the measured Doppler shift, and therefore must be a straight line, although given that a taper was used in Doppler processing, it is heavily weighted towards the middle of the CIT. It remains an outstanding exercise to simulate these kinds of phase interactions, for example, using ray tracing [Västberg & Lundborg, 1997; Yau, 2003].

The remainder of this chapter will present small-scale variability statistics in the form of histograms, derived from the intra-dwell CSF data from 5–17 September 2015. The aim being to characterise the typical magnitude of variations in each of the peak parameters (i.e. power, steer angle, and Doppler), and dependencies on time-of-day, frequency, and bandwidth (delay resolution). Relationships with other measures of wavefront non-planarity and mode complexity are also explored, to better identify how and when the

ELOISE AoA estimates are likely to be degraded. It should be noted that while the focus here is on the CSF data, such issues also pertain to the ionogram observations, although their impact is expected to be less severe, owing to the shorter CIT and better delay resolution.
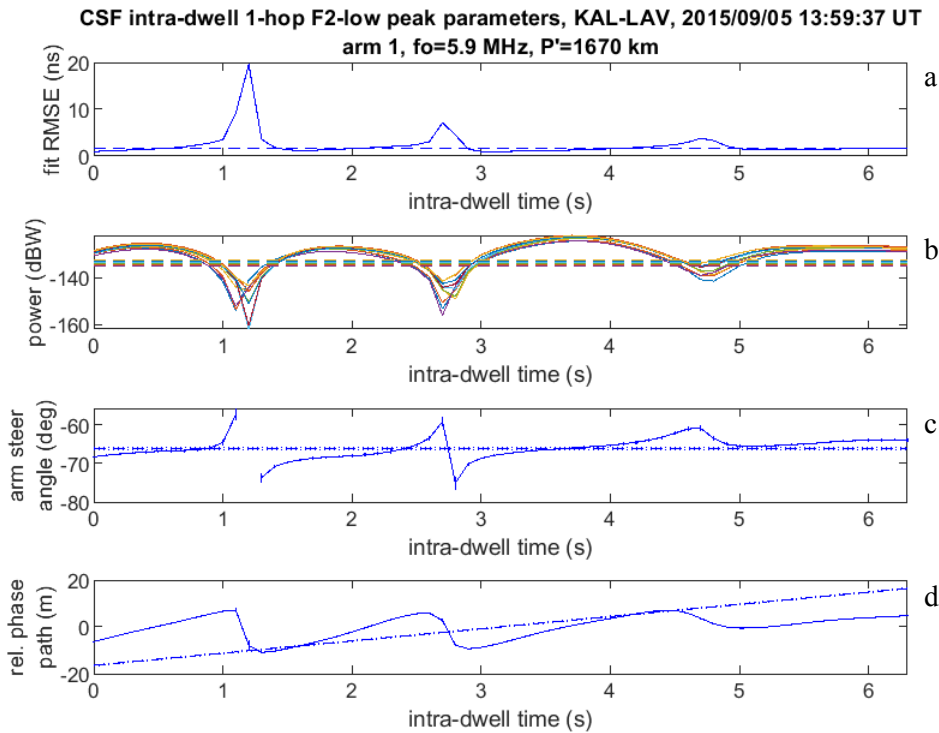
Figure 6.2. As in **Figure 6.1**, but for a less stable, composite peak exhibiting the characteristics of O/X mode mixing.

6.2 Intra-dwell channel scattering function statistics

Over the period from 5–17 September 2015, a total of 69,898 dwells from the six Laverton CSF paths were analysed for their intra-dwell temporal statistics. Coondambo CSF paths were excluded as the observations were predominantly daytime only. For each CSF observable, the two metrics considered were the time-averaged mean and standard deviation across the 64 single-sweep estimates. The mean is designed to be compared against the corresponding full-length estimate from standard CSF processing, to identify where unresolved Doppler components may be an issue, while the standard deviation is a measure of temporal variability. Note the distinction between this

temporal standard deviation and the spatial standard deviation that captures the AoA fit uncertainty over the full CIT.

Although these two intra-dwell metrics cannot fully capture the small-scale characteristics of parameter variability, including autocorrelation scale, time rate of change, and any beat period, they were chosen in the first instance as they are simple to interpret and require no assumptions about the form of the variation or stationarity of samples. So as not to mask the geometrical dependence on the asymmetric array, each arm was analysed separately, giving independent steer angles that must be combined to produce unambiguous bearing and elevation angles. At zero elevation, these steer angles are measured clockwise with respect to the individual arm boresights, being 125 and 215 °T for arms 1 and 2, respectively.

Given the volume of data, and therefore breadth of ionospheric propagation conditions, results were grouped into several broad sets:

- Low- versus high-ray F2 modes.

- CSF frequency ratio, with respect to the F2 maximum observed frequency (MOF).

- CSF waveform bandwidth (30, 15 or 7.5 kHz).

- Daytime versus night-time observations.

A multi-mode flag was also developed to identify and potentially remove peaks where the loss of Doppler information that comes with single-sweep processing would lead to unresolved multi-mode (merged peaks). **Figure 5.4** in **Section 5.2** contains one such example of multiple distinct CSF peaks at almost the same group delay. The special flag is simply based on counting the number of CSF peaks that overlap with the current peak in delay but are resolved in Doppler. However, so few examples of this condition were detected that the impact was believed to be fairly inconsequential, constituting just 5% of F2-low and 4% of F2-high peaks. Far more likely is unresolved multi-mode caused by insufficient delay resolution at the lower CSF bandwidths.

Figure 6.3 shows histograms of the intra-dwell means and standard deviations in the peak power, steer angle, and Doppler measurements, separated into low- and high-ray modes. For the Kalkarindji to Laverton path depicted, the median of the standard
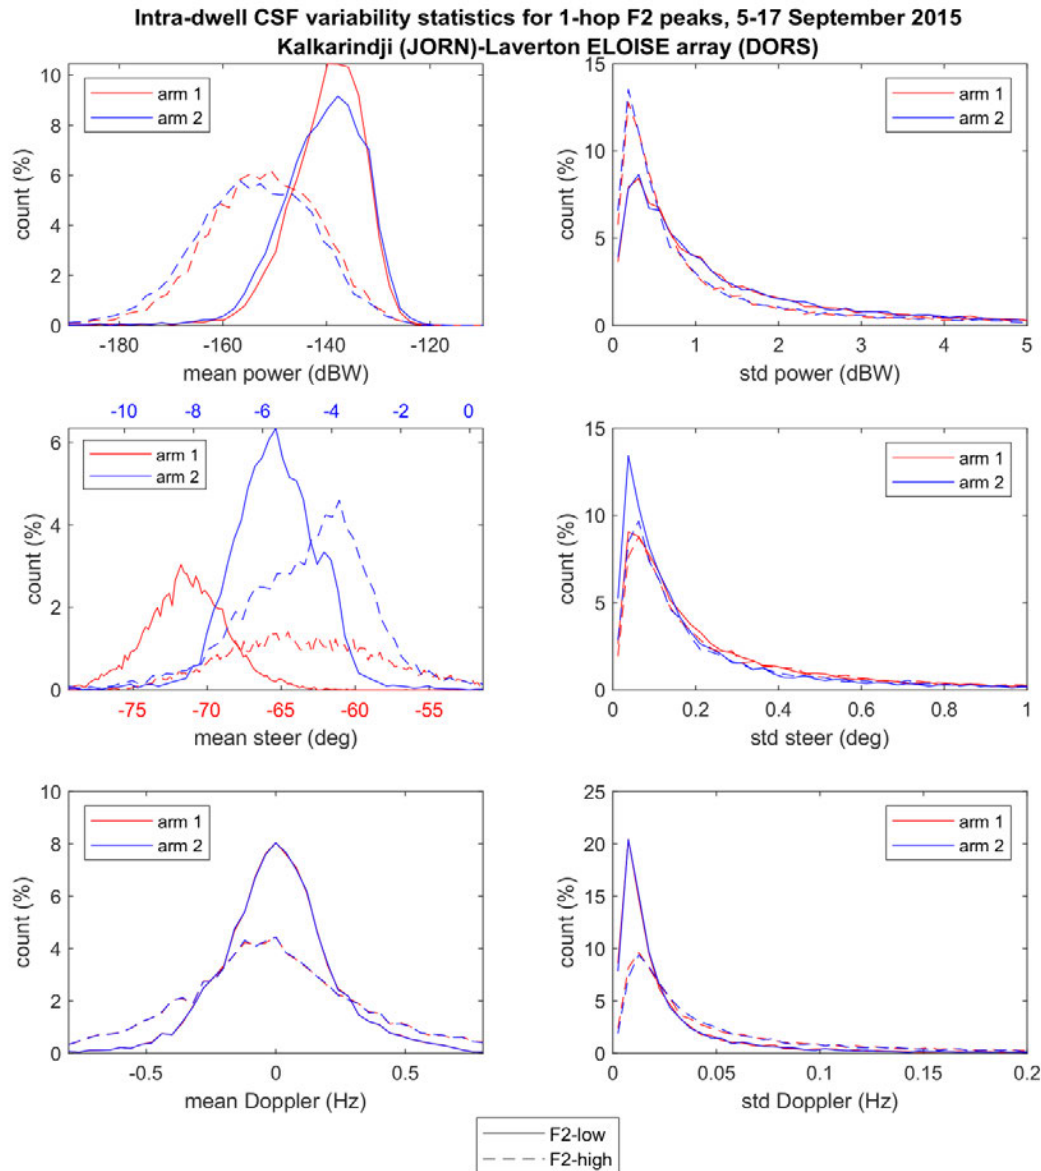
Figure 6.3. Intra-dwell CSF statistics for the Kalkarindji to Laverton path, split into 1-hop F2-low and F2-high modes (solid and dashed lines, respectively). The left column of histogram plots show the mean peak parameter values, for power (top), steer angle (middle), and Doppler (bottom), while the right column of plots show the corresponding standard deviations. The count for a given bin is expressed a percentage of the total CSF dwells. Each arm is analysed and plotted separately, in red and blue, and apart from the expected geometry dependence in the mean steer angles, the results are largely consistent. Note that in the left panel of the middle row, the mean steer angles are plotted with a separate colour-coded axis for each arm. Since the end-fire arm (red) is shown on a broader scale than the broad-side arm (blue), the area under the distribution appears to be less, but in fact the sum of all bins is the same (100%).

deviation distributions in power, steer, and Doppler are in the range of 0.5–1 dB, 0.1–0.2°, and 0.02–0.03 Hz, respectively. Not surprisingly, the high-ray peaks have consistently lower mean power (by ~15 dB), due to defocusing effects as a uniform fan of rays are spread more sparsely over the ground [Davies, 1990, p. 161], and mean steer angles closer to zero, due to increased beam coning at higher elevations. In addition, the high rays span a wider range of mean steer angles and Doppler values. This is because relatively small changes to the electron density profile during disturbed conditions (such as spread-F, solar terminator effects, and the passage of TIDs) give rise to large fluctuations in the AoA and phase path of high rays [Davies, 1990, p. 161], owing to their deeper penetration into the ionosphere.

The standard deviations indicate that these high-ray dynamics also translate to more variability in Doppler and, to a lesser extent, steer angle over the intra-dwell period; in other words, the assumption of stable propagation over the length of the CIT does not hold up as well for the high ray as for the low ray, leading to larger standard deviations in general. Interestingly the standard deviations in power show the reverse situation, with F2-high power varying less than F2-low over the CIT. This is contrary to what might be expected given the above but is probably due to the fact that F2-high is less likely to contain unresolved O/X multi-mode.

The steer angles for arm 1 (in red, close to endfire) naturally have a broader distribution than for arm 2 (in blue, close to broadside) as they are more strongly influenced by elevation changes, rather than the relatively stable bearing angle, according to the standard coning formula [Fabrizio, 2013, eq. (4.73); also quoted in the caption of **Table 3.3**]. Given that a larger aperture leads to more precise AoA estimation in general, with less sensitivity to receiver phase noise, the longer length of arm 1 is believed to mostly compensate for the reduced AoA resolution at endfire. This results in steer angle standard deviations that are mainly driven by intra-dwell ionospheric variability, rather than measurement noise, and are thus only marginally worse than arm 2.

Shifting attention now to the F2-low mode only, being the primary mode of choice for HF propagation under most conditions, the CSF statistics can be grouped by frequency. **Figure 6.4** shows histograms of the intra-dwell variability for three frequency ratio sets, centred on 90.0%, 81.0% and 72.9% of the F2 MOF (i.e. integer powers of 0.9), with a 4% tolerance. These correspond to the target frequencies chosen during CSF scheduling

(see **Section 4.2.1**), although for the purpose of this analysis the MOF is derived from the same epoch, rather than the previous epoch as in real-time scheduling. The larger frequency ratios penetrate slightly deeper into the F2 layer and, as such, may be expected to encounter more F2 variability, albeit to a lesser extent than the high ray. However, the larger ratios are also subject to less influence from the underlying E and F1 ionisation, with refractive effects from these layers diminishing as the wave frequency becomes increasingly greater than their respective critical frequencies (as captured in the Appleton-Hartree formula [Davies, 1990, ch. 3.1]).
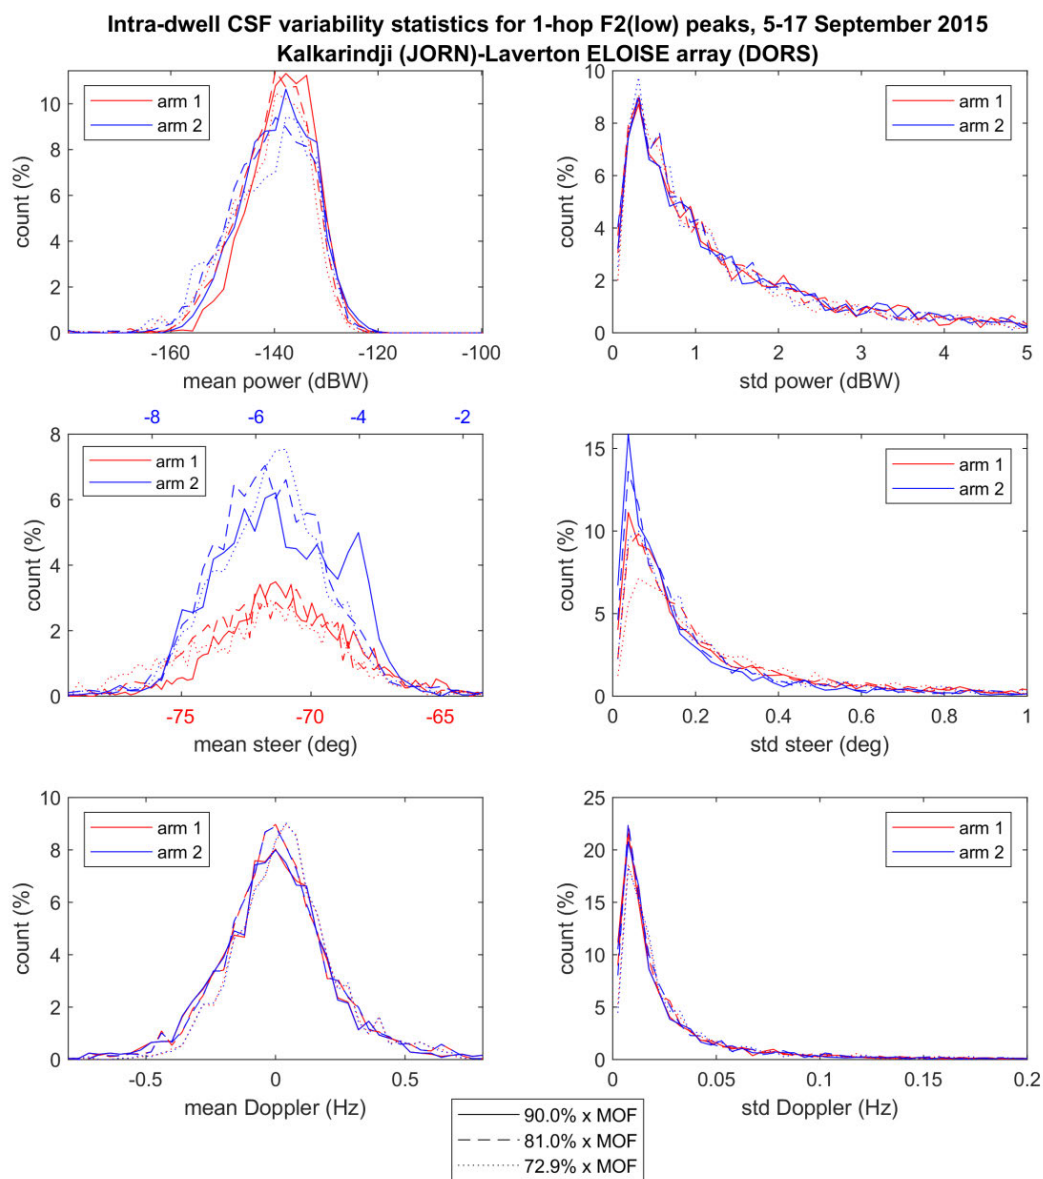
Figure 6.4. The same statistics as in **Figure 6.3**, except split according to the frequency ratio with respect to the F2 MOF (90.0%, 81.0% or 72.9%).

In practice, the differences appear to be relatively small, with only the standard deviations in steer angle (middle-right panel of **Figure 6.4**) exhibiting any distinct frequency dependence. Given that these standard deviations become more tightly distributed with increasing frequency ratio, representing an overall decrease in intra-dwell steer variations, this indicates that the greater influence from underlying E and F1 ionisation at lower frequencies may be important. The effect may also be a result of increased levels of unresolved multi-mode components at lower frequencies, although it is noted that the point at which the O/X ionogram traces intersect is frequently above 90% of the MOF during the day and below 72.9% of the MOF during the night; that is, outside the bounds of the chosen CSF frequencies. Nevertheless, unresolved O/X is still likely due to the limits on CSF delay resolution. The very slight increase in mean power at the highest frequency ratio (90.0%) is presumed to be due to focusing, although the effect is subtle after aggregating over a large number of observations and with the implicit but significant modification by antenna patterns [Dyson & Bennett, 1988].

As a final set of comparisons, **Figure 6.5** and **Figure 6.6** show the intra-dwell variability histograms broken down into waveform bandwidth and day/night sets, respectively. These are inherently related, given the diurnal pattern of HF spectrum availability that forces the use of narrower CSF bandwidths at night (cf. **Figure 4.4** in **Section 4.2.1**). Both show the same characteristic increase in mean power at night (lower bandwidths) due to reduced D region absorption, and slight shifts in the mean steer and Doppler estimates due to different day/night gradients and dynamics. In particular, the distributions of arm 2 steer angles are offset by approximately 1° between day (30 kHz bandwidth) and night (15 and 7.5 kHz bandwidths), indicative of a distinct change in path bearing across the two periods. Such large-scale persistent phenomena will be examined in further detail in **Chapter 7**.

All standard deviation measures are found to increase by varying extents at night (lower bandwidths), with the steer angles being most affected. Again, unresolved multi-mode, which can be more prevalent during night-time spread-F, is thought to contribute to this, as well as the fact that the array aperture in wavelengths, and hence AoA precision, is reduced for the lower frequencies (longer wavelengths) required to operate at night. A closer inspection of the statistics, replotted in **Figure 6.7** as a function of hour-of-day, reveals that the pre-dawn period of 18–22 UT typically contains the highest intra-dwell standard deviations; again this is most pronounced in the steer angles on both arms.
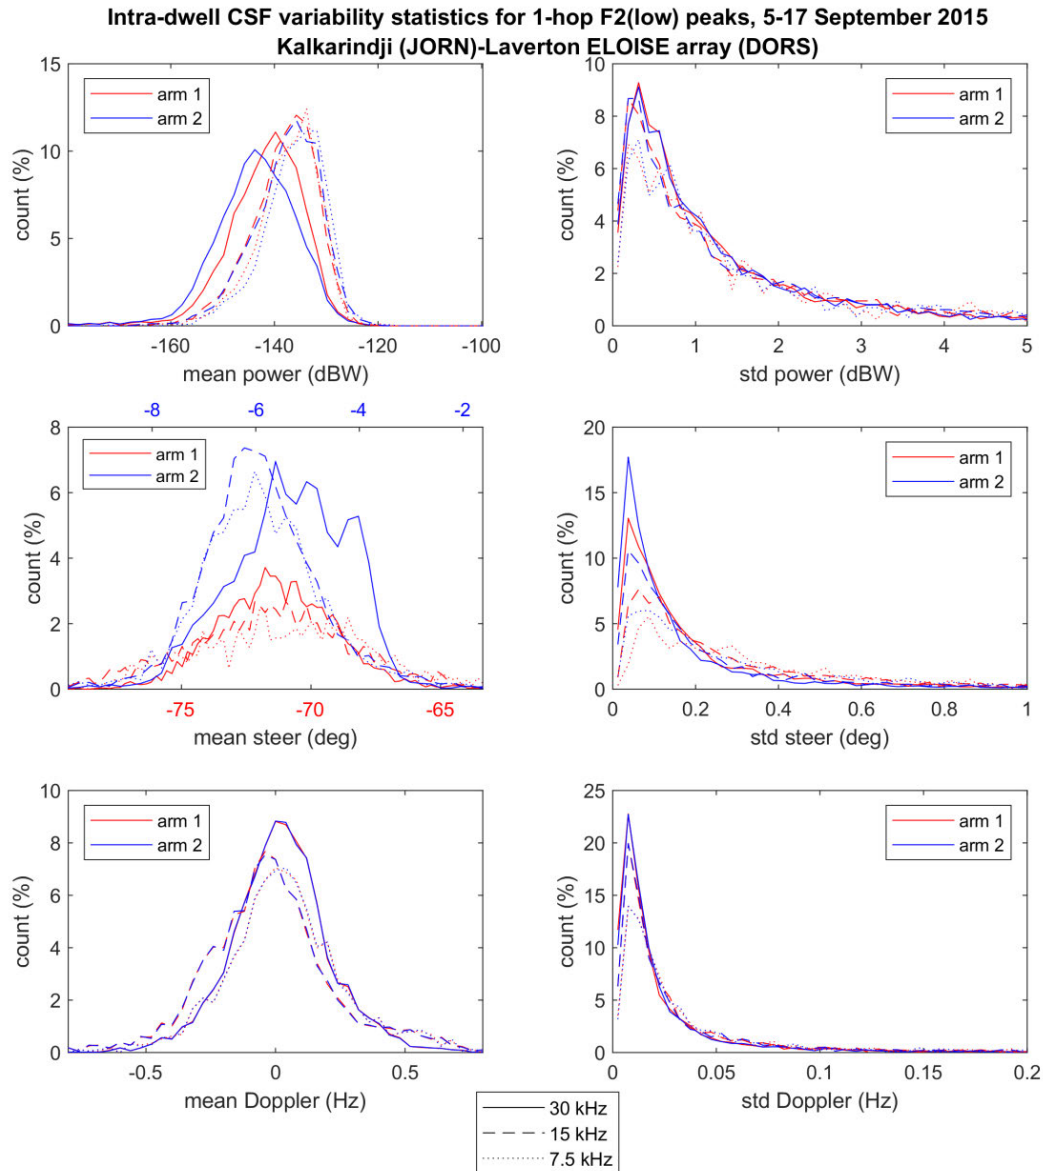
Figure 6.5. The same statistics as in **Figure 6.3**, except split according to the waveform bandwidth (30 kHz, 15 kHz or 7.5 kHz).

While not presented here, the other Laverton CSF paths tell much the same story. The path geometries, with respect to the two array arms, and day/night changes in ionospheric propagation tend to dominate the results. Of particular note is that the two paths with the most northerly 1-hop midpoints, being Humpty Doo and Scherger to Laverton (geomagnetic latitudes 29–30 °S), have the greatest differences in small-scale temporal variability between day and night. For instance, median intra-dwell standard deviations at night are around 2–2.5 times those during the day for steer, and 1.5–2

times for Doppler. This speaks to the influence of the shoulder of the equatorial anomaly, and the deposition of plasma structure associated with the fountain effect around the evening pre-reversal enhancement [Rishbeth, 2000; Balan et al., 2018], reaching even comparatively mid-latitudes.
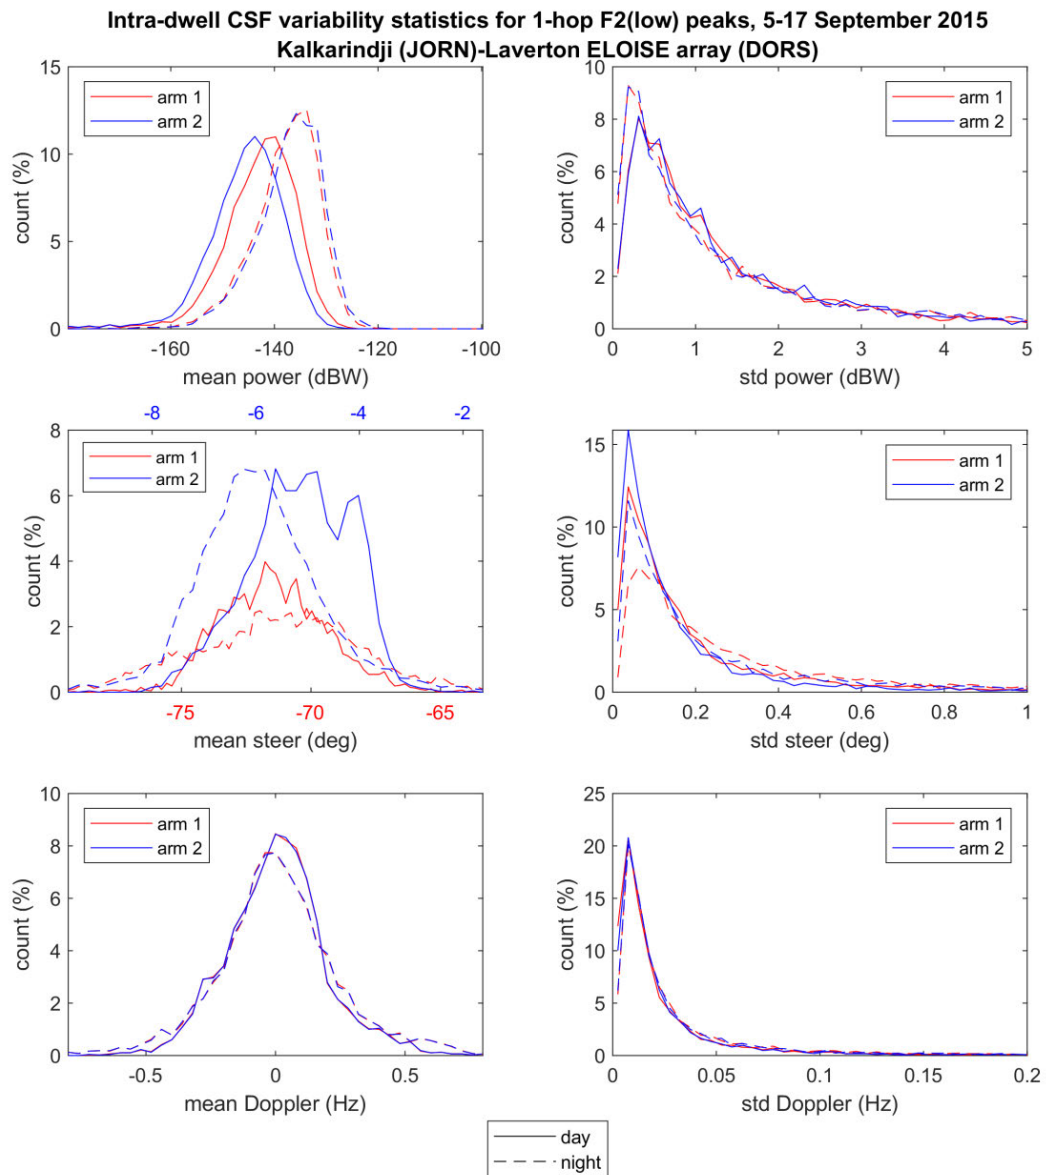
Figure 6.6. The same statistics as in **Figure 6.3**, except split according to daytime and night-time, based on the solar zenith angle at the Earth's surface.
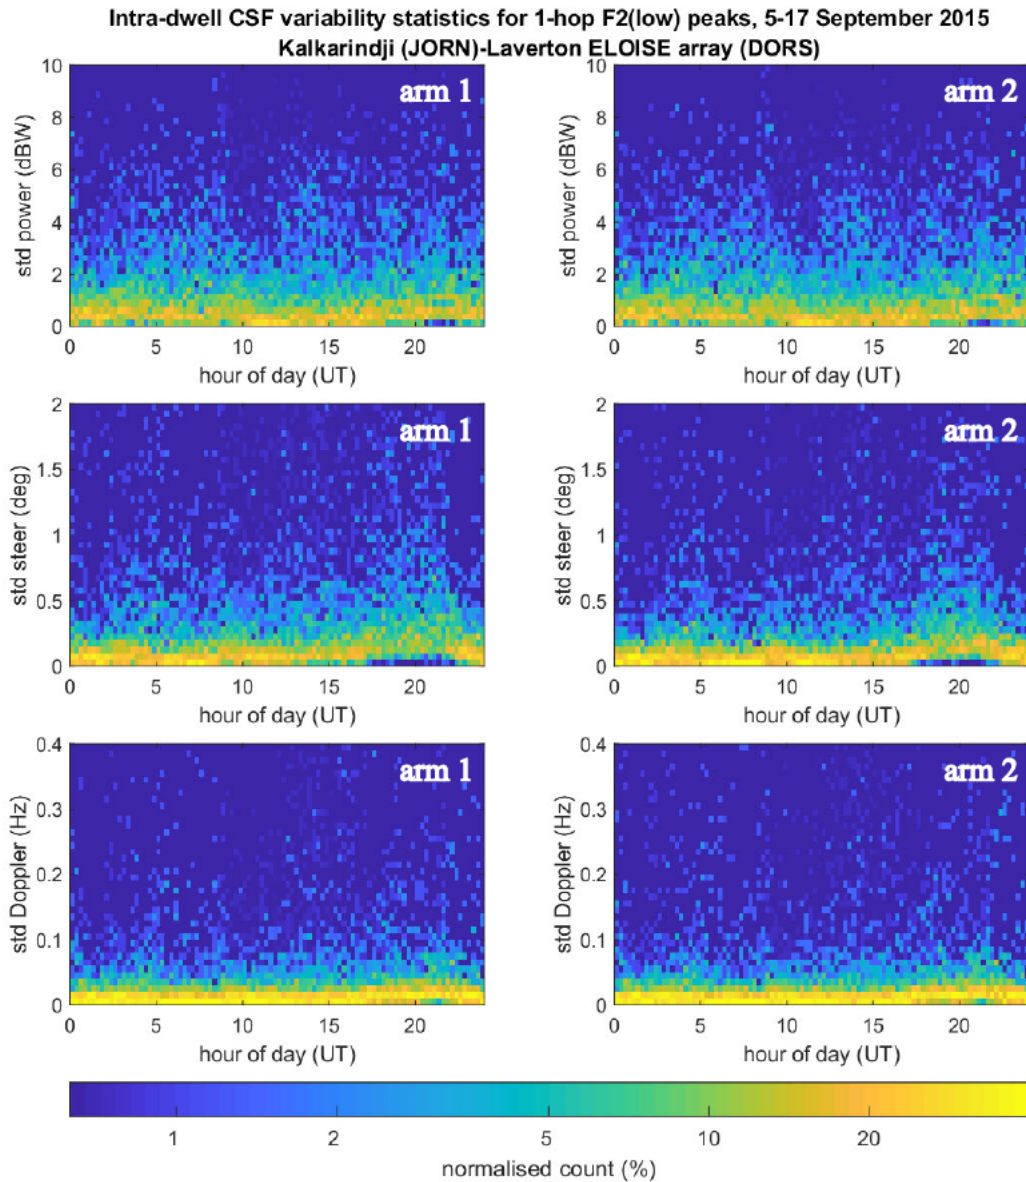
Figure 6.7. The same standard deviation statistics as in **Figure 6.3**, separated into individual (3.75 min) time bins, showing the hour-of-day dependence for each arm (left and right columns) and peak parameters (power, steer and Doppler, in top, middle and bottom rows, respectively). Each vertical image slice is a single histogram that sums to 100%. A logarithmic colour scale is used to better expose the tails of the distributions.

6.3   Non-planarity and spatial mode complexity

The above results demonstrate that small-scale intra-dwell variability, due to both intrinsic ionospheric path changes and mode resolvability, routinely leads to fluctuations in power, steer, and Doppler of around 0.5–1 dB, 0.1–0.2°, and 0.02–

132

0.03 Hz, respectively. For the fraction of the temporal variability caused by changes in the interference field impinging on the array (i.e. driven by unresolved multi-mode), it is of interest to identify whether this correlates with non-planarity ("crinkliness") in the time-averaged spatial pattern. This section considers two metrics, the time-averaged fit RMSE and the rank of the sample spatial covariance matrix, to see if there is a spatial analogue to the measures of temporal variability reported in the previous section. Issues of non-planarity and the space-time characteristics of HF sky-wave channels, including simulations, have previously been explored in some detail by Fabrizio [2000].

The example in **Figure 6.2** clearly suggests that the fit RMSE (both single-sweep and time-averaged) becomes elevated during periods of mode mixing, indicating a degree of non-planarity in the wavefront. However, in other cases, the RMSE can remain close to zero, indicating a planar wavefront, despite the apparent AoA varying over the course of the CIT. Furthermore, array manifold errors, due to mutual coupling, and measurement noise can have the effect of increasing the RMSE, despite these processes being potentially stationary and causing no time variation in the AoA.

The rank of the spatial covariance matrix, on the other hand, is an intrinsic measure of the number of incoherent directional sources impinging on the array within the narrowband group delay cell [Krim & Viberg, 1996]. While the size of the array still places a fundamental limit on the ability to separate signals in the presence of noise, with such a measure it is possible to detect the presence of multiple signals that cannot otherwise be resolved by conventional beamforming. Distinct multi-mode paths through the ionosphere will in general be statistically independent and, hence, incoherent, although O and X polarisation components may be partially coherent. In practice, the spatial covariance matrix must be estimated as a time-average of samples (sweeps); that is,

$$\widehat{\boldsymbol{R}}_x \equiv \frac{1}{P} \sum_{p=1}^{P} \boldsymbol{x}_p \boldsymbol{x}_p^H, \tag{6.1}$$

where $\boldsymbol{x}_p$ are column vectors containing the complex CSF cell values (after delay processing) measured across all array elements at the peak group delay, for sweep $p = 1, \dots, P$, and $\boldsymbol{x}_p^H$ denotes their Hermitian transposes. For the ELOISE CSF data on a

single arm, x_p is a 10×1 vector and the total number of sweeps $P = 64$. Note that the signal itself is expected to be highly correlated over consecutive sweeps.

The assumption of stationarity that is intrinsic to the estimator in equation (6.1) will be violated if the signal paths are changing with time. The resultant sample covariance matrix is therefore not always an accurate representation of the true covariance matrix at any instant, and the effective rank, obtained via eigenvalue decomposition of $\widehat{\boldsymbol{R}}_x$, will reflect a combination of the actual spatial mode complexity and degree of non-stationarity (i.e. small-scale temporal variability causing subspace leakage).

Consequently, for better or worse, the effective rank of the sample spatial covariance matrix is expected to be a rather reasonable indicator of temporal variation in the apparent AoA, associated with interactions between multiple directional components, but not necessarily a good measure of wavefront planarity. Certain classes of crinkly (non-planar) wavefronts, due to either the external propagation channel or calibration errors, that have a fixed spatial structure (i.e. effectively "frozen" on the scale of the CIT [Fabrizio, 2013, ch. 13.1.1]) and are thus fully correlated between receivers can nevertheless have a rank-1 covariance matrix and fixed AoA. As such, the rank metric may in fact be more closely aligned to the intra-dwell standard deviation across single-sweep AoAs, reported in the previous section, than the fit RMSE. Given that any real data will contain some level of uncorrelated noise, the sample covariance matrix is technically always full-rank, so eigenvalue decomposition was employed as a practical approach for evaluating the number of directional components and, thus, developing an effective proxy for the rank.

Returning to the two earlier examples in **Section 6.1** (i.e. **Figure 6.1** and **Figure 6.2**, representing peak estimates with stable and unstable intra-dwell AoAs, respectively), **Figure 6.8** demonstrates the impact of unresolved modes on the eigenspectra, derived from the sample spatial covariance matrices. Each 10×10 covariance matrix estimate is a 64-sweep average over the CSF CIT for array arm 1, as defined by equation (6.1), and is characterised by ten eigenvalues, which have been sorted from largest to smallest. The point at which each spectrum flattens out to become noise-like marks the boundary between the signal and noise subspaces, and the number of eigenvalues to the left of this may be interpreted as the number of signals. There appear to be about two signals for the stable (blue) peak and three signals for the unstable (orange) peak, indicating that

even temporally stable AoA characteristics do not necessarily imply a single pure mode. The relative power in the first principal component (i.e. sorted eigen-index 1) provides a good quantitative sense of how single-moded any given peak is. In these examples, the second component (i.e. eigen-index 2) is a huge 51 dB down from the first component for the stable peak in **Figure 6.1**, but a more modest 21 dB down for the unstable peak in **Figure 6.2**.
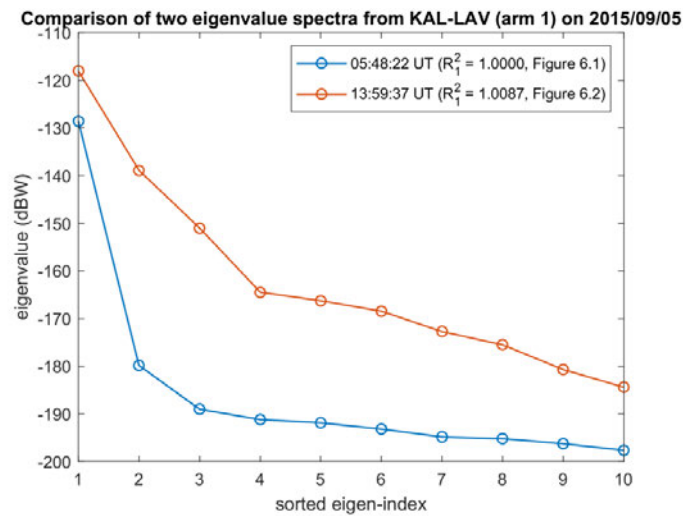
Figure 6.8. Eigenvalue spectra for the stable (blue) and unstable (orange) 1-hop F2-low peak examples in **Figure 6.1** and **Figure 6.2**, respectively. These derive from eigen-decomposition of the 10×10 sample spatial covariance matrix estimated for array arm 1.

To explore this relationship further, **Figure 6.9** shows hundreds of superimposed eigenvalue spectra, spanning 24 hours of F2-low CSF peak data on the Kalkarindji to Laverton path. The blue lines indicate the median and lower/upper deciles of the normalised spectra. Under quiet to unsettled geomagnetic conditions, typical of the ELOISE core period, it is apparent that the vast majority of the power is found in the first principal component. The median of the next strongest eigenvalue is roughly three orders of magnitude weaker, across day/night and both arms, implying a near-singular (rank-1) sample covariance matrix for most peaks; in other words, containing a single dominant propagation mode and AoA, to within the limits of the array resolution and SNR. The relative strength of the second eigenvalue increases slightly at night (mainly

135

in the upper decile). This again may reflect on increased multi-mode complexity and resolvability issues, especially given the reduced night-time delay resolution associated with narrower CSF bandwidths, or subspace leakage due to increased intra-dwell dynamics (non-stationarity). While the clear channel evaluation algorithm actively seeks to avoid co-channel interferers, because of the raised spectral congestion at night, such sources may also add to the signal subspace.

Considering just the first principal component, which is taken to represent the dominant propagation mode, **Figure 6.10** plots the proportion of variance (power) as a collection of cumulative density functions (CDFs) for the full period from 5–17 September 2015. This metric can be interpreted as a maximised R-squared statistic in the context of multivariate regression analysis [Härdle & Simar, 2015, ch. 11.3], and is defined as

$$R_1^2 \equiv \frac{\lambda_1}{\sum_{i=1}^{10} \lambda_i} = \frac{\lambda_1}{tr(\widehat{\boldsymbol{R}}_x)}, \tag{6.2}$$

where λ_i are the sorted eigenvalues ($i = 1, \dots, 10$) of the 10×10 sample covariance matrix $\widehat{\boldsymbol{R}}_x$.
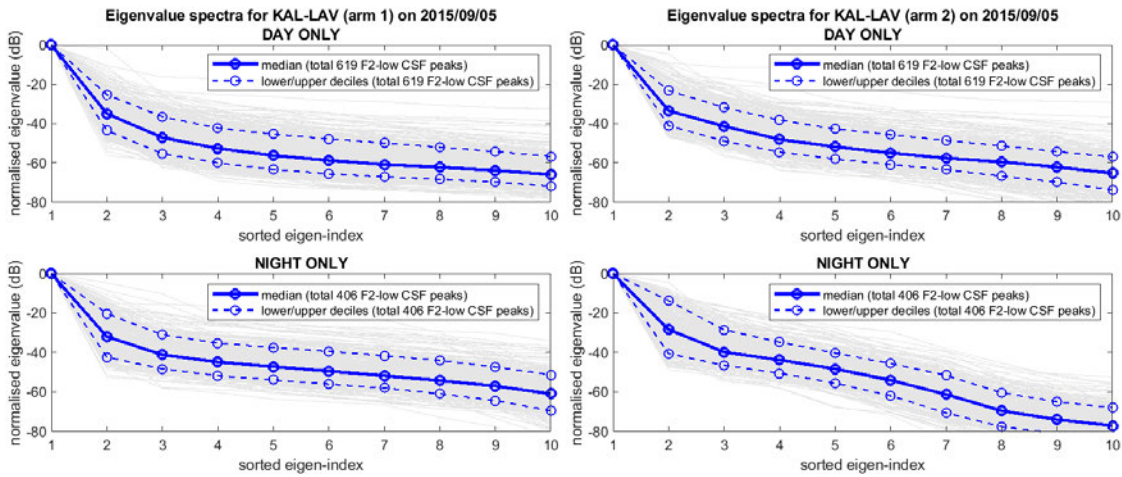
Figure 6.9. Normalised eigenvalue spectra for 24 hours of 1-hop F2-low peaks from Kalkarindji to Laverton on 5 September 2015. The left and right columns represent the two array arms, each characterised by a 10×10 spatial covariance matrix, while the top and bottom rows represent daytime and night-time observations, respectively. Spectra estimated from the individual F2-low CSF peaks are shown in grey, with the median and lower/upper deciles overlaid in blue.
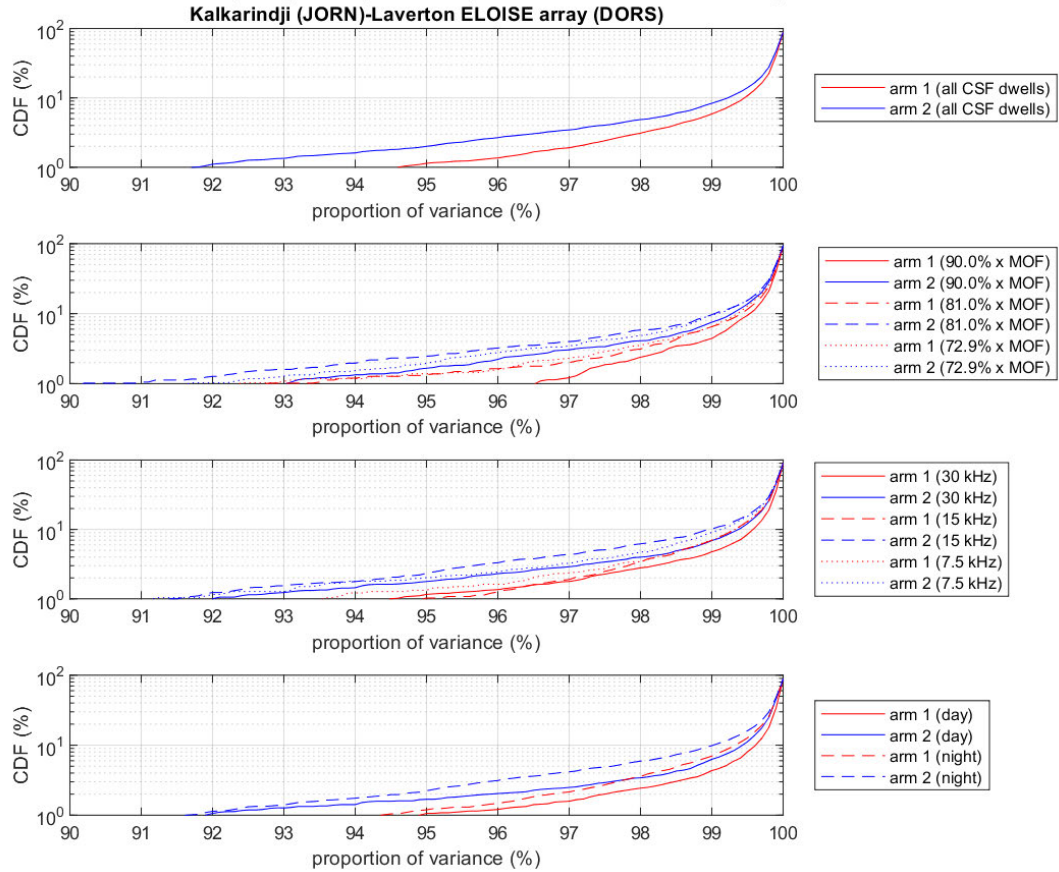
Figure 6.10. Cumulative density functions showing the proportion of variance in the first principal component of the sample spatial covariance matrix, aggregated over all F2-low CSF peaks from Kalkarindji to Laverton, 5–17 September 2015. The top panel shows all dwells combined, while in the subsequent panels results have been grouped into several different sets (the same as in the previous section), to highlight the dependence on frequency, bandwidth and time of day.

Each panel in **Figure 6.10** presents alternative set groupings on the two array arms, as described in the previous section, to establish how single-moded (and implicitly stationary) the CSF peaks tend to be under different operating conditions. Note that the median line depicted in **Figure 6.9** is equivalent to a CDF value of 50%, which is only reached towards the very right of the axes in **Figure 6.10**, when about 99.9% of the variance is in the first eigenvalue alone. The logarithmic CDF scale highlights the tail of the distribution, which varies more considerably across the different sets. Higher frequency ratios and bandwidths, during the day, tend to be closer to the single-mode

(rank-1) limit; perhaps not surprisingly, these are much the same factors that lead to reduced intra-dwell temporal variability.

To quantify these similarities, correlation coefficients were calculated between pairs of the four spatio-temporal variability metrics, namely:

1   the standard fit RMSE (from delay-Doppler processing over the full CIT);

2   the median intra-dwell fit RMSE (from delay-only processing of 64 single-sweep estimates);

3   the reciprocal of the proportion of variance in the first principal component (i.e. $1/R_1^2$, being a rank-like measure of the mode complexity); and

4   the intra-dwell (temporal) standard deviation in the estimated steer angle.

The correlation coefficients are summarised in **Figure 6.11**, with printed numbers indicating the mean coefficients over all six ELOISE CSF paths into Laverton and both array arms. As might be expected, given their closely related definitions for characterising wavefront planarity (or linearity per arm), the strongest non-trivial correlation is between metrics 1 and 2, with a mean coefficient of 0.70. A moderate (positive) correlation is also observed between metrics 3 and 4 (mean coefficient of 0.63), being the two more direct measures of temporal variability over the course of the CSF dwell, and to a lesser extent, between metric pairs 2/3 and 2/4 (mean coefficients of 0.57 and 0.55). In contrast, only weak correlations (albeit still positive) were found between metric pairs 1/3 and 1/4.

The interpretation of these results is that there definitely exists a connection between small-scale spatial and temporal variability measures. However, the correlation is not sufficiently strong to conclude that spatial variability (i.e. non-planarity) always accompanies temporal variability (i.e. intra-dwell mode beating and AoA perturbations). Therefore, in practice, it may be desirable to use a combination of metrics to flag peak data that do not adhere to the fundamental assumptions of the ELOISE AoA processing; these being based on planar wavefronts from a single stationary propagation mode.
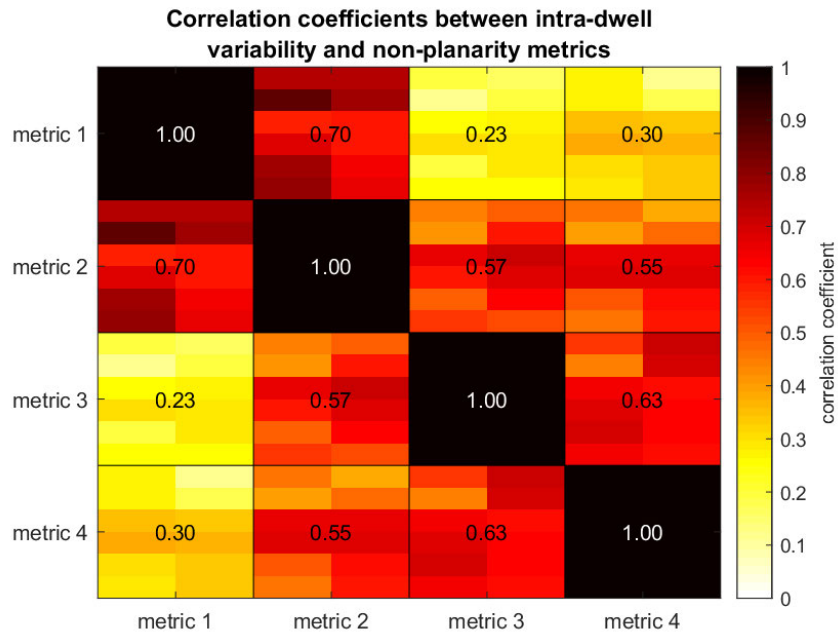
Figure 6.11. Correlation coefficients between pairs of the four spatio-temporal variability metrics. All six ELOISE CSF paths into Laverton have been included, amounting to a total of 60,032 F2-low peaks from 5–17 September 2015. Within each square on the 4×4 grid are 12 (6×2) coloured cells, representing the 6 paths and 2 arms individually. The printed numbers are the mean correlation coefficients over all paths and arms.

6.4 Future work

Although the ELOISE CSF data set offers many insights into small-scale propagation variability, it would be remiss to close this chapter without commenting on the need for further experimentation to better understand the physical nature of such disturbances and their impact on the standard 3.75 min AoA sampling. The analysis presented in this chapter is very much an opportunistic one, which was not envisaged as part of the ELOISE design. Should future observations be possible, more continuous ("stare") monitoring of AoA perturbations on timescales between the CSF CIT and the sounder revisit interval would be highly desirable. For disturbance periodicities in this range, there is an expectation of aliased contributions to the ELOISE AoA power spectra; these are impossible to resolve from the medium- to large-scale components discussed in the next chapter. Fading effects could also be better characterised through the use of a dual-polarisation antenna array.

Resolvability of modes remains an important issue in the ELOISE analysis, particularly for CSF data with its narrower processing bandwidths. To properly address this, an alternative 2D array, with a difference co-array more suited to imaging [Hoctor & Kassam, 1990], would be highly desirable; the existing ELOISE layout was simply not designed for this purpose, and as such, its co-array lacks many of the key vector spacings, as described in **Appendix B**. If the incidence of unresolved modes can be reduced, it is expected that the intrinsic small-scale variability that remains due to the ionosphere will be much better characterised.

6.5   Chapter summary

- The impact of unresolved modes and small-scale temporal variability on ELOISE AoA and Doppler estimates has been explored, by looking at examples and statistics from intra-dwell CSF processing (i.e. within the 6.4 s standard processing interval). In more severe cases, 20–30 dB dips in peak power and sudden AoA jumps of up to 20° were recorded.

- Intra-dwell variability statistics, presented as histograms of the means and standard deviations in CSF peak power, AoA (steer) and Doppler, indicate that small-scale effects are reasonably modest on the whole, with median standard deviations in the range of 0.5–1 dB, 0.1–0.2°, and 0.02–0.03 Hz, respectively. However, the distributions do distinctly shift from day to night (coinciding with narrower bandwidths) and between the F2 low and high rays, with more variability detected in the estimated steer angles for F2-low at night. This suggests that unresolved multi-mode is a major contributing factor, not just intrinsic ionospheric variability alone.

- No single metric fully captures both spatial and temporal variability in the propagation mode peak over the intra-dwell period. A combination of metrics, including those based on the RMSE of the AoA fit, principal components of the sample spatial covariance matrix, and intra-dwell standard deviation in AoA, may therefore offer the best ability to detect the small subset of peaks that fail the assumptions of the ELOISE AoA processing.

7 Characterisation of medium- to large-scale ionospheric variability

7.1 Overview of ELOISE peak data and ionospheric models

Medium- to large-scale spatial and temporal variability in the ionosphere exists on scales of hundreds to thousands of kilometres and tens of minutes to hours, and accounts for the largest component of variability by amplitude. Routine physical drivers such as thermospheric neutral wind dynamics (including solar tides), passage of the solar terminator, changes in solar flux, and geomagnetic storms can all cause horizontal gradients and propagation path changes, which manifest as off-angle returns and Doppler signatures, respectively. These types of variations may not be entirely predictable by climatology but are reasonably well sampled by conventional ionosonde networks. Travelling ionospheric disturbances (TIDs), in particular, occupy this part of the variability domain at F2 heights, although a more detailed treatment of this class of disturbance is reserved for **Chapter 8**.

The ELOISE angle-of-arrival (AoA) peak data set, with its 3.75 min revisit interval, captures the multi-dimensional propagation effects of medium- to large-scale variability across five key observables: signal amplitude, group delay, Doppler, bearing, and elevation. Given that signal amplitude is not strictly required for understanding path geometry changes caused by electron density gradients and perturbations, the focus of this chapter will be on the latter four observables. The information provided by AoA measurements is essential for gradient characterisation, since anti-symmetric gradients (such as a tilt about the midpoint) introduce only second-order effects in group delay [Bennett & Dyson, 2002] and thus may not otherwise be detected. As a general rule, a purely down-range tilt will impose an elevation offset, while a cross-range tilt will predominantly cause a bearing offset. It is noted again that Doppler is only available from the channel scattering function (CSF) observations.

As an example, an 8-hour night-time sequence of ELOISE AoA peak data is shown in **Figure 7.1**. The five rows represent each of the observables, for both ionogram peaks (left column) and CSF peaks (right column) at a frequency of around 8 MHz. For the most part, the ionogram and CSF data sets capture the same ionospheric variability, although the higher delay resolution of the ionograms means that O- and X-mode components can more often be resolved (e.g. between 1200 and 1330 UT). The vertical error bars, derived from the fitted width of the peaks and scaled to one standard deviation, assuming a normal distribution, capture both this instrumental resolution and ionospheric spread. Some but not all of the variations in time are correlated or anti-correlated (by eye) between the different observables, and some but not all are quasi-periodic (TID-like). A significant proportion of the seemingly random "noise" from sample to sample falls outside the width of the uncertainty bars, and is believed to be due to unresolved mode interactions and aliased small-scale effects, as discussed in **Chapter 6**.

Common periodicities across group delay, Doppler and AoA are a good indicator of TID activity, although it is important to note that not all quasi-periodic signatures are TIDs; for example, mesoscale plasma structure deposited by the equatorial fountain or localised patches of enhanced ionisation moving back and forth over the sounder midpoint may produce similar effects (e.g. Lynn et al. [2013]; Lynn et al. [2016]). Spatial lags between multiple sounder paths and a phase pattern descending as a function of height and plasma frequency are additional characteristics that may be used to identify atmospheric gravity waves (AGWs) manifested as TIDs [Hines, 1960].

In the case of **Figure 7.1**, a long sequence of medium-scale TIDs was clearly observed as "kinks" in the traces of the individual ionograms, although a distinct period was lacking from the time series itself, likely indicative of multiple interacting wave components. This was typical of the majority of observations during the ELOISE campaign. While these TIDs did not manifest as large perturbations in the group delay of the fitted peaks, the effects are more pronounced in Doppler, elevation and, to a lesser extent, bearing. During this period, the elevation differs from that predicted using equivalent geometry (i.e. based on group delay only, as shown by the dashed line in **Figure 7.1**) by up to ±5°. The bearing also appeared to be shifted slightly away from the equator throughout the night, which may be attributable to persistent large-scale horizontal gradients in the background ionosphere.
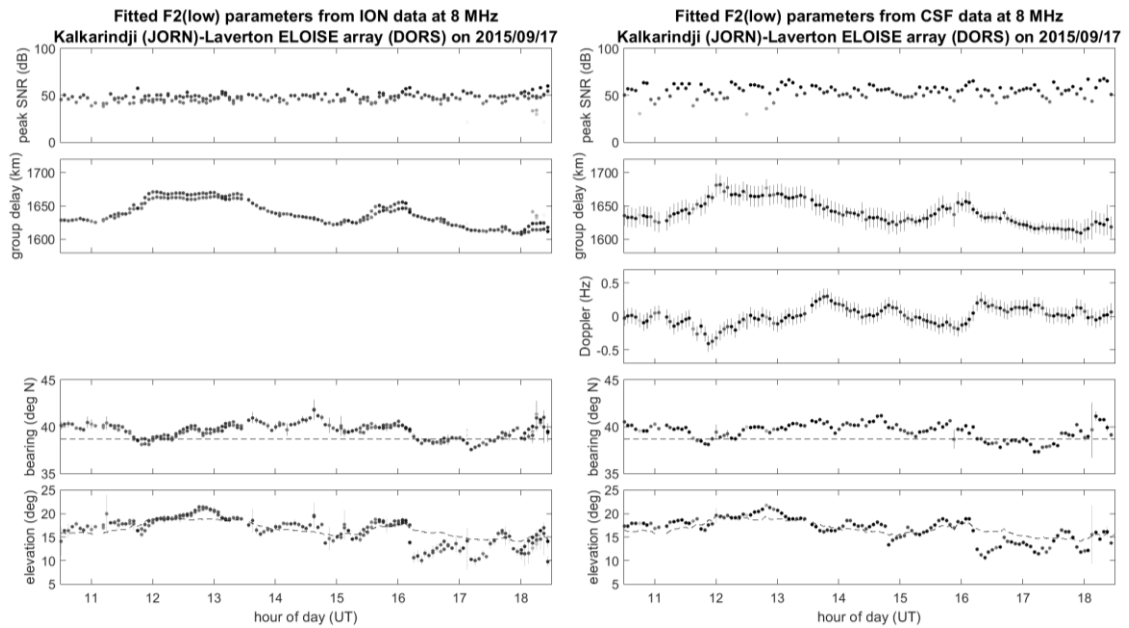
Figure 7.1. Time series of multi-dimensional AoA peak data, measuring the night-time 1-hop F2-low mode at ~8 MHz from Kalkarindji to Laverton (1507 km). The left column shows the ionogram peaks (without Doppler information), and the right column shows the interleaved CSF peaks. Although the CSF target band was fixed for this period, the frequency allocation did change slightly over an interval of ~0.5 MHz based on clear channel availability. From top to bottom the panels show peak SNR, group delay, Doppler, bearing, and elevation, with AoA bias corrections from **Section 5.5** applied. Markers are shaded according to their peak SNR, with uncertainty bars measuring one standard deviation. The dashed lines in the bottom two panels represent the equivalent great circle mirror geometry (i.e. predicted from group delay alone).

Without routine and widespread AoA observations in Australia, it currently falls to conventional sounder-based models of the ionospheric electron density to capture this class of propagation variability and its impact on HF radar and communication systems. The JORN Real-Time Ionospheric Model (RTIM) [Barnes et al., 2000; Gardiner-Garden, Heitmann, et al., 2008; Gardiner-Garden et al., 2011; Gardiner-Garden et al., 2018] is the most detailed "nowcast" model of the Australian ionosphere, and since 2017 has been accepting both OIS and VIS midpoint profiles as inputs [Ayliffe et al., 2019]. A stand-alone, batch-mode implementation of the RTIM's spatial and temporal mapping principles has been developed and tested on the ELOISE data set by Dr Robert Gardiner Garden [Gardiner-Garden et al., 2019], and is used as the benchmark empirical

143

model in this analysis. It is herein referred to as the Gardiner-Garden Ionospheric Model (GGIM).

The purpose of this chapter is to investigate what can be inferred about the medium- to large-scale ionospheric variability from AoA and Doppler measurements on the oblique ELOISE paths, and how the effects of spatial gradients (tilts) and the time rate of change of phase path (Doppler) compare to equivalent estimates from profile fitting and spatial mapping across a dense network of conventional OIS and VIS paths. Both regular patterns of daily variability and irregular isolated disturbances will be studied, across both the observable and tilt domains. Horizontal tilt estimation and drift analysis are proposed as mechanisms to compare disturbance signatures across different paths, with quite different geometries.

7.2   Diurnal patterns in angle-of-arrival observables

While the F2 layer varies quite substantially from day to day, there are nevertheless consistent solar-driven diurnal patterns that clearly manifest in the ELOISE AoA observables. This component of the variability should ideally be captured in the monthly-median behaviour of any model. To visualise such patterns, all F2-low CSF peaks from 5–17 September 2015 were aggregated into time-varying histograms for each Laverton CSF path, and plotted in the form of heat maps (normalised at each hour-of-day epoch). The result for Humpty Doo to Laverton is shown in **Figure 7.2**. At F2 heights over the path midpoint, sunrise and sunset times are around 1030 and 2030 UT.

Note that these plots combine multiple operating frequencies, so Doppler frequency shifts (in Hz) have been rescaled to Doppler velocities (i.e. rates of change of phase path, in m/s) to remove the explicit frequency dependence. As expressed in the following equation (e.g. Davies et al. [1962]), this includes a change of sign:

$$v_{id} \equiv \frac{dP}{dt} = -\frac{cf_{id}}{f}, \qquad\qquad (7.1)$$

where v_{id} is the ionospheric Doppler velocity, P is the phase path, c is the speed of light, f_{id} is the ionospheric Doppler frequency shift, and f is the operating frequency.
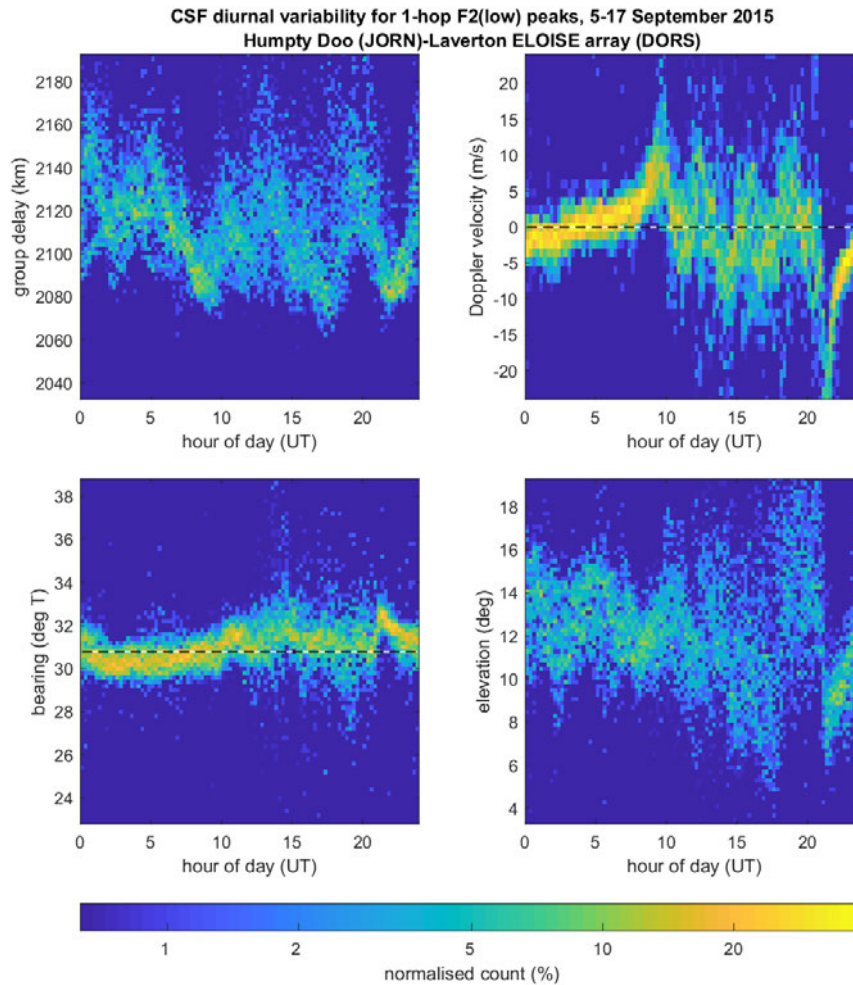
Figure 7.2. Diurnal variability in the 1-hop F2-low CSF peak estimates, for the Humpty Doo to Laverton path (1989 km). The top row of panels shows the peak group delay (left) and Doppler velocity (right), while the bottom row shows the peak bearing (left) and elevation (right). Dashed lines represent the unperturbed state (i.e. zero Doppler and great circle bearing). The histogram count is normalised independently at each epoch; that is, such that each vertical image slice sums to 100%. The local solar time at the path midpoint is approximately UT + 8.5 hr.

For a simple spherical mirror model, the Doppler velocity is equal to twice the vertical velocity component of the reflecting surface.

Some of the noteworthy features in **Figure 7.2** include:

- The presence of significant day-to-day variability in group delay and elevation, whereas Doppler and bearing are much more tightly constrained. The distributions,

even for Doppler and bearing, become quite dispersed at night when spread-F irregularities and unresolved modes are more prevalent.

- A tendency for group delay to increase over several hours after dawn, and decrease over several hours before dusk, with a faint V-shaped reversal pattern within an hour or so of the terminators (reaching a minimum value of ~2080 km at ~0900 UT and ~2200 UT). It is believed this is the influence of changes to the underlying E and F1 layer ionisation, which forms and decays very predictably from day to day, governed by the intensity of solar ionising radiation [Kouris & Muggleton, 1973; Muggleton, 1975; DuCharme et al., 1971; 1973]. A similar but weaker signature is also present in elevation, with an increase after dawn; not surprising since delay and elevation are related to first order by equivalent mirror geometry.

- Another more subtle descending signature in group delay, reaching a local minimum of ~2070 km at ~1730 UT, which resembles the characteristic midnight decrease of the F2 layer altitude initially identified in incoherent scatter radar observations from Arecibo, Puerto Rico [Nelson & Cogger, 1971], and subsequently studied more widely at sites including Townsville, Australia [Dandenault & Richards, 2015]. Such behaviour, attributable to tidal components in the neutral winds, is usually followed by a pre-dawn height rise, with a corresponding collapse in F2 ionisation levels (i.e. rapid recombination at lower altitudes, causing the F2 trace to recede). An inspection of the typical night-time F2 height variation sketched in Figure 1 of Nelson & Cogger [1971] reveals strong similarities with the diurnal variability in delay shown here.

- Sharp negative and positive Doppler velocity excursions at the dawn and dusk terminators, respectively, with a slow and steady transition in between, crossing zero Doppler at around 0400–0500 UT, just after local noon. That is to say that the phase path tends to decrease in the morning and increase in the afternoon, with the rate of change becoming larger in magnitude nearer the terminators, when the F2 height is known to rapidly fall (at dawn) and rise (at dusk). The zero-Doppler crossing at noon is exactly as predicted by basic Chapman layer theory [Pickering, 1975]. It is interesting that this is contrary to the rate of change of group delay for most of this time.

- Systematic offsets of almost ±2° in bearing across day and night, with a small but distinct shift towards north over the period surrounding local noon. Such bearing offsets can potentially be explained by persistent horizontal electron density gradients perpendicular to the path; in this case, from the north-west to the south-east. During the day, the electron density usually increases from the south to the north of Australia (on one side of the anomaly crest), so in theory, the apparent reflection point is expected to be displaced in the direction of the higher electron density and lower reflecting height to the north.

- A sudden jump in the bearing, away from north and towards the east, shortly after dawn, followed by a gradual retracing (over ~3 hours) back to the great circle path bearing (dashed line). This is almost certainly due to the strong east-west (zonal) gradient in electron density established between the daytime and night-time longitudinal sectors, with a typically lower F2 height of reflection in the eastern daytime sector.

The analysis was subsequently repeated for all F2-high peaks on the same Humpty Doo to Laverton path, with an equivalent set of panels shown in **Figure 7.3**. The characteristic Doppler and bearing features previously identified in the F2-low plots of **Figure 7.2** are still clearly present, and appear to be scaled up in magnitude by roughly a factor of two (the vertical axes have been adjusted accordingly). The typical daily range of Doppler velocities spans ±40 m/s for F2-high (cf. ±20 m/s for F2-low), while the range of bearing offsets spans ±4° (cf. ±2° for F2-low). This simply reflects the fact that the high ray propagation path undergoes more refraction and phase path retardation than the low ray, and is susceptible to gradients and irregularities over a larger spatial volume. Group delay and elevation are largely featureless in **Figure 7.3** due to the significant spread in values from day to day.

Similar diurnally consistent features present themselves on the other ELOISE CSF paths as well, although the bearing perturbations were found to be more pronounced on paths with closer to north-south alignment (e.g. Humpty Doo at a notional 31 °T with respect to Laverton), rather than east-west (e.g. Cloncurry at 70 °T). To illustrate this, a further set of F2-low plots are presented in **Figure 7.4** for Cloncurry to Laverton. A possible interpretation is that large-scale zonal (east-west) gradients are having a greater impact than meridional (north-south) gradients, noting again that it is the cross-plane gradient

that affects bearing. However, this does not then support the fact that the largest perturbations in bearing on the Humpty Doo path (**Figure 7.2**) occur around local noon, when zonal gradients should be relatively small (i.e. with the rate of change of solar zenith angle at a minimum). It would be of interest to repeat such an analysis on additional paths in future, to better understand this aspect dependence, and rule out instrumental effects.
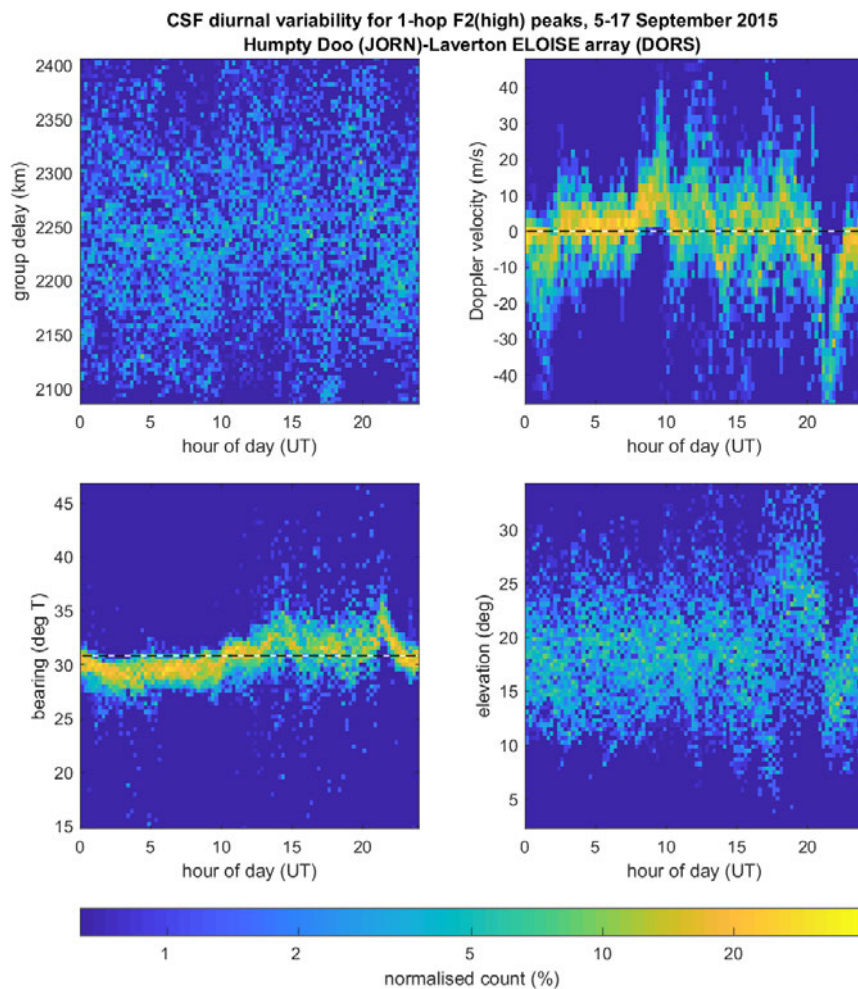
Figure 7.3. As in **Figure 7.2**, but for the F2-high peaks. Note that all four vertical scales have been reduced down by a factor of two to accommodate the much broader range of values in the high-ray observables.

CSF diurnal variability for 1-hop F2(low) peaks, 5-17 September 2015
Cloncurry (DSTO)-Laverton ELOISE array (DORS)
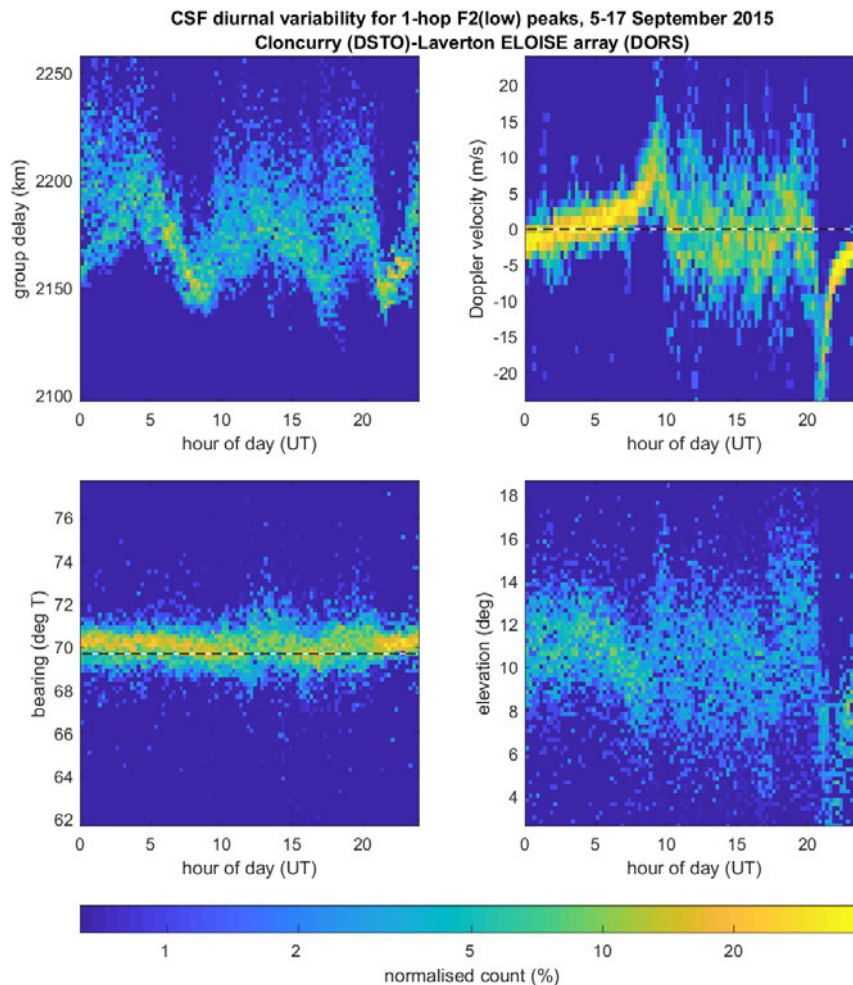
Figure 7.4. As in **Figure 7.2**, but for the Cloncurry to Laverton path. This path is similar in length to Humpty Doo to Laverton (both approximately 2000 km), but with a distinctly different great circle bearing (31 °T for Humpty Doo versus 70 °T for Cloncurry, calculated at the receiver site).

One potentially important factor that is masked in the above diurnal variability plots is frequency dependence across the three CSF channels, which sample different sections of the F2 trace. **Figure 7.5** presents the same F2-low peak data as in **Figure 7.2**, but rebinned in terms of the CSF frequency ratio. Each image panel has been normalised according to the distribution of all frequency ratios for this path, shown in **Figure 7.6**. The obvious group delay and, to a lesser extent, elevation dependencies in **Figure 7.5** are of no surprise, and simply reflect the curvature of the F2-low trace up to the nose of the ionogram (at which the frequency ratio equals one). There are also weaker but intriguing dependencies in the Doppler and bearing plots.

The trend towards positive/negative Doppler velocities at the upper/lower frequency ratios is observed on all Laverton paths and may well be an artefact of the CSF scheduling algorithm. That is, ratios above the highest target ratio (0.900) are more likely to be measured when the MOF is decreasing, which corresponds to declining ionisation levels, an increasing reflection height, and a negative/positive Doppler shift/ velocity; the opposite applies for ratios below the lowest target ratio (0.729). Such an effect has the potential to be exacerbated by the real-time CSF scheduling latency. In practice, there are comparatively few estimates at these extreme ends (as evident in **Figure 7.6**), so the impact on the earlier diurnal variability plots will be minimal.
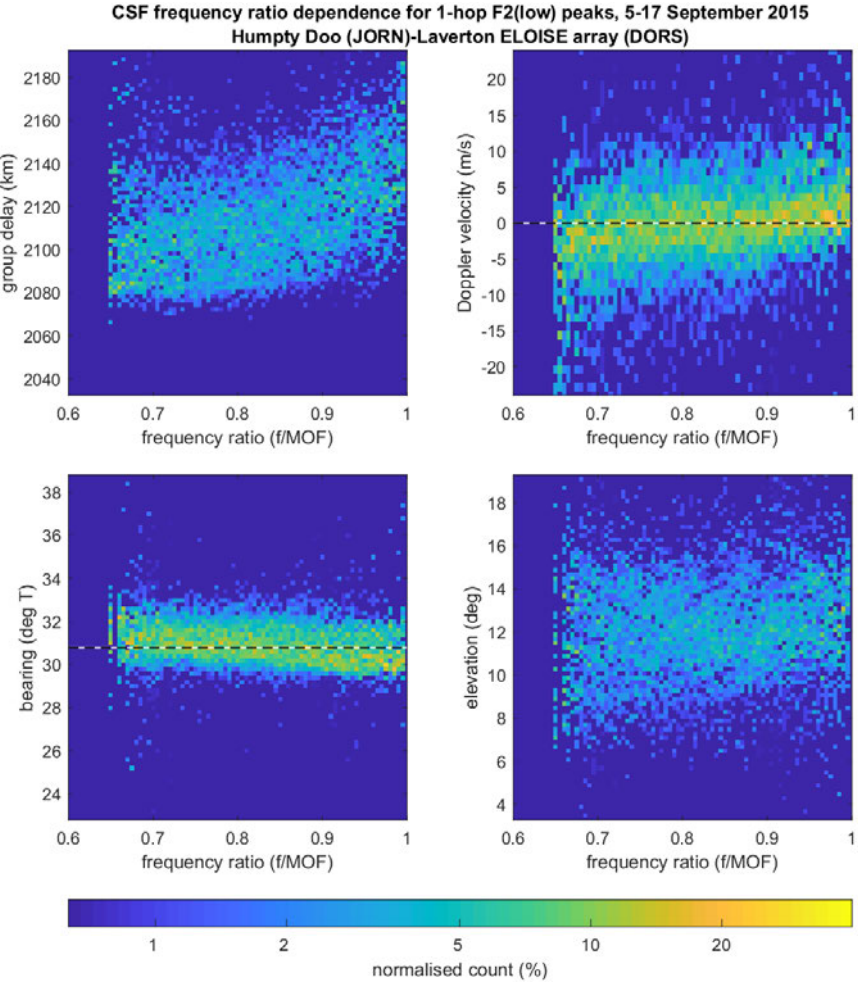
Figure 7.5. As in **Figure 7.2**, but rebinned as a function of the CSF frequency ratio (i.e. operating frequency divided by the maximum observed frequency). Although CSF target frequencies are fixed ratios of the (previous) MOF, channel availability and persistence rules in the scheduling algorithm (see **Section 4.2.1**) mean that a continuum of ratios between about 0.65 and 1 are recorded in practice.
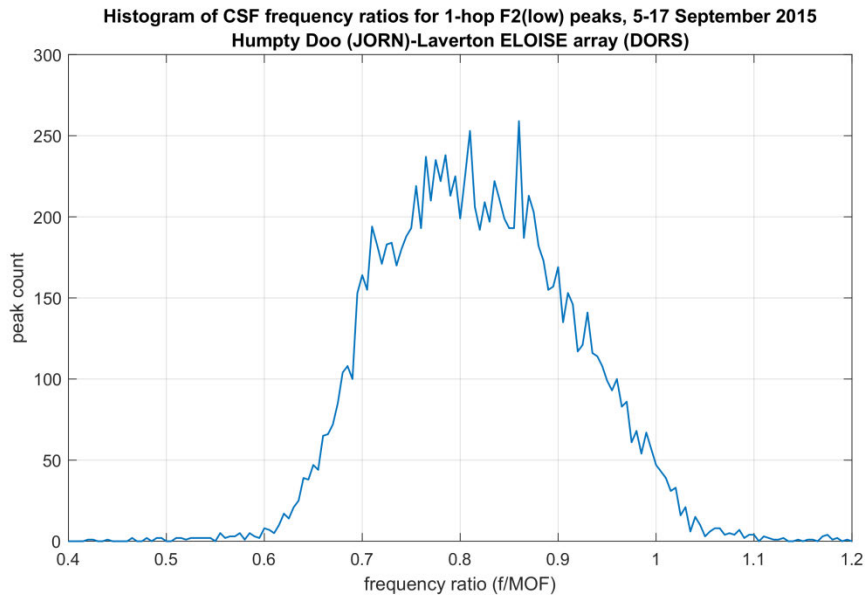
Figure 7.6. Distribution of CSF frequency ratios in the peak data set contributing to **Figure 7.5**.

The variation in bearing with frequency ratio, while very subtle, is only evident on the paths that are closest to north-south alignment. Like the pattern of diurnal variability, it is difficult to rule out instrumental (array calibration) effects on this scale, but a geophysical cause such as the anisotropy due to the geomagnetic field is suspected. For example, higher ratios beyond the O/X trace cross-over point are increasingly likely to represent the X-mode, which in theory should be deflected towards the magnetic equator [Davies, 1990, p. 200]. The opposite is true for the O-mode at lower ratios.

7.3   Comparisons against conventional ionospheric modelling

Conventional ionospheric sounders have long been used to construct parameterised models of ionospheric electron density for the purpose of propagation analysis and predictions. Such models are reasonably well regarded and can be considered a suitable benchmark for comparison against the ELOISE AoA observables, in terms of their representation of medium- to large-scale ionospheric variability. As in the ELOISE on-board signal processing, a vertical electron density profile is typically fitted to each conventional (amplitude-only) ionogram image, either directly or via trace extraction,

151

and this becomes the primary input to data-assimilative models such as the JORN Real-Time Ionospheric Model (RTIM). The key assumption is that horizontal gradients have a limited effect on group delay, and so, in the absence of AoA, each ionogram can be fitted assuming a spherically-symmetric profile in the vicinity of the midpoint. Horizontal gradients are thus only introduced into the model as a result of spatial mapping across a wide network of sounder sites.

The ionospheric model considered in this study is the Gardiner-Garden Ionospheric Model (GGIM). In short, the GGIM takes fitted profiles from oblique and vertical incidence sounders as its input, and outputs a set of spatial maps representing each of the ten JORN multi-segment quasi-parabolic (MQP) profile parameters, and their uncertainties, on a regional latitude/longitude grid or sampled at the sounder sites themselves. The JORN MQP profile is described further in **Appendix D**; see also Gardiner-Garden et al. [2018]. Each spatial map is independently constructed as the summation of an empirical (monthly-median) climatology field, a large-scale anomaly field, and a residual (spatially-uncorrelated) anomaly field.

The GGIM inputs are a combination of DORS (Digital Oblique Receiving System) OIS [Ayliffe et al., 2019], PRIME (Portable Remote Ionospheric Monitoring Equipment) VIS [Harris et al., 2016], and DPS-1 (Digisonde Portable Sounder) VIS [Reinisch et al., 1997]. The OIS are scaled using the DST-IIP (Ionogram Image Processing) algorithm [Heitmann & Gardiner-Garden, 2019], while the VIS are scaled using a combination of ARTIST-4 [Reinisch et al., 2005], PRIME on-board software [Harris et al., 2016], and DST-IIP. After temporally smoothing the input parameters, the monthly-median climatology is subtracted off, and the large-scale anomaly is estimated by fitting a planar surface to each gridpoint (i.e. assuming locally linear gradients) based on an inverse-distance weighted combination of all sounder sites over a ±15 min time window [Gardiner-Garden et al., 2019]. Successive model updates are carried out on the same cadence as the 3.75 min sounder revisit interval. While functionally similar to the JORN RTIM, the GGIM algorithm has been developed specifically for offline, batch-mode analysis, in the absence of real-time data latencies.

A single geographic snapshot of the full GGIM (i.e. climatology + large-scale anomaly + residual anomaly) is shown in **Figure 7.7**. In this depiction, all ten parameter maps have been combined to construct a grid of electron density profiles, with the true height

at a fixed plasma frequency of 2.5 MHz converted to an equivalent virtual height of reflection by way of a simple no-field analytic ray trace. During the ELOISE experiment, the density of sounder sample points over Australia (indicated by the red dots) was substantially higher than usual, enabling a greater proportion of medium-scale ionospheric variability to be captured by a conventional model such as the GGIM. As for this particular example, the dominant feature is a large-scale horizontal gradient running from north to south.

In the first sub-section that follows, the input MQP profile fits will be evaluated for their ability to predict (via ray tracing) the CSF Doppler variations observed, on an individual sounder basis. The second sub-section will then look at the GGIM output (i.e. incorporating up to 68 conventional sounder inputs) and compare both Doppler and AoA predicted estimates against observations from the ELOISE arrays.
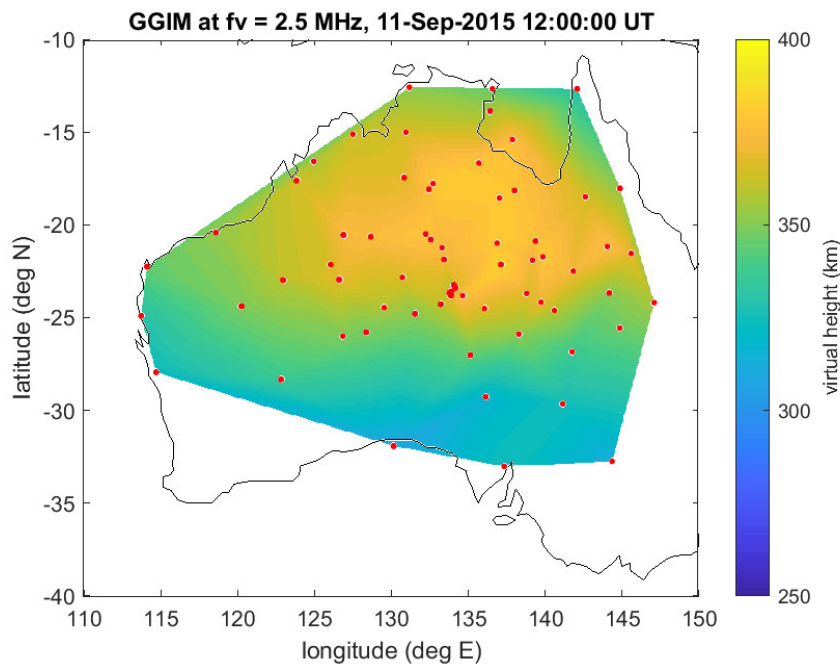
Figure 7.7. A GGIM map of the iso-ionic (no-field) reflecting surface for the night-time F2 layer at a plasma frequency of 2.5 MHz. The input sounder sites are indicated by the red dots; this may include up to 68 VIS locations and QVIS/OIS midpoints, with a median separation distance of 270 km (roughly 3° in latitude/longitude) with respect to the nearest neighbour. This places a lower limit on the spatial scale of irregularities that can be characterised by the model (notionally about 500 km), with aliased energy from smaller-scale disturbances contributing to the unresolved residual anomaly.

7.3.1 *Parameterised profile fits*

The Doppler shifts captured in the ELOISE CSF observations are a measure of temporal variability in the ionosphere, being proportional to the time rate of change of the phase path. This relationship can be expressed as

$$f_{id} = -\frac{f}{c}\frac{dP}{dt},$$

(7.2)

which is just a manipulated form of equation (7.1).

Non-zero Doppler shifts arise due to both time evolution of the ray path (i.e. motion of the reflection point) and temporal changes in the refractive index along the ray [Bennett, 1968; Dyson, 1975]. The former component is often assumed to dominate when there is bulk motion of the ionospheric structure without a substantial change in shape, such as in the presence of TIDs, while the latter component accounts for the underlying electron density fluctuations, which are usually more slowly varying except around dawn and dusk. The electron density itself may change from any combination of plasma transport, production, and loss; in the situation where production and loss processes are in equilibrium (e.g. in the middle of the day or night), this leads to the concept of a Doppler-derived "drift" velocity, which will be considered later in **Section 7.5**.

For a single quasi-parabolic (QP) layer [Croft & Hoogasian, 1968], the Doppler component due to electron density changes only (assuming spherical symmetry) can be parameterised in terms of the rates of change in the individual layer parameters:

$$f_{id}^{*} = -\frac{f}{c}\left(\frac{\partial P}{\partial f_c}\frac{df_c}{dt} + \frac{\partial P}{\partial h_m}\frac{dh_m}{dt} + \frac{\partial P}{\partial y_m}\frac{dy_m}{dt}\right),$$

(7.3)

where f_c, h_m and y_m represent the critical frequency, peak height and semi-thickness of the layer, respectively. Analytical expressions for the partial derivatives of phase path with respect to each QP parameter have been derived by Boldovskaya [1982]. Along the F2-low trace, the second and third terms in equation (7.3) are the largest contributors to f_{id}^{*}, whereas along the F2-high trace, the first term becomes increasingly important. This is illustrated in **Figure 7.8** for a simulated QP layer with f_c = 6 MHz, h_m = 300 km, y_m = 100 km, $\frac{df_c}{dt}$ = 0.5 kHz/s, $\frac{dh_m}{dt}$ = -20 m/s, and $\frac{dy_m}{dt}$ = -10 m/s, evaluated

154

over a 1-hop path length of 2000 km. Such an example is broadly representative of an F2 layer around dawn. As the F2 height and thickness parameters are usually correlated (in a positive sense), the second and third terms (red and yellow lines) tend to have opposite signs and thus partially cancel each other. Refer to Lynn [2009] for an application of this single-layer Doppler estimation technique to TID characterisation.

In reality, of course, a single layer is insufficient to describe the daytime ionosphere. Instead, having fitted a three-layer (E/F1/F2) electron density profile to the group delay trace in each ELOISE ionogram (part of the on-board processing described in **Section 4.2**), it is possible to synthesise the phase path using QP analytic ray tracing (ART) equations for the full multi-segment profile [Dyson & Bennett, 1988; Bennett et al., 1991]. Thus, the (spherically-symmetric) ionospheric Doppler f_{id}^* between successive 3.75 min ionogram observations was able to be estimated and compared to the measured Doppler. A weighted smoothing filter (±7.5 min) was applied to each QP parameter prior to ray tracing, in order to suppress fitting noise and interpolate across short outages or where a suspect F2 fit status was reported [Heitmann & Gardiner-Garden, 2019].
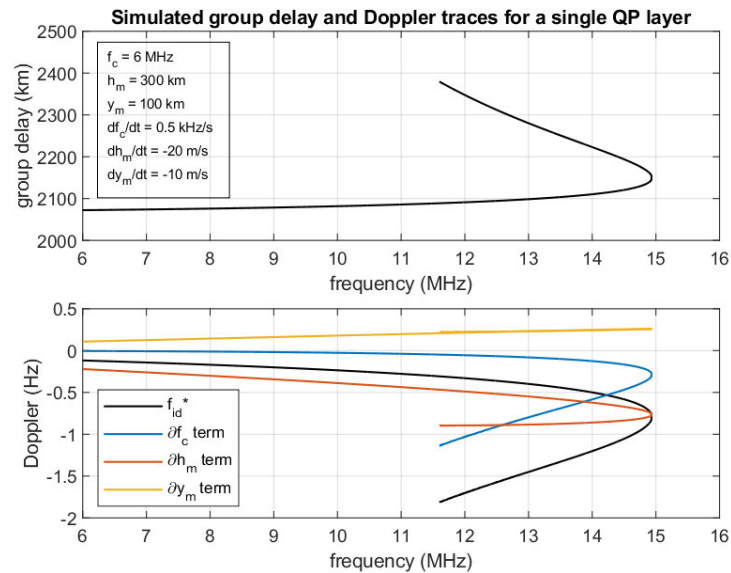
Figure 7.8. Simulated ionogram traces for a single time-varying QP layer. The top panel shows the group delay trace, assuming spherical symmetry and no geomagnetic field, while the bottom panel shows the Doppler variation, both as a total sum (black line) and separated out into each of the three partial terms in equation (7.3) (coloured lines).

Day- and night-time examples of synthetic trace slices and Doppler estimates from this quasi-parabolic spherically-symmetric model (QP-SSM) are provided in **Figure 7.9** and **Figure 7.10**, respectively. These cases were selected for having a particularly stable CSF frequency, making F2 high-ray dynamics simpler to interpret. The latter figure, in particular, is for the same path and period as depicted in **Figure 7.1**. Both F2-low (blue lines) and F2-high (red lines) are plotted, with Doppler computed as a centred finite difference of the phase path over ±3.75 min. Overlaid on these are the observed CSF peak parameters at the same time-varying frequency, classified as F2-low (blue crosses) or F2-high (red crosses) based on their group delay relative to the synthetic trace. Many of the jumps in the group and phase delay of CSF peaks are caused by small changes to the CSF frequency, though the Doppler synthesis compensates for this (i.e. each finite difference is calculated at a single frequency only).

Note that in keeping with wider usage, the terms "group path" and "phase path" notionally refer to synthetic estimates, derived from ray tracing integrals, while "group delay" and "phase delay" refer to observations, based on measured time or phase offsets. However, each pair fundamentally represents the same physical quantity, so the distinction is largely unnecessarily.

It is of interest to point out the following:

- Under the assumption of a mirror reflector, phase path and group path are the same as the geometric path, but according to Fermat's principle this is not true in general [Bennett, 1967; Nickisch et al., 2006; 2007]. Phase path (delay) is always less than group path (delay) for a medium with a refractive index less than one; the difference is especially large for the high ray (red) in **Figure 7.9** and **Figure 7.10**. It is also plain to see that the rate of change of group path (range rate) and phase path (Doppler) are not the same for ionospheric propagation, as would be the case for a mirror model or, indeed, line-of-sight backscatter off an air or surface target. Note that phase delay is only measurable in a relative sense, with an integer wavelength ambiguity, and is therefore not plotted.

- While F2-low Doppler shifts are confined to ±0.5 Hz in these two examples, the F2-high modes are frequently at least double this. The same relationship was seen earlier in comparing **Figure 7.2** (F2-low) and **Figure 7.3** (F2-high). When the CSF frequency is furthest from the MOF (e.g. around 0130–0330 UT in **Figure 7.9**),

F2-high mode Doppler amplitudes in the range of 1–2 Hz are both estimated and observed (at least, where the peak SNR remains detectable). At 13.5 MHz, this corresponds to Doppler speeds of 20–40 m/s, or vertical ionospheric motion of 10–20 m/s assuming a spherical mirror model. Such behaviour is exactly as illustrated in **Figure 7.8**.

- The two magneto-ionic (O/X) polarisation components are synthesised using a first-order model for geomagnetic field effects [Bennett et al., 1991; Chen et al., 1990; 1992], assuming an ideal Earth-centred dipole. This is applied as a frequency and delay perturbation to the no-field solution from the ART algorithm and serves to capture the bulk of the dependence on path location and geometry in a fast, analytical approach. In theory, the use of an effective no-field frequency will give rise to a frequency displacement in Doppler variations for the synthesised O/X components [Lynn, 2008], yet in general the O/X Doppler separation in these figures is much less apparent than in group or phase delay. On the high ray, this displacement should manifest as reduced Doppler amplitudes for the O-mode compared to the X-mode, given that the effective frequency will land further up the trace, at a greater apogee height for the X-mode (and noting the frequency-dependence in **Figure 7.8**). Indeed, this is clearly observed at the more extreme Doppler values (e.g. 0100–0200 UT in **Figure 7.9**), where the larger Doppler amplitudes, both positive and negative, are consistently from the X-mode.

Although it is not at all surprising to see excellent agreement between synthetic and observed group delay, since this is the domain in which the profile is fitted, it is pleasing to also see good agreement in Doppler. At times when there is a poorer match in Doppler, this may come down to either inadequacies in the profile parameterisation (e.g. inability to represent TIDs under the assumption of spherical symmetry), or smaller-scale variability that is under-sampled by the 3.75 min revisit interval. As an example of the former case, the synthesised F2-high Doppler excursion from 1630–1700 UT in panel (d) of **Figure 7.10** (red lines) is barely recorded in the observations (red crosses), and appears to be the artefact of a TID-induced ripple that descends through the F2 trace over a sequence of ionograms; a slight oscillation in the observed group delay (red crosses) of panel (b), completely missed by the fitted profile (red lines), is evidence of this. It should be emphasised that this is not a direct failure of the

fitting process itself, but simply reflects the limited number of MQP profile parameters, deliberately chosen to ensure reliable reproduction of the broader trace features only [Heitmann & Gardiner-Garden, 2019].
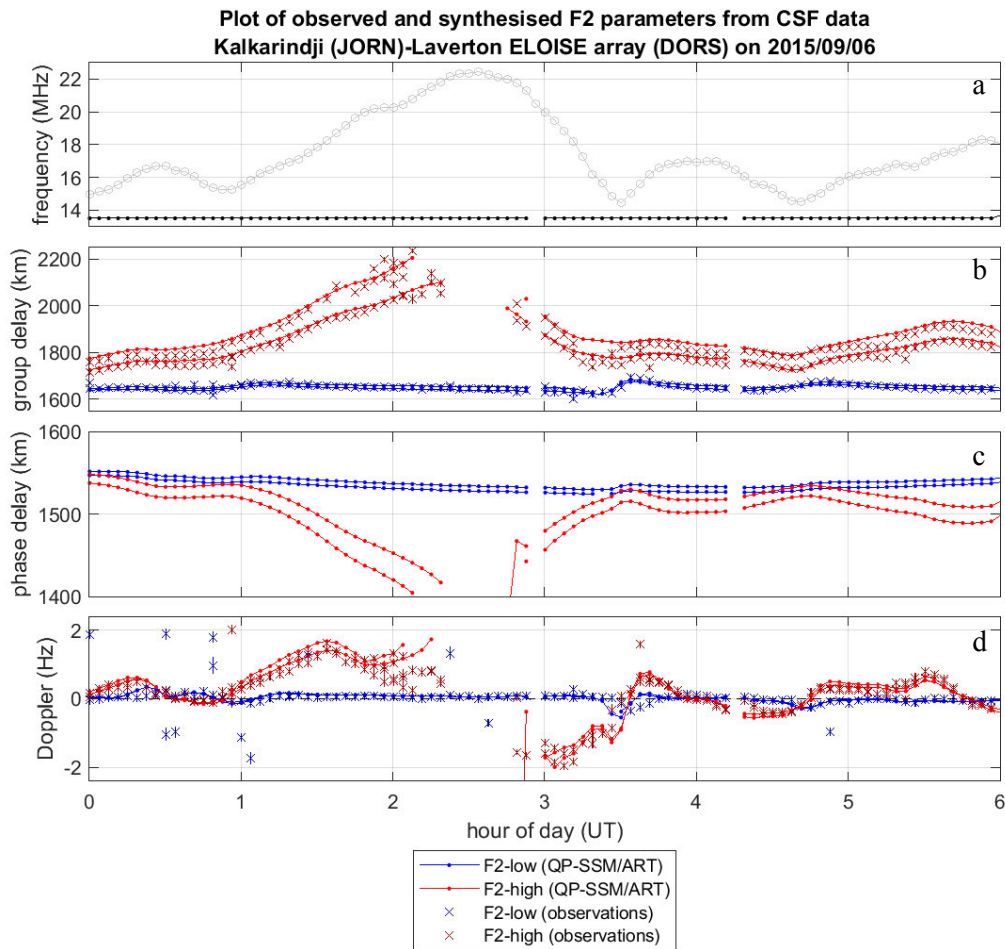
Figure 7.9. Daytime comparison of synthetic F2 trace slices and Doppler estimates against CSF observations. Synthetic results have been derived from the quasi-parabolic spherically-symmetric model (QP-SSM), fitted to each ionogram, in conjunction with analytic ray tracing (ART). The four panels, from top to bottom, show (a) the CSF frequency (with F2 MOF in grey), (b) group path (delay), (c) phase path (delay), and (d) Doppler, each as a function of time. The blue and red lines/crosses represent F2-low and F2-high, respectively.
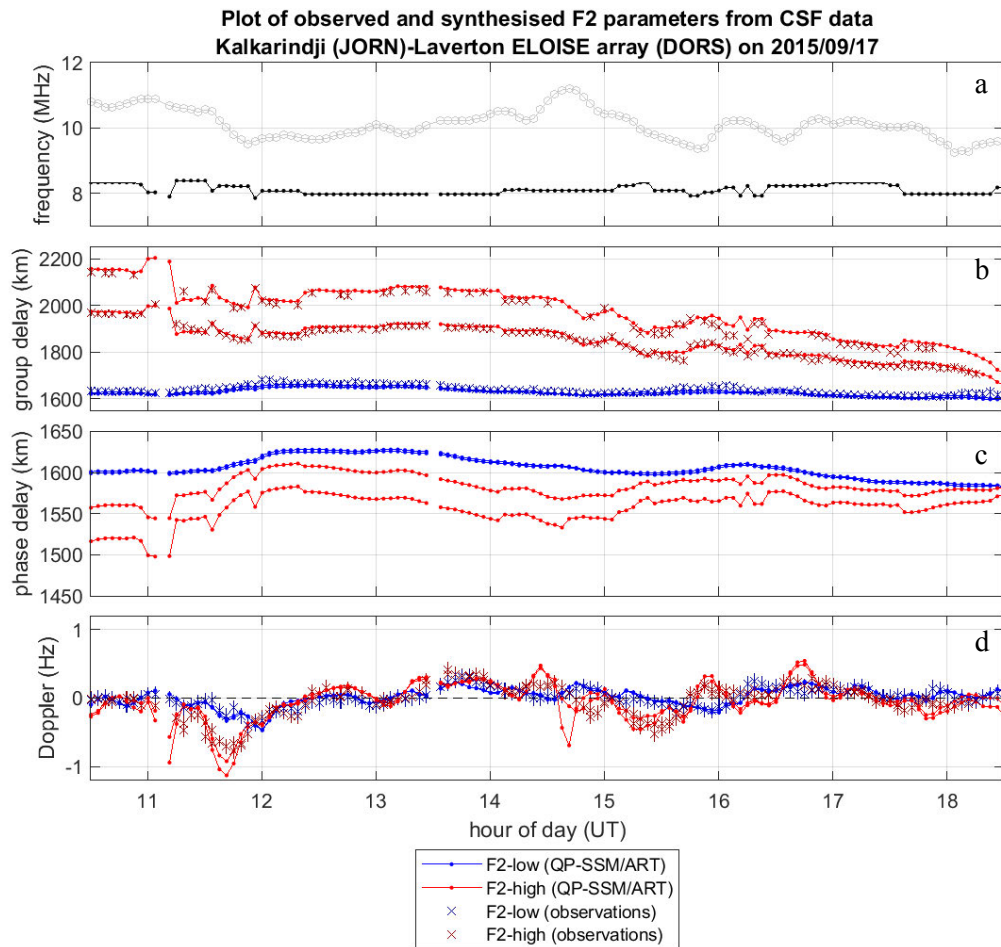
Figure 7.10. Night-time comparison of synthetic F2 trace slices and Doppler estimates against CSF observations. The plot layout is the same as in **Figure 7.9**, and depicts the same observations as in **Figure 7.1**, but for both F2-low and F2-high.

Turning now to aggregated results over the entire period 5–17 September 2015, **Figure 7.11** (Humpty Doo to Laverton) and **Figure 7.12** (Cloncurry to Laverton) show the synthesised diurnal variability in group delay and Doppler velocity for F2-low in the same format and for the same paths as the top two panels of **Figure 7.2** and **Figure 7.4**, derived from CSF observations. The key features from the observations, particularly the Doppler excursions around dawn and dusk, are clearly reproduced by ray tracing in both group delay and Doppler velocity. This confirms that the MQP profiles that form the inputs to the GGIM spatial maps reflect reality in these two dimensions. Naturally a spherically-symmetric model cannot accurately predict AoA variations, so the bearing and elevation panels are omitted for the synthetic QP-SSM results presented here, but will be considered in the next sub-section using the GGIM.
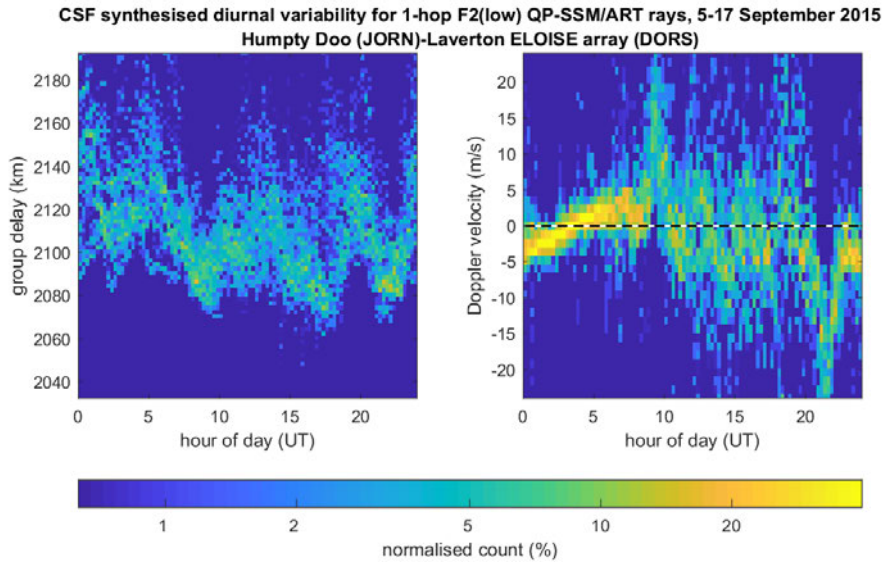
Figure 7.11. Diurnal variability in the 1-hop F2-low synthetic estimates (QP-SSM/ART), for the Humpty Doo to Laverton path (cf. observations in **Figure 7.2**).
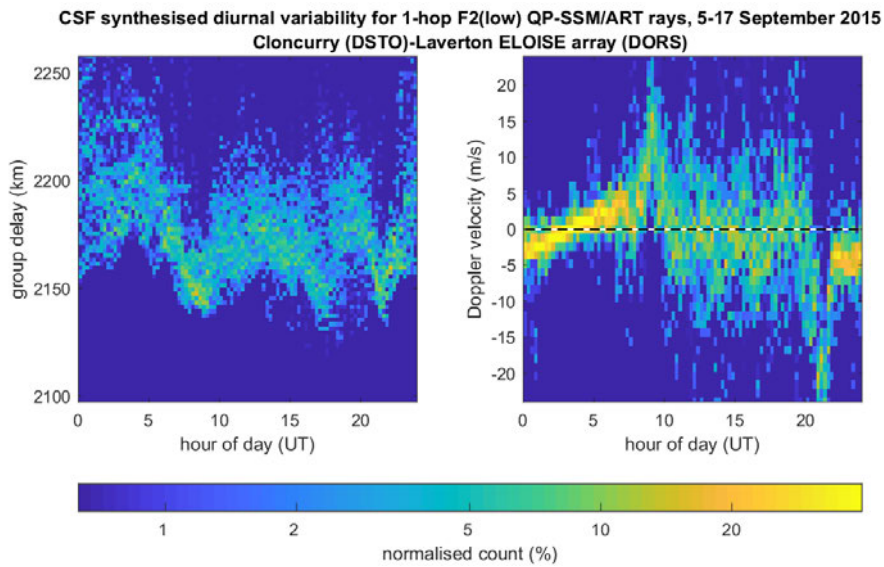
Figure 7.12. Diurnal variability in the 1-hop F2-low synthetic estimates (QP-SSM/ART), for the Cloncurry to Laverton path (cf. observations in **Figure 7.4**).

7.3.2 Large-scale spatial maps

As outlined earlier, the GGIM takes MQP profiles from up to 68 conventional VIS and OIS paths across Australia and constructs large-scale spatial maps of the ten QP parameters describing the E, F1 and F2 layers. In conjunction with 3D magneto-ionic numerical ray tracing (NRT) and ray homing, off-angle perturbations due to horizontal gradients (on the scale of the sounder midpoint separations) can be synthesised on the ELOISE oblique paths of interest and compared to the AoA observations. The HF propagation toolbox PHaRLAP, by Dr Manuel Cervera [Cervera & Harris, 2014; Cervera, 2019], was used for this purpose, incorporating the International Geomagnetic Reference Field (IGRF) with 2015 coefficients [Thébault et al., 2015]. To capture the typical range of 1-hop AoA variations, NRT rays were fired off over a $\pm 8°$ azimuthal fan, centred on the great circle bearing, and a broad elevation fan of at least $60°$, depending on the path length. Rays were uniformly distributed in solid angle, using the method described by Cervera & Harris [2014], with roughly one degree separating adjacent rays in azimuth and elevation.

Note that the GGIM is not a direct interpolation scheme and the model output at the ELOISE path midpoint will not necessarily be identical to the input sounder data before or after temporal filtering, especially if there is another path midpoint close by. It incorporates fixed weighting functions derived from intersite parameter covariances and forms a linear combination of the MQP profile data from all sites. Uncertainties are mapped in a similar way, from the input parameter errors. The full GGIM field, used throughout this chapter, includes the residual anomaly component, and thus retains far more of the medium-scale temporal variability present in the input fits, considerably improving the match to observations; however, estimates of the local horizontal gradients are derived from the large-scale anomaly only.

A time series of CSF observables, spanning a complete 24 hours, is shown in **Figure 7.13**, with synthetic estimates from both QP-SSM (see previous sub-section) and GGIM models overlaid in blue and cyan, respectively. Only the F2 low ray is plotted, in part for sake of clarity, and also due to the fact that the high ray is more difficult to interpolate without a significant increase in the ray sampling density and, hence, memory and computational requirements. Again, bearing and elevation estimates are omitted for the QP-SSM results, as spherical symmetry allows for no representation of
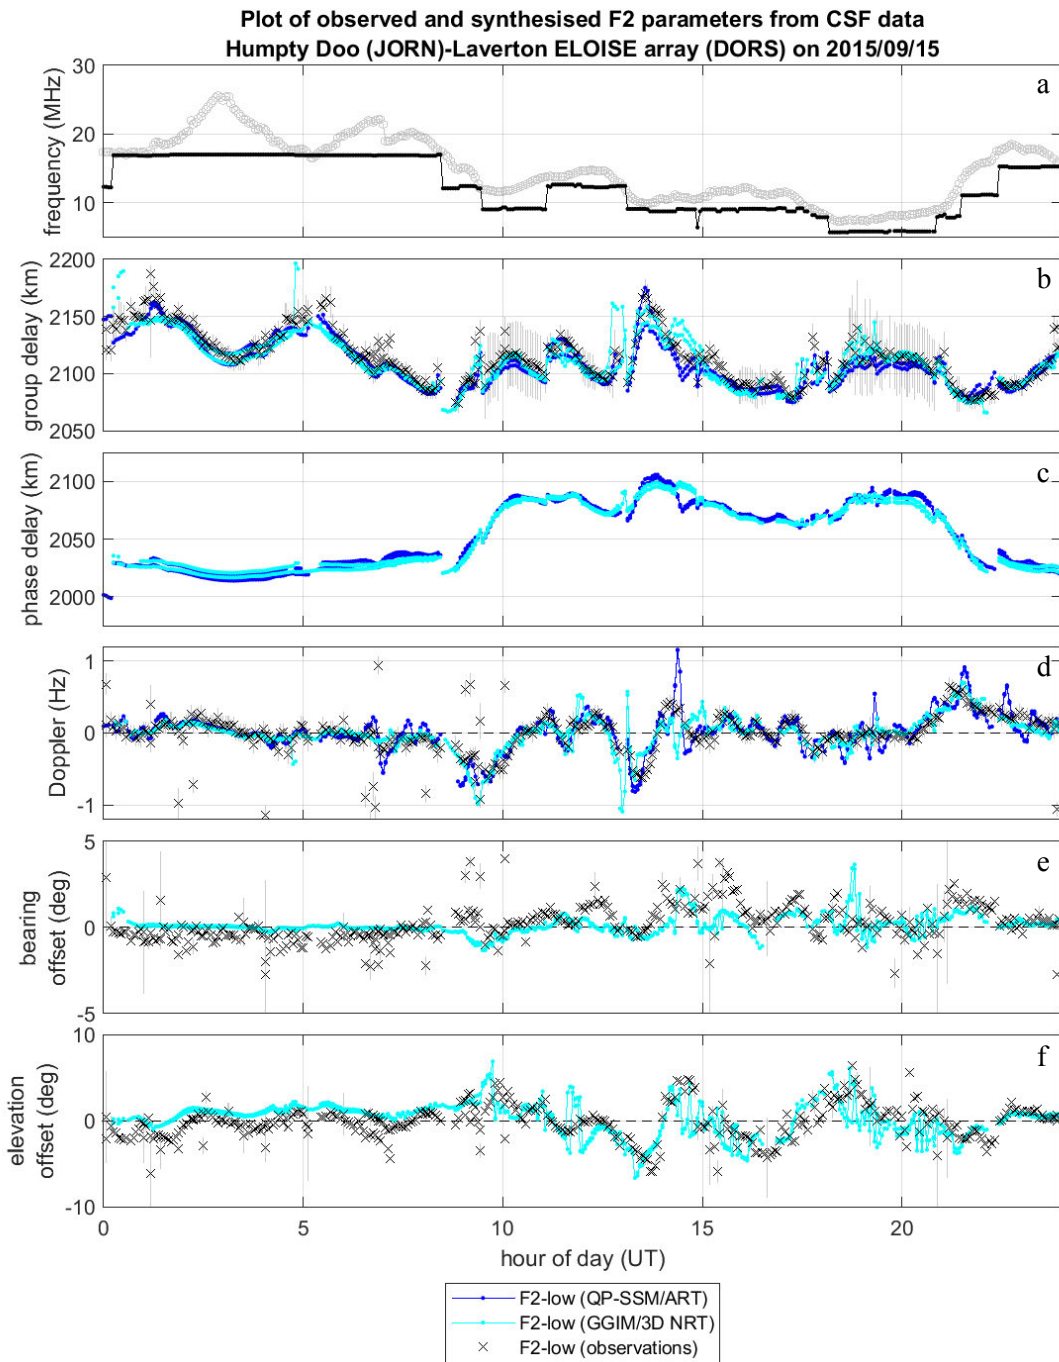
Figure 7.13. Time series of multi-dimensional AoA peak data, on a single frequency-agile channel measuring the 1-hop F2-low mode from Humpty Doo to Laverton. The true observations (fitted peaks) are shown as black crosses with uncertainty bars (plus/minus one standard deviation), while synthetic estimates are shown in blue (for QP-SSM/ART model) and cyan (for GGIM/3D NRT model). The two models are described in **Sections 7.3.1** and **7.3.2**, respectively. From top to bottom the panels show (a) CSF frequency (with F2 MOF in grey), (b) group delay, (c) phase delay (synthetic only), (d) Doppler, (e) bearing offset (with respect to the great circle bearing), and (f) elevation offset (with respect to an equivalent untilted spherical mirror).

162

gradients. With numerous CSF frequency changes over the course of a day, there are naturally some abrupt jumps in the data sequences, although the GGIM estimates tend to be more susceptible as a result of numerical sensitivity to inhomogeneities in the residual anomaly.

On the whole, the observed pattern of hour-to-hour variability (black crosses) is captured reasonably well by the synthetic GGIM estimates in **Figure 7.13**. This includes not only the fitted group delay observable, but also the unfitted Doppler, bearing and elevation observables. Both QP-SSM and GGIM successfully reproduce the large-scale Doppler excursions (e.g. 0800–1400 UT), along with many of the medium-scale quasi-periodic patterns (e.g. 1500–1800 UT). There is a tendency for the QP-SSM fits to overestimate the Doppler magnitudes, but encouragingly, this seems to be largely absent from the GGIM-derived results (i.e. after spatial mapping). In terms of AoA offsets, the GGIM mostly tracks the observed perturbations, at least in sign, although often with reduced amplitudes in the bearing offset. This may reflect an under-representation of small- to medium-scale gradients in the GGIM.

Note that bearing and elevation have been presented as offsets, to draw out the influence of horizontal gradients, rather than just being overwhelmed by the first-order elevation dependence on group delay (i.e. from the equivalent mirror geometry). Elevation offsets are expected to be a little larger when reflection occurs near the peak of a layer (i.e. in the vicinity of a critical frequency cusp in the ionogram), since the mirror model approximation is not so good under these circumstances [McNamara, 1991, ch. 12.3; Dao et al., 2016]; this includes operating frequencies close to the F2 MOF. However, the larger excursions of many degrees are almost certainly due to gradient effects.

Diurnal variability plots for the GGIM with 3D NRT are shown in **Figure 7.14** (Humpty Doo to Laverton) and **Figure 7.15** (Cloncurry to Laverton), for comparison against the CSF observations in **Figure 7.2** and **Figure 7.4**, respectively. The similarities are quite remarkable in light of the fact that the GGIM does not assimilate any Doppler or AoA observations. Features surrounding the dawn/dusk terminators, along with the steady increase in Doppler velocity throughout the day, are almost exact replicas of the observations, and even the relative levels of scatter in the estimates from day to night are largely consistent.

As in **Figure 7.13**, it is the synthesised bearing in **Figure 7.14** and **Figure 7.15** that is most at odds with the observations. There is far less day-to-day variability (vertical scatter) in the GGIM (although some of the extra scatter may be a manifestation of CSF measurement errors), and the observed midday shift towards north on the Humpty Doo path is not at all present. However, the fact that the more east-west Cloncurry path encounters less variation in bearing across the aggregated day than the Humpty Doo path is reproduced by the model to some extent. The dawn/dusk bearing excursions, in particular, are noticeably smaller for Cloncurry (at 70 °T); this is to be expected, as the dominant east-west gradient at these times will be better aligned with the down-range direction of the path than for Humpty Doo (at 31 °T), and thus not so likely to impact the bearing (only the elevation angle).
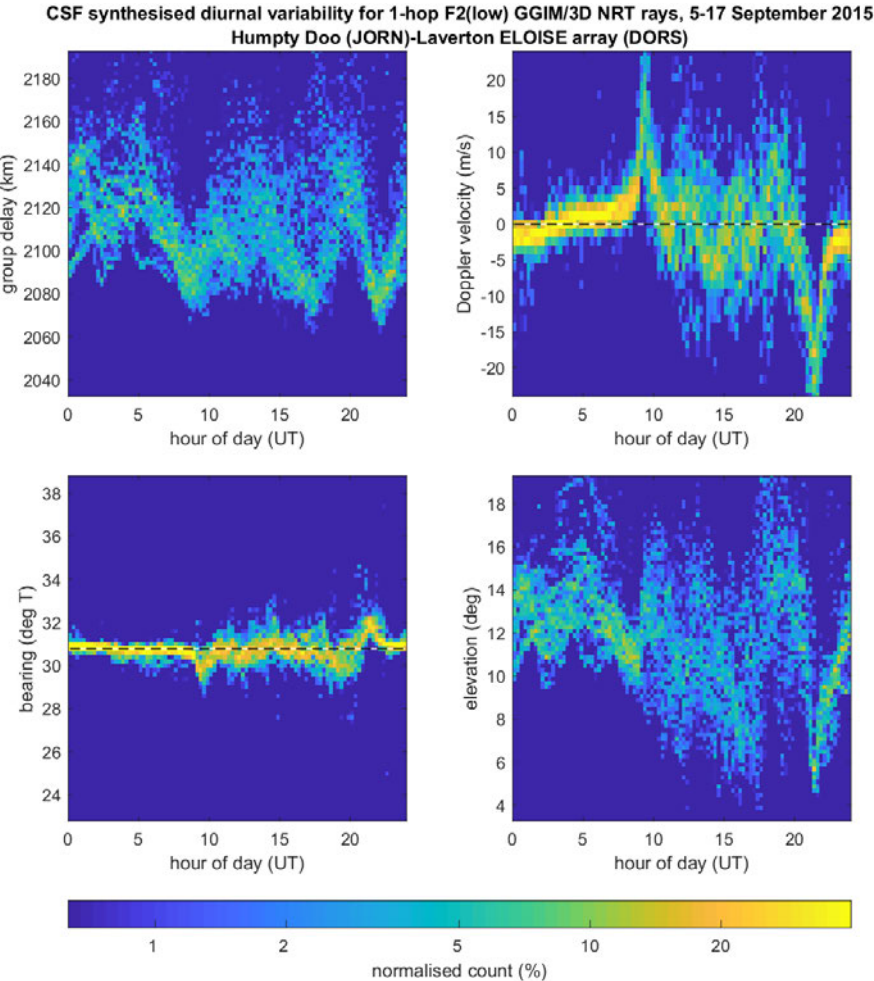
Figure 7.14. Diurnal variability in the 1-hop F2-low synthetic estimates (GGIM/3D NRT), for the Humpty Doo to Laverton path (cf. observations in **Figure 7.2**).
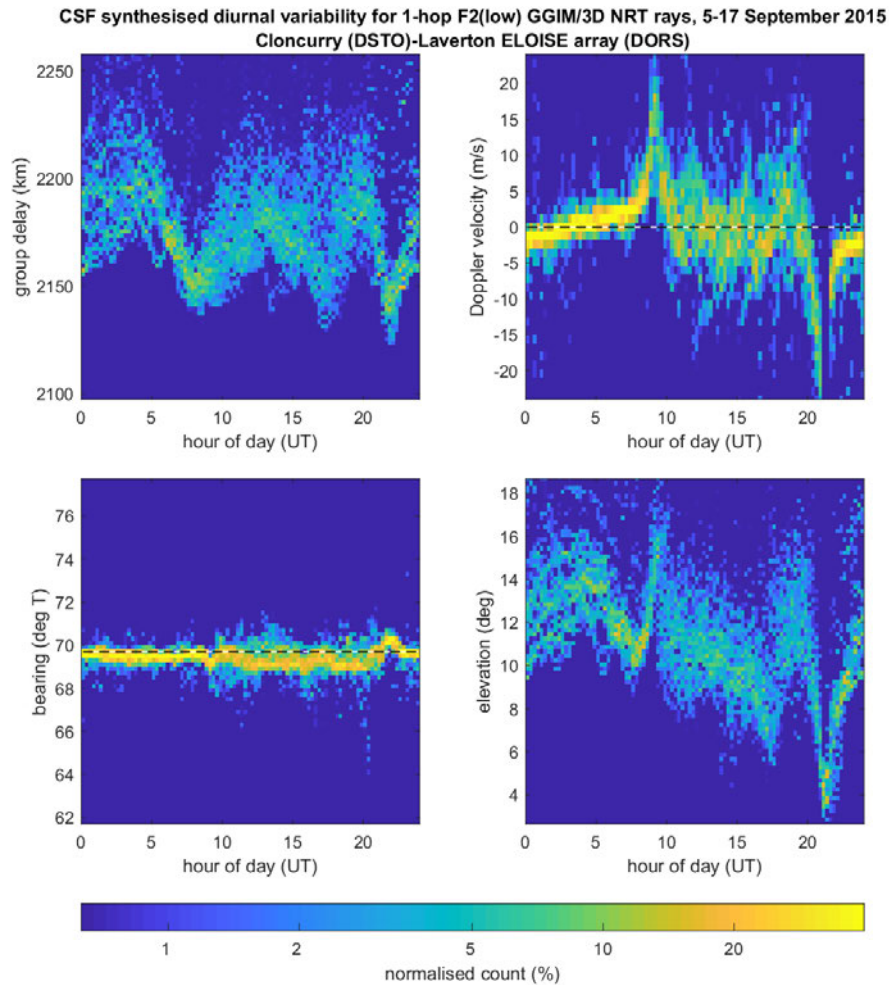
Figure 7.15. Diurnal variability in the 1-hop F2-low synthetic estimates (GGIM/3D NRT), for the Cloncurry to Laverton path (cf. observations in **Figure 7.4**).

7.3.3 Statistical model evaluation

While a comprehensive investigation of ionospheric modelling performance is outside the scope of this work, it is nevertheless of interest to compare the available data set of ELOISE CSF observations against the QP-SSM and GGIM models in a statistical sense. Of course, this not only captures modelling deficiencies, but also measurement errors; for example, due to unresolved multi-mode and array calibration issues. In order to provide a robust set of matched mode pairs, only the dominant (largest amplitude) 1-hop F2-low peak from each CSF dwell is retained, and this is differenced against the average of synthetic O and X components with the shortest group path, on the basis that O/X are rarely resolvable in the CSF observations.

165

Figure 7.16 and **Figure 7.17** present histograms of the differences in each of the CSF observables for the Humpty Doo and Cloncurry to Laverton paths; these notionally represent modelling errors if the peak estimates could indeed be considered the "truth", but in practice both model and measurement contain errors to some degree. Although implicitly dependent on the CSF waveform and processing parameters, recall from **Section 4.4** that typical peak estimation uncertainties in group delay, Doppler, bearing and elevation were on the scale of 8–15 km, 0.12 Hz, 0.1–0.2° and 0.2–0.4°,
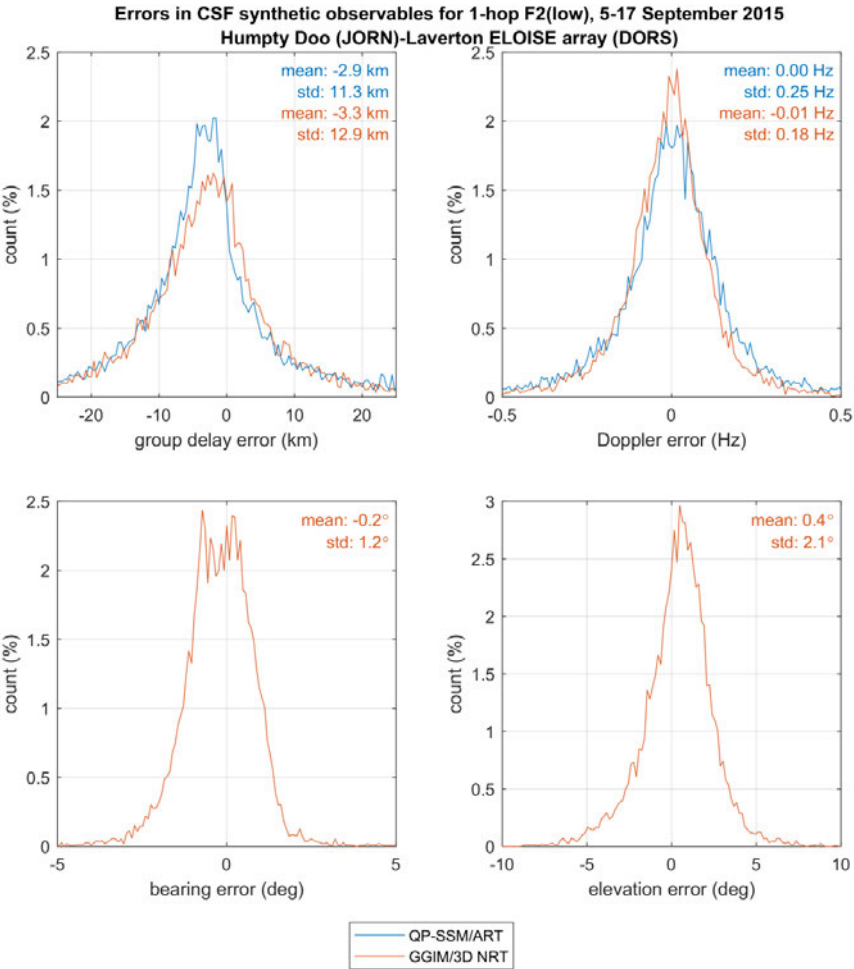
Figure 7.16. Histograms of differences between the 1-hop F2-low CSF observables and their synthetic counterparts, for the Humpty Doo to Laverton path. The "errors", plotted for both QP-SSM (blue) and GGIM (orange) over the period 5–17 September 2015, are defined as synthetic estimates minus observations, after AoA bias corrections. The top row of panels shows the error distributions of group delay (left) and Doppler (right), while the bottom row shows bearing (left) and elevation (right). Mean and standard deviation statistics are listed in the top-right corner of each panel.
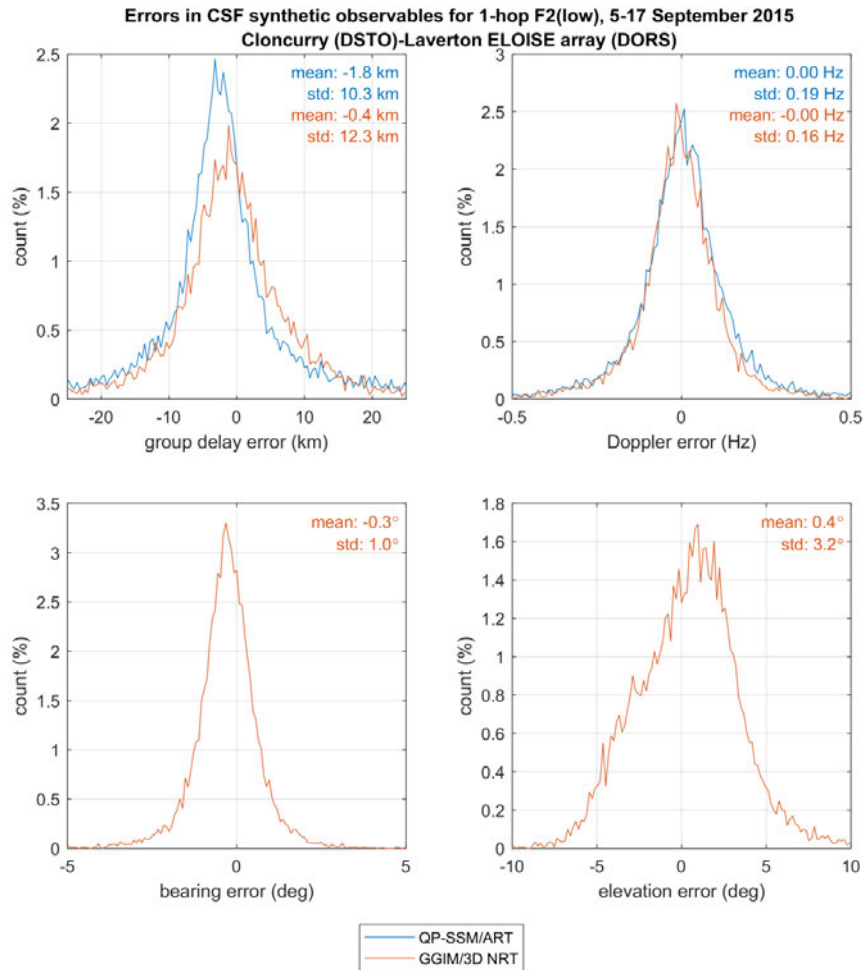
Figure 7.17. As in **Figure 7.16**, but for the Cloncurry to Laverton path.

respectively, for CSF data. For group delay and Doppler, these numbers are conspicuously similar to the standard deviations listed in the top-right corner of each histogram panel, indicating that the differences may in fact be as much measurement error as modelling error. In terms of model performance, this is quite an encouraging result. For bearing and elevation, the standard deviations are indisputably larger than the peak estimation uncertainties, pointing to a more definite modelling error that, despite being fairly modest in size, cannot be dismissed.

Although the AoA bias corrections derived in **Section 5.5** have been applied, residual elevation biases may still be present in the observations as a result of the inherent equivalent geometry assumption, so reported mean differences of a few tenths of a degree are not entirely unexpected. However, the bias of up to 3 km between modelled

and observed group delay is somewhat puzzling and deserves further investigation. Given it is comparable in size to the ionogram delay or virtual height resolution, it may be an artefact of using the edge rather than the centre of the ionogram traces to fit the electron density profiles. This approach is common to both ARTIST-4 and DST-IIP algorithms and affects the synthetic QP-SSM and GGIM observables. The possibility of some CSF peaks merging with additional off-angle returns at higher delays may also be a contributing factor.

7.4   Evaluation of ionospheric tilts

Horizontal gradients routinely exist in the ionosphere due to a number of factors, including spatial variations in solar forcing (particularly around the dawn and dusk terminators), ion diffusion along geomagnetic field lines (for example, giving rise to the equatorial anomaly), atmospheric winds and tidal oscillations, travelling ionospheric disturbances, and localised effects of extreme space weather events. In this context, "large scale" typically refers to spatial scales on the order of 500–1000 km, detectable by conventional ground-based ionosonde networks, while "medium scale" refers to structure over a more modest 100–500 km, often near the limits of the network's spatial sampling resolution. The impacts of gradients on HF AoA measurements are well understood in the HF direction-finding community [McNamara, 1991], with roughly one-third of the typical AoA estimation variance attributed to systematic (solar-related) gradients and two-thirds attributed to TIDs [Jones & Reynolds, 1975].

Ionospheric gradients can be characterised either by the "slope" of the iso-ionic contours (expressible as a tilt angle by taking the arctangent), or by the horizontal displacement of the ray apogee. The former is the more intrinsic ionospheric property, while the latter describes the measurable effect on HF propagation. Both will be a function of frequency, or equivalently height, with certain classes of disturbance such as TIDs having clearly defined signatures in these dimensions (i.e. manifesting as height-dependent phase shifts). Mid-latitude tilt angles will ordinarily be on the scale of less than 3° [Gardiner-Garden et al., 2019], which translates to apogee displacements on the scale of <20 km for cross-plane tilts or <300 km for in-plane tilts on any given 1-hop path, depending on the height of reflection. A specific gradient realisation can thus lead to vastly different apogee displacements for near orthogonal path bearings.

For a spherical Earth and stratified ionosphere, a first-order representation of gradients in the electron density iso-surfaces is based on shifting the centre of the spherical ionosphere away from the centre of the Earth (e.g. Folkestad [1968]; Dyson & Bennett [1992]). This scheme lends itself well to both analytic ray tracing through a non-concentric but spherically-symmetric quasi-parabolic segment profile, as well as the simpler calculation of geometric mirror reflection by a set of spherical (virtual) shells. To avoid making any assumptions about the underlying E/F1 electron density profile, the latter mirror approximation is adopted for interpreting the ELOISE AoA observations; that is, with the reflection point mapping approach for tilt estimation described earlier in **Section 4.5**.

Each fitted F2 peak from the AoA ionogram data (both low and high ray) is mapped independently, yielding its own off-centre ionospheric mirror. Local gradients in these mirror estimates are then binned by equivalent vertical frequency (i.e. notional plasma frequency), and averaged in the form of trimmed and weighted means. Corresponding standard deviations capture uncertainty in the gradient estimates. At any point on the Earth, the partial derivatives of the overhead mirror surface in the zonal (east-west) and meridional (north-south) directions give rise to a pair of tilt angles, measured with respect to the local horizontal. Although the tilt parameters are assumed to be global under this approach, in practice ionospheric refraction effects for a given 1-hop oblique path are constrained to a limited region surrounding the apogee, especially for low-ray propagation, so this does not generally pose a major concern.

Geomagnetic splitting is also neglected in this mirror approximation, noting that O/X returns are not automatically separated in the ionogram images. These unmodelled field effects have the potential to masquerade as tilts [Dao et al., 2016], but as illustrated in the simulations of **Figure 7.18**, the resultant AoA errors are fairly modest on the ELOISE oblique paths of interest, compared to typical measurement uncertainties of a few tenths of a degree (e.g. see **Table 4.1** in **Section 4.4**). In the top row of plots, a spherically-symmetric (tilt-free) profile model is specified (left) and synthetic F2-low group delay traces (right) are generated using 3D NRT. In the bottom row of plots, bearing (left) and elevation (right) offsets show the differences between the 3D NRT estimates and those assuming equivalent no-field propagation via an Earth-centred mirror ionosphere (i.e. in the same sense as the AoA offsets defined earlier). Note that

the worst-case scenario is encountered on very short paths, near the geomagnetic equator, directed at 45° to the field [Dao et al., 2016].

A similar AoA-based tilt study has recently been reported by Paznukhov et al. [2020], using Digisonde drift analysis data at vertical and near-vertical incidence [Reinisch et al., 2005; Kozlov & Paznukhov, 2008]. However, in this experiment, observations were at a single frequency only, and therefore lacked any information about the height dependence of tilts. The VIPIR/Dynasonde VIS system offers a more complete, height-dependent specification of tilts over a 3D electron density distribution as part of its NeXtYZ ionogram inversion technique [Zabotin et al., 2006], using a series of stratified plasma "wedges", and likely represents the leading approach for vertical-incidence data.
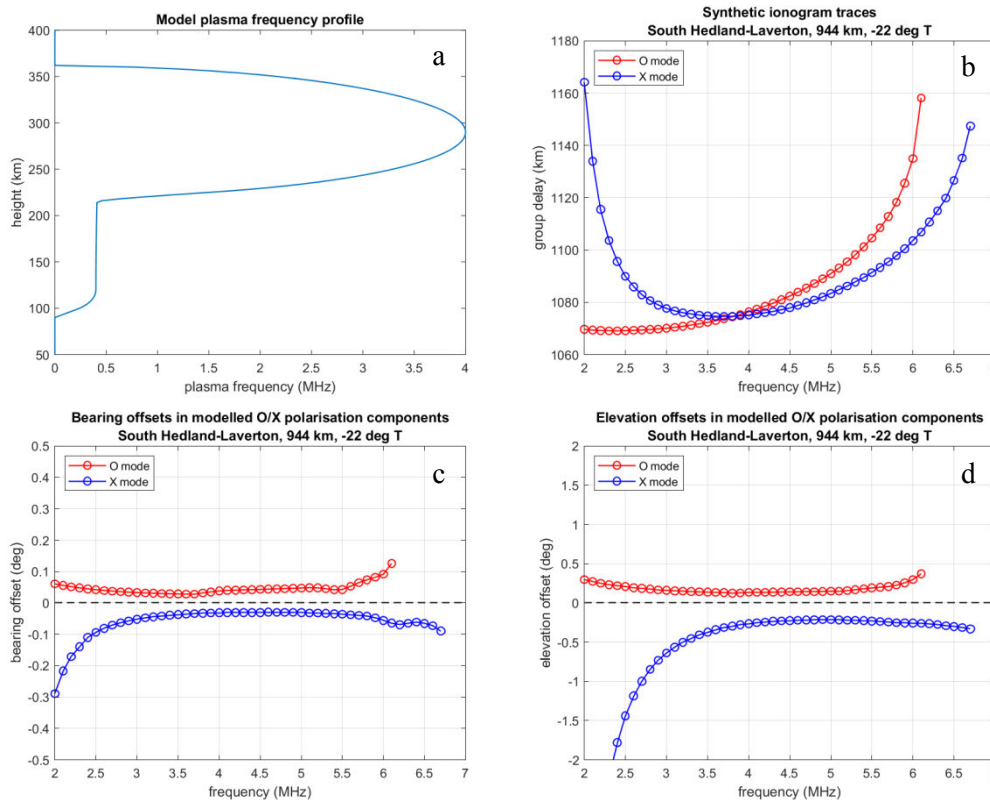
Figure 7.18. Synthetic 1-hop F2-low ionogram traces calculated by 3D numerical ray tracing through a representative night-time profile on the South Hedland to Laverton path (944 km). In theory, this is the worst case of all ELOISE AoA paths for geomagnetically induced deflections in AoA, owing to a combination of its short length and alignment to the field. Panel (a) shows the model plasma frequency profile, while (b), (c) and (d) show group delay, bearing offset and elevation offset, respectively, as a function of frequency for both O (red) and X (blue) polarisation components.

A sample of F2 tilt estimates from a complete day of ELOISE AoA ionograms is shown in **Figure 7.19**, alongside the midpoint virtual height profiles obtained from the peaks. The corresponding tilt uncertainties are shown in **Figure 7.20**; these are typically on the scale of 1° or less. It is noted that the uncertainties are simply a measure of sample variance in the gradient estimates, and do not include the effects of residual array calibration or mirror model errors. Instances of greater uncertainty (e.g. 15–16 UT) potentially flag when the peaks are too far horizontally dispersed to truly represent a common gradient. Nevertheless, the scale of uncertainties on the whole indicates a strong level of consistency across multiple mapped peaks. Although CSF data can be transformed in the same way, the broader set of frequencies sampled by the ionograms every 3.75 min lends itself better to populating all equivalent vertical frequency bins.
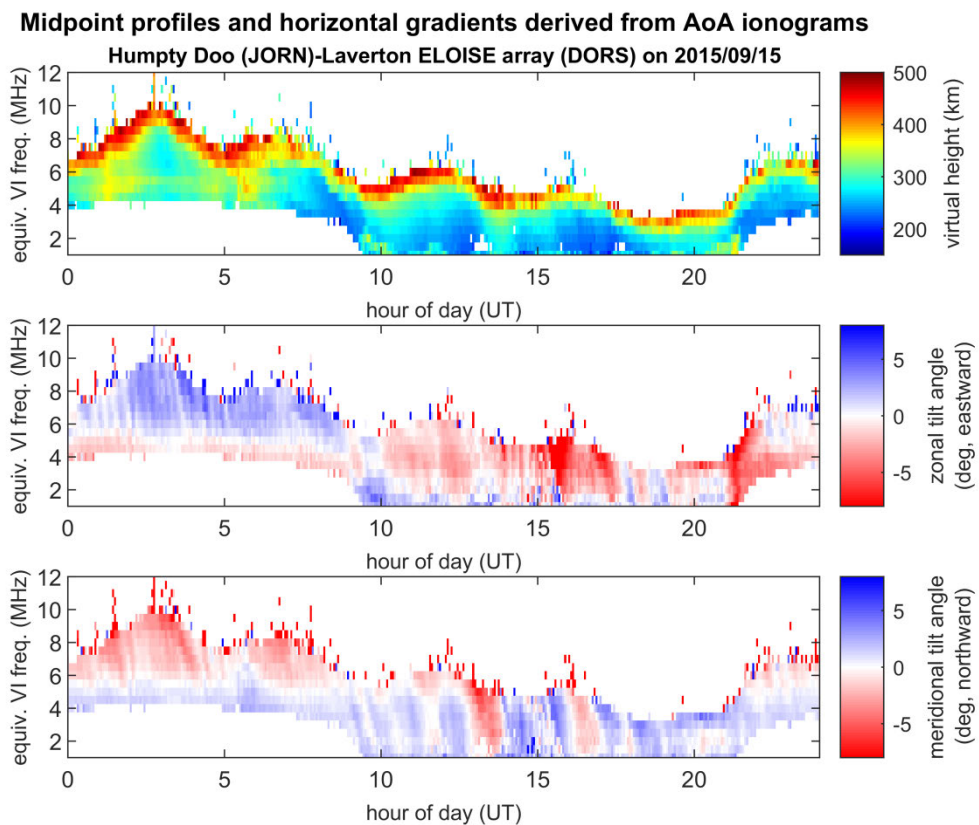
Figure 7.19. Time series of midpoint F2 virtual height profiles and horizontal gradients (tilt angles) derived from fitted ionogram F2-low peaks on the Humpty Doo to Laverton path. The path and period of observation are the same as for the CSF peak data in **Figure 7.13**. From top to bottom the panels show virtual height, zonal tilt and meridional tilt, each as a function of hour of day and equivalent vertical-incidence (VI) frequency (notionally representative of the plasma frequency).
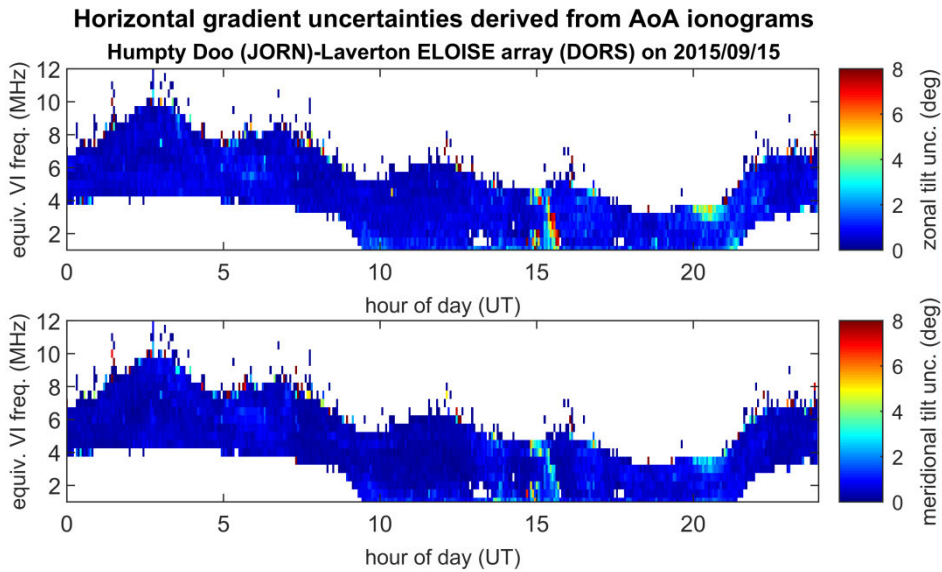
Figure 7.20. Time series of zonal and meridional F2 tilt uncertainties (top and bottom panels, respectively), corresponding to the gradient plots in **Figure 7.19**. These are derived from SNR-weighted sample standard deviations across the tilted mirror model estimates of all F2-low peaks whose equivalent VI frequency falls within a given 0.5 MHz bin.

Geophysical features in **Figure 7.19** are dominated by the terminator effects, particularly shortly after dawn (~21 UT) and, to a lesser extent, before dusk (~10 UT), when the zonal tilt angles tend to swing to negative and positive values, respectively. This is a statement that the iso-ionic F2 contours increase in height from the daytime to the night-time longitudinal sector. The tilt estimates also contain significant evidence of other medium- to large-scale variability that changes from day to day, with much of this resembling the classic TID-like signature of quasi-periodic (20–60 min) perturbations, descending in phase from upper to lower heights or equivalent VI frequencies. These features are most pronounced in the meridional tilt angles at night.

Just as in **Section 7.3** above, the GGIM offers a comparative representation of ionospheric gradients, derived from spatial mapping across a network of conventional OIS and VIS ionosondes. The resultant tilt angles, extracted from the GGIM at the AoA path midpoint, are plotted in **Figure 7.21**; this includes both F2 and the underlying E and F1 layers (not shown in **Figure 7.19**). Note that the GGIM tilt angles are calculated over a local Cartesian latitude/longitude grid with a finite difference approximation in
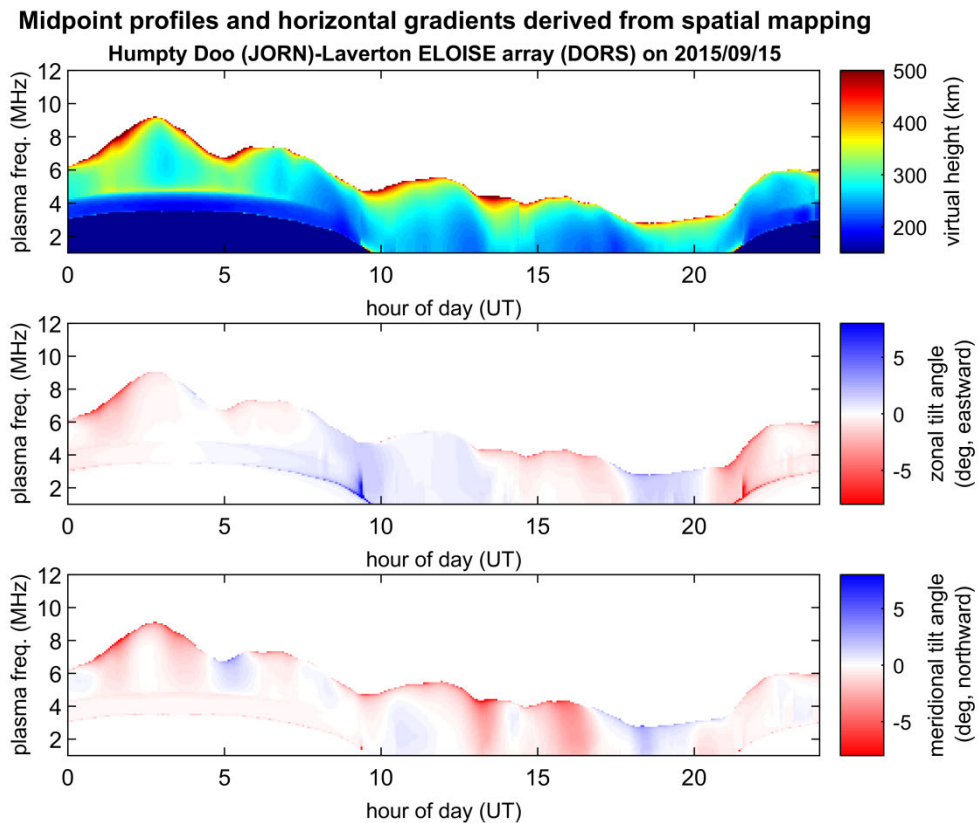
Figure 7.21. Time series of E/F1/F2 virtual height profiles and horizontal gradients derived from the GGIM spatial maps at the midpoint of the Humpty Doo to Laverton path (cf. observations in **Figure 7.19**, for F2 only).

true height, whereas those from the AoA observations use a spherical coordinate system and compute the slope of the tangent to the mirror (virtual height) surface at the midpoint. Despite these differences in their construction, one might still expect the two sets of tilt estimates to be reasonable analogues. In practice, however, the agreement between **Figure 7.19** and **Figure 7.21** is not so good, even in a broad qualitative sense.

In the case of the observations, it is suspected that the relatively long (1989 km) path length accentuates both calibration and mirror model errors, leading to apparent biases, especially during the daytime, that are not present in the GGIM. The impact of tilts will generally be more pronounced and thus measurable in the AoA of shorter paths [McNamara, 1991, p. 189]. The systematically positive (blue) zonal tilt angles and negative (red) meridional tilt angles from 0–9 UT at the higher daytime frequencies in **Figure 7.19** should therefore be viewed with some suspicion. The observations are also

affected to some extent by any variations in the gradients away from the midpoint, unlike the GGIM representation in **Figure 7.21**; this is one reason why the earlier comparison against ray-traced synthetic observables for the same path and period (**Figure 7.13**, at a single time-varying CSF frequency only) may have appeared more convincing. In addition, experience suggests that the medium-scale portion of the TID spectrum is only partially captured by conventional models such as the GGIM, due to the restricted number of profile parameters (i.e. degrees of freedom in the vertical dimension), critical sampling limits of the sounder network (in both space and time), and generous smoothing filters. This problem is further exacerbated by the fact that the residual anomaly is omitted when calculating the GGIM gradient estimates; despite being essentially uncorrelated between the input sample points, the residual contains much of the medium-scale physical variability that remains just above this lower sampling bound. Hence, it is not surprising to see that fewer phase-descending features are present in the GGIM tilt plots of **Figure 7.21**.

Considering the two intersecting (near-orthogonal) ELOISE AoA paths of different lengths, namely, Lynd River to Laverton (2598 km) and Kalkarindji to Coondambo (1593 km), **Figure 7.22** draws some of the strengths and limitations of the tilt estimation technique into starker relief. In theory, the zonal and meridional tilt angles should be agnostic to the path geometry and very close in value, given that the 1-hop midpoints are separated by only 100 km. However, differences well in excess of the estimated uncertainties are present at times, even in the virtual height profiles (e.g. from 1–2 UT), making it difficult to declare when results are physically meaningful. Despite these shortcomings, the meridional tilts in **Figure 7.22** do show an acceptable level of agreement from 2–5 UT, during what appears to be the passage of a TID sequence.

By performing the same tilt estimation technique on the synthesised GGIM/3D NRT observables, as described in **Section 7.3.2**, it is possible to evaluate the spatial degradation caused by the transformation process itself; that is, from the modelled electron density distribution to the HF observables (via ray tracing) and back again (via a tilted mirror approximation). **Figure 7.23** demonstrates that some, but certainly not all, of the midpoint tilt angle fidelity from **Figure 7.21** is lost through the combined effects of ray tracing/interpolation (i.e. integrating over more than just the midpoint ionosphere), O/X mode merging, and the inverse reflection point mapping process. The tilt magnitudes tend to be reduced after the forward and back transform, especially for

those rays that penetrate more deeply into the F2 layer, presumably due to the spatial integration that occurs. Furthermore, even the sign (direction) of the tilts changes at times when close to zero, possibly due to the neglect of geomagnetic splitting. This leaves a certain degree of doubt over the validity of the reflection point transformation and its assumptions.

To conclude this section, and tie together tilt estimates across multiple ELOISE AoA paths, an example of three collinear paths is presented in **Figure 7.24**. The active storm day on 11 September 2015 contained one of the clearest instances of a night-time TID seen during the ELOISE experiment, manifesting as a series of strong phase-descending fringes, with a period of ~50 min (overlaid dotted lines), between about 9 and 12 UT. For sake of brevity, only the stronger meridional component of the tilt pattern is shown.
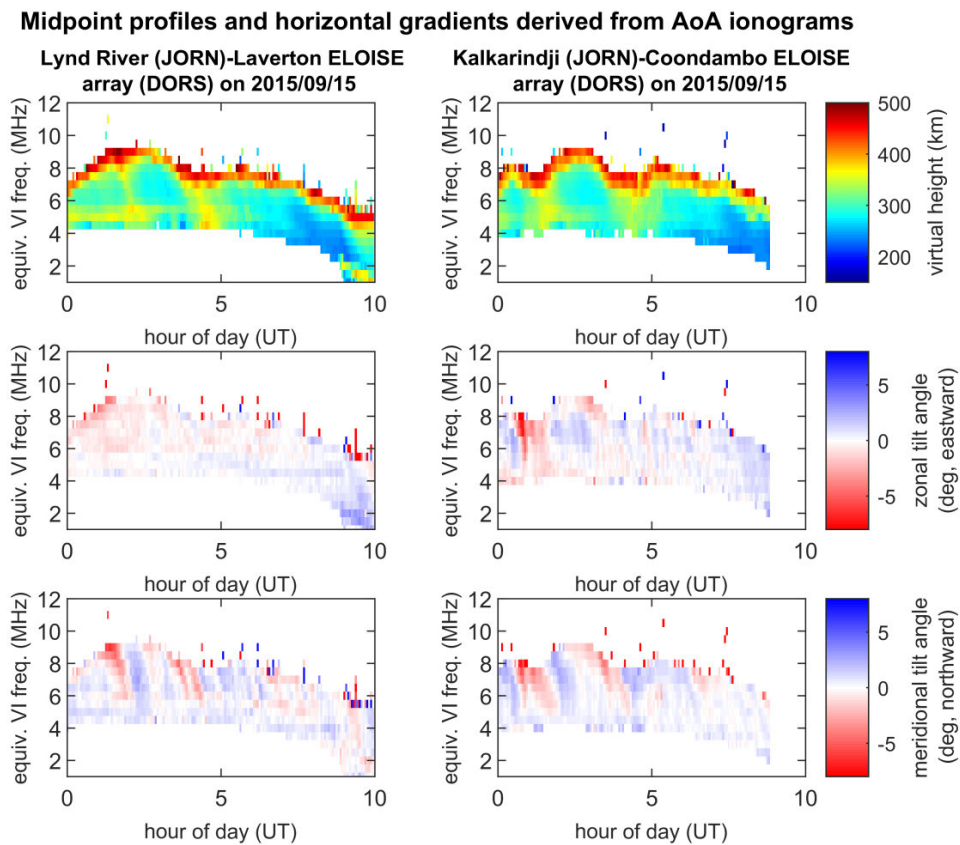
Figure 7.22. Comparison of F2 virtual height profiles and horizontal gradients for two intersecting oblique paths: Lynd River to Laverton (left column) and Kalkarindji to Coondambo (right column). This is the same day as in **Figure 7.19**, but for different paths. As Coondambo was operated on a supervised shift basis only, observations were not available past ~9 UT on this day.
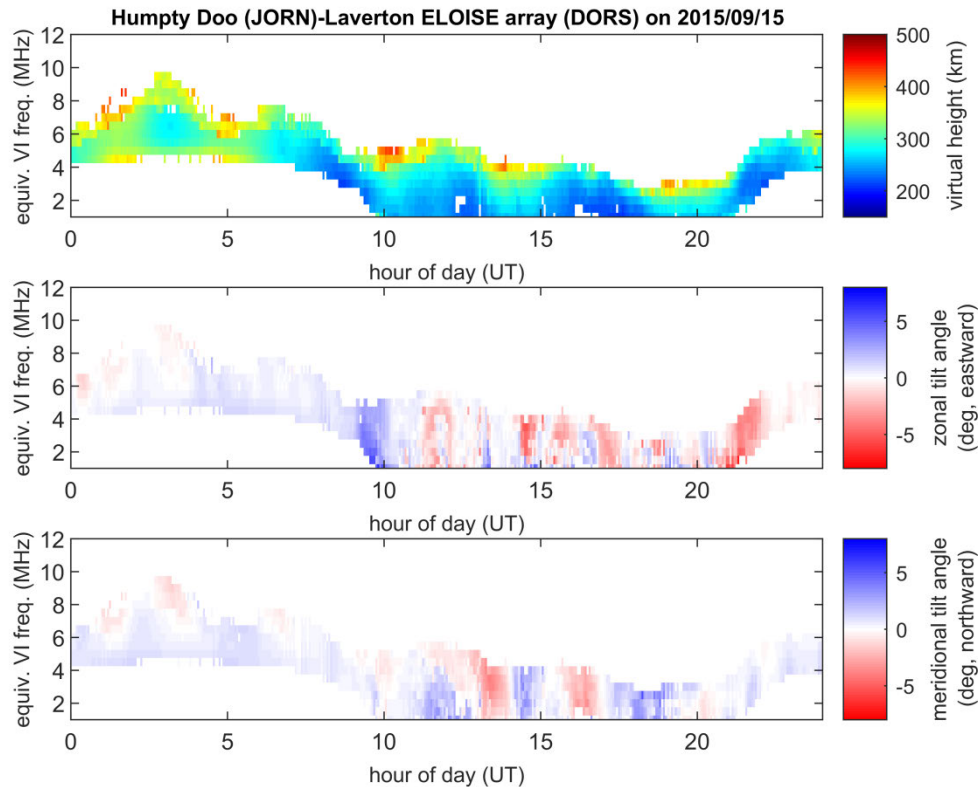
Figure 7.23. Inverted F2 virtual height profiles and horizontal gradients derived by applying reflection point mapping and tilt estimation to interpolated ray data from the GGIM with 3D NRT (cf. true GGIM midpoint tilts in **Figure 7.21**). Rays are sampled every 40 kHz at the same oblique frequencies as in the ELOISE AoA ionograms, to produce a synthetic equivalent of **Figure 7.19**, including inherent mirror model errors.

It is interesting to note that the three paths of different lengths yield remarkably similar tilt estimates (even more so than the intersecting paths of **Figure 7.22**), and that there was no significant lag between adjacent midpoint locations separated by 270–340 km. At most, there appears to be a very slight lag from west (Harts Range) to east (Lynd River) on the same scale as the 3.75 min revisit interval, a result best illustrated in **Figure 7.25** after computing cross-correlation functions between each pair of paths. As TIDs have horizontal speeds of less than 1000 m/s [Hunsucker, 1982], for the largely east-west alignment of midpoints and dominant meridional tilt component, this points to a disturbance propagating in a primarily north-south direction, likely from the auroral or polar source regions towards the equator [Francis, 1975]. Regrettably the near-orthogonal paths into Coondambo were not observed during this period to help verify

such an interpretation, or enable a legitimate determination of the horizontal wavelength of the disturbance.

Figure 7.26 shows the GGIM representation of the meridional tilts estimated in **Figure 7.24**. In this case, the overall level of agreement is quite encouraging, with peaks and troughs closely aligned, although the TID-like features are weaker in the model than the AoA observations suggest, probably for the reasons explained earlier. Nickisch et al. [2016] have previously demonstrated that TIDs and their corresponding tilts can be modelled quite successfully from a network of CSF beacons without assimilating AoA (only delay and Doppler), and these GGIM results tend to concur, even without Doppler observations.
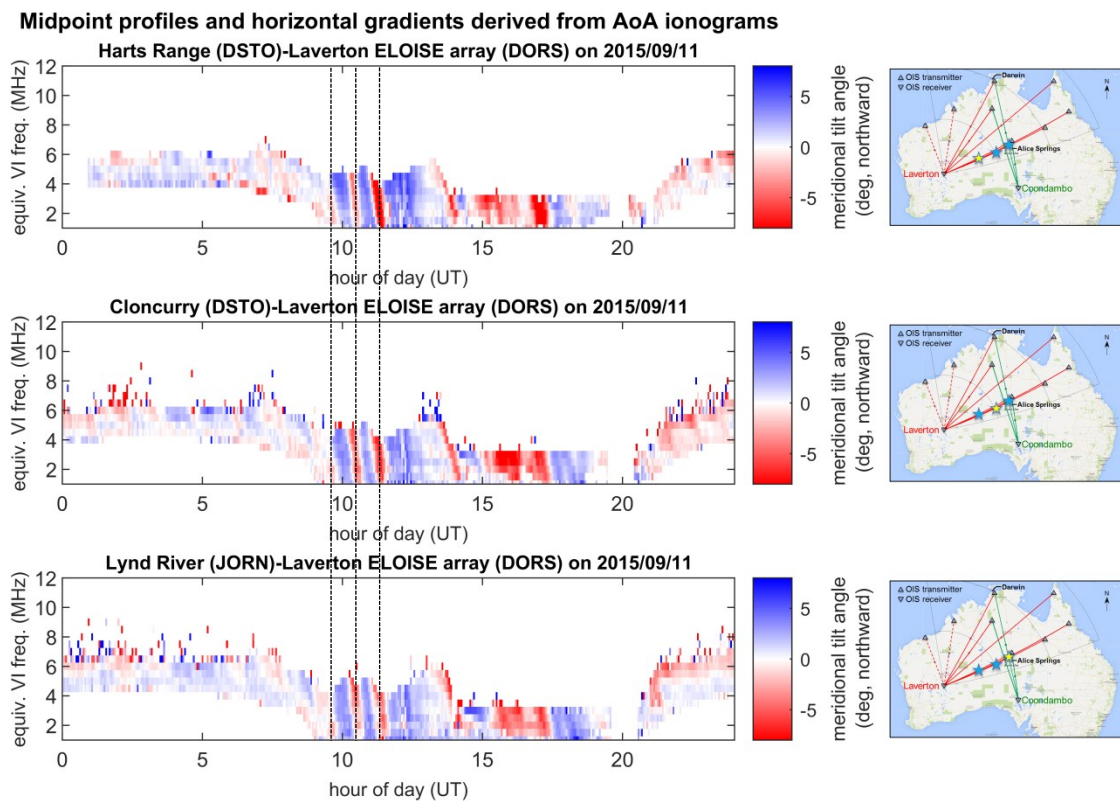
Figure 7.24. Example of meridional F2 tilt estimates for three collinear paths into Laverton. From top to bottom, the panels show results from Harts Range (1383 km), Cloncurry (2055 km), and Lynd River (2598 km). The 1-hop midpoint locations are marked with a yellow star in the thumbnail maps to the right of each row; see **Figure 3.1** in **Section 3.1** for a full-size map of the ELOISE AoA paths. Dotted black lines running vertically through all three panels have been added to assist in interpreting the quasi-periodic TID-like pattern.
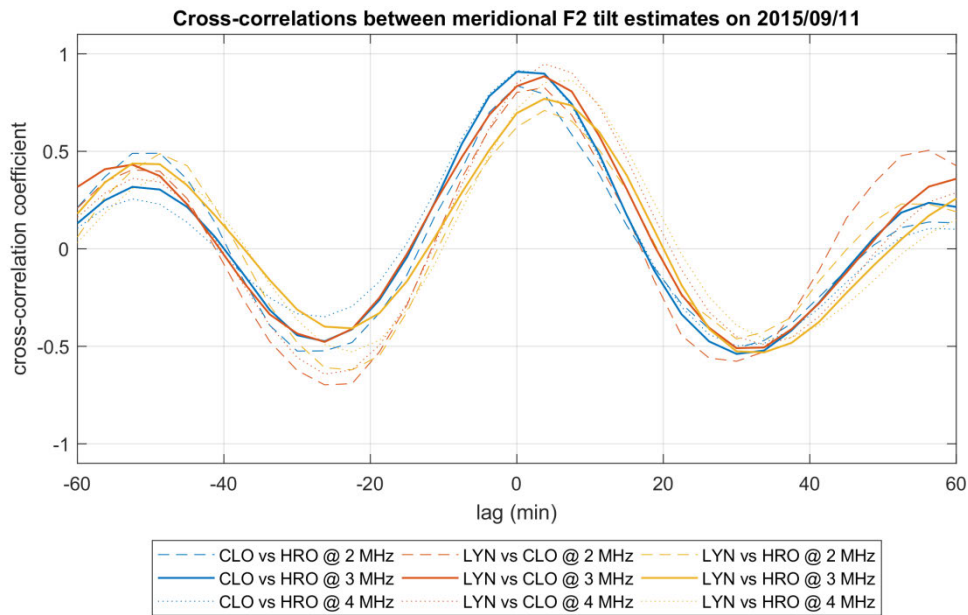
Figure 7.25. Cross-correlation functions between meridional tilt estimates on pairs of the three collinear paths in **Figure 7.24**. A 3-hour time period (0915–1215 UT) has been selected to draw out potential lags in the characteristic TID-like pattern identified. Results calculated at three different equivalent VI frequencies (2, 3 and 4 MHz) each exhibit a similar spatial relationship, with the path from Harts Range (HRO) leading in time, followed by Cloncurry (CLO), then Lynd River (LYN), all within the span of about 5 min.

In terms of the ability to accurately estimate ionospheric tilts from AoA observations, the above results are evidently a mixed bag. While key assumptions of the estimation technique are expected to hold up somewhat better on shorter paths and for night-time observations, there are still many features that cannot be reconciled against the GGIM model. While the hope for this analysis was to combine estimates from across the ELOISE AoA network, the limited night-time data from Coondambo meant that there was insufficient path diversity over many of the most suitable data sets. Originally the intent of this work was to subtract off the large-scale tilts to better reveal TIDs, as a precursor to **Chapter 8**, but with the estimated tilts evidently including substantial contributions from TIDs themselves, this approach was not ultimately pursued.
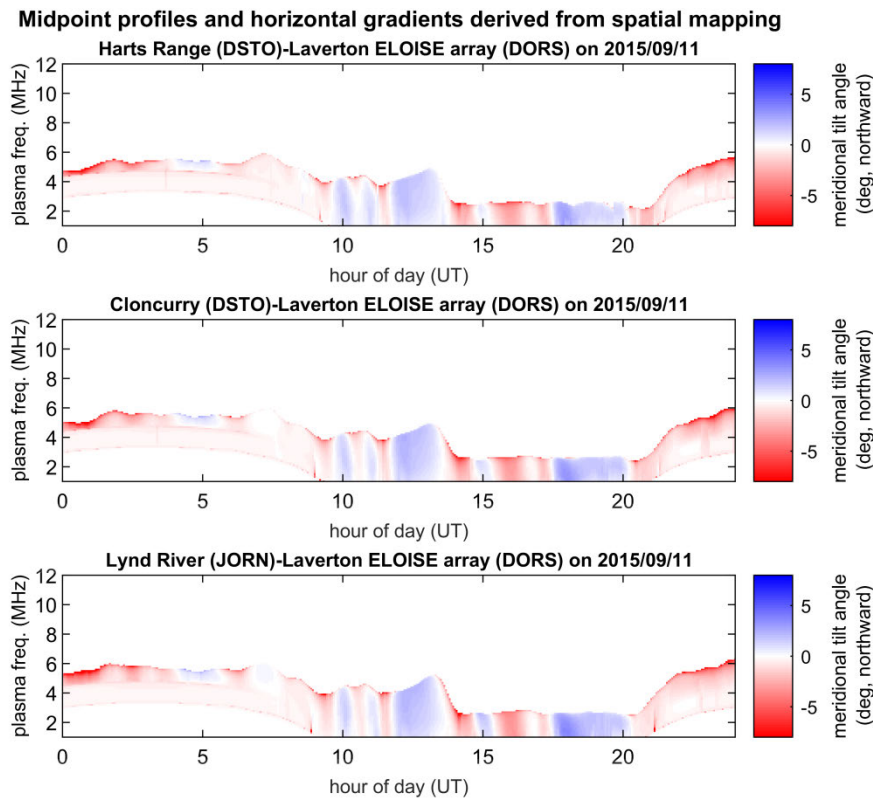
Figure 7.26. Meridional E/F1/F2 tilts derived from the GGIM spatial maps at the midpoints of the three collinear paths in **Figure 7.24**.

7.5   Evaluation of ionospheric drifts

As established in earlier CSF examples, AoA perturbations caused by medium- to large-scale ionospheric variability are typically accompanied by characteristic Doppler signatures, representing changes in the ray path and electron density. Whereas the tilt estimation technique described above contains an inherent directional ambiguity when gradients are associated with horizontal motion in the disturbance field, the Doppler information from CSF observations has the potential to resolve this. Under the assumption of a mirror reflector, measured Doppler shifts correspond to the radial component of motion in the reflecting surface [Dyson, 1975]. Hence, using the AoA-derived path geometries, it is possible to calculate an ionospheric drift velocity vector that fits the Doppler measurements over a given time interval. In theory, these drifts can be attributed to a combination of thermospheric winds and electrodynamic (E×B) transport, as well as gravity wave forcing. The horizontal drift components are effectively the temporal analogue to the (spatial) horizontal tilts, being a first-order

179

representation of the ionospheric structure and dynamics that give rise to the HF propagation observables.

Ionospheric drift analysis is of course not a new problem, and many alternative methodologies have been developed since the mid-20[th] Century [Davies, 1990, ch. 7.9; Rawer, 1993, ch. 12.3; Reinisch et al., 1998]. With conventional (non-Doppler) ionosondes at vertical or quasi-vertical incidence, spaced-sensor experiments are often used, which relate phase lags in the cross-correlations or cross-spectra of height measurements at three or more non-collinear sites to the horizontal component of motion (e.g. Munro [1950]; Harris et al. [2012]). A similar style of analysis can also be applied to other observables besides virtual height or delay, including Doppler itself (e.g. Crowley & Rodrigues [2012]; Chum et al. [2014]).

However, these techniques are not well suited to the ELOISE oblique paths, given that reflection points are often horizontally displaced by more than 100 km from the great circle midpoint (e.g. see **Figure 5.18** in **Section 5.5**); that is, on the same scale as the midpoint separation between adjacent paths. A faster revisit rate would also be more favourable for estimating spatial lags. The technique implemented here instead relies on the measured CSF AoA and Doppler for a single path only, averaged over time and space, and is most similar to the Digisonde Drift Analysis or "sky map" approach [Reinisch et al., 1998; Wright & Pitteway, 1994; Parkinson et al., 1997], albeit at oblique rather than vertical incidence and with a much longer integration period (tens of minutes or hours, rather than seconds). The basic principles are also akin to the full correlation analysis technique of Briggs et al. [1950] (see also Briggs [1984] and Briggs & Vincent [1992]), which relates the velocity of the diffraction pattern on the ground to the velocity of scatterers in the sky; this is most suited to a vertical-incidence (monostatic) arrangement where a receiver per element array is not practical [Briggs, 1980; Whitehead et al., 1983].

Bullett [1994] and Scali et al. [1995] have demonstrated that low-cost ionosonde drift measurements can compare favourably to those from incoherent scatter radar for understanding ionospheric plasma convection, although the drift estimates may not always be representative of the bulk ionospheric motion as assumed [Dyson, 1975; Parkinson et al., 1997]. Furthermore, variations in the velocity profile with height may have a smoothing effect on the drift estimates [Dyson et al., 2001]. Cannon et al. [1991]

discuss some of the limitations of the analysis technique, yet they are still able to find good agreement with satellite measurements of average polar cap convection.

Pfister's theorem [Pfister, 1971; Brownlie et al., 1973; Dyson, 1975; From et al., 1988] relates the Doppler "velocity" (rate of change of phase path) v_{id} to the apparent velocity of the reflecting surface $\boldsymbol{v}$, assuming all Doppler contributions arise from changes in the geometric path length to a mirror reflector in free space or, more generally, from apparent motion of the ionosphere without change of shape [Bennett & Dyson, 1993]. In the general bistatic (oblique incidence) case, Pfister's theorem can be expressed as

$$v_{id} \equiv \frac{dP}{dt} = (\hat{\boldsymbol{s}}_R + \hat{\boldsymbol{s}}_T) \cdot \boldsymbol{v}, \qquad (7.4)$$

where $\hat{\boldsymbol{s}}_R$ and $\hat{\boldsymbol{s}}_T$ are unit vectors directed along the ray path ("line of sight") to the reflection point from the receiver and transmitter, respectively. In the monostatic (vertical incidence) case, phase path changes on both the upward and downward ray segments are equal, and the right side of equation (7.4) reduces to simply $2\hat{\boldsymbol{s}} \cdot \boldsymbol{v}$ (where $\hat{\boldsymbol{s}} = \hat{\boldsymbol{s}}_R = \hat{\boldsymbol{s}}_T$). By letting $\bar{\boldsymbol{s}} \equiv (\hat{\boldsymbol{s}}_R + \hat{\boldsymbol{s}}_T)/2$, the bistatic case takes on a similar form, although ray bisector $\bar{\boldsymbol{s}}$ need not be a unit vector.

This approach comes with a number of serious caveats:

- In practice, the ionosphere is a refracting medium with time-varying electron density, resulting in Doppler contributions that are unrelated to motion in the reflection points and geometric path. Furthermore, it is the projection of motion on the wave normal that matters, not the ray direction; these are not the same in a general anisotropic medium, particularly in the vicinity of the ray apogee. A full expression for the Doppler shift, of which equation (7.4) is an approximation, can be found in Bennett [1968, eq. (13)]. The approximation is expected to hold up best when plasma transport processes dominate over net production and loss (i.e. for near-equilibrium conditions), allowing ionospheric structure to move without an apparent change of shape, and there is minimal refraction from underlying ionisation (e.g. night-time F2-low).

- Like many drift analysis techniques, the validity of results can suffer in presence of multiple propagating disturbances [Hines & Rao, 1968]. The estimated drift velocity

is effectively a superposition of motion across these many wave components, rather than truly representing bulk motion of the reflecting surface as a whole. In some cases, however, one disturbance component may dominate. Note that the latest Generalized Digisonde Drift Analysis [Reinisch et al., 1998; Reinisch et al., 2005; Kozlov & Paznukhov, 2008] has relaxed this assumption and can now resolve several velocities from simultaneous reflection points.

- With only three CSF dwells every 3.75 min, each on its own discrete frequency, it is necessary to perform a significant amount of temporal averaging to achieve a semi-meaningful result. This is done by combining peaks across many separate dwells into the one set of equations, and is in addition to the spatial averaging over the reflecting region, both horizontally and vertically. While a time-varying but frequency-independent vertical drift component is supported, the horizontal drift velocity is assumed to be constant for all such peaks, thus neglecting any non-stationarities and vertical shears. Given scope for a limited number of degrees of freedom, the justification for allowing the vertical velocity component to vary is that this is expected to oscillate more than the horizontal component over a TID period, both due to the passage of the wave itself and background ionospheric changes.

To obtain the drift velocity, equation (7.4) is posed as a linear least-squares matrix problem, in a local North-East-Down (NED) Cartesian coordinate system centred on the midpoint, and solved using all F2-low peak data from a given time interval. That is,

$$S\boldsymbol{u} = \frac{1}{2}\boldsymbol{v}_{id}, \tag{7.5}$$

where the $K \times (L+2)$ design matrix

$$S = \begin{bmatrix} \bar{s}_{N,1} & \bar{s}_{E,1} & \bar{s}_{D,1} & 0 & \cdots & 0 \\ \bar{s}_{N,2} & \bar{s}_{E,2} & \bar{s}_{D,2} & 0 & \cdots & 0 \\ \bar{s}_{N,3} & \bar{s}_{E,3} & \bar{s}_{D,3} & 0 & \cdots & 0 \\ \bar{s}_{N,4} & \bar{s}_{E,4} & 0 & \bar{s}_{D,4} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \bar{s}_{N,K} & \bar{s}_{E,K} & 0 & 0 & \cdots & \bar{s}_{D,K} \end{bmatrix} \tag{7.6}$$

contains the components of the ray bisector in NED coordinates for each mapped peak $(1, ..., K)$, drift velocity vector $\boldsymbol{u} = \left[v_N, v_E, v_{D,t=1}, v_{D,t=2}, \cdots, v_{D,t=L} \right]^T$ contains the NED drift velocity components (including time-varying v_D at time epochs $t = 1, ..., L$), and Doppler velocity vector $\boldsymbol{v}_{id} = \left[v_{id,1}, \cdots, v_{id,K} \right]^T$ contains the measurements. Note that there are ordinarily about three F2-low peak measurements for each time epoch separated by 3.75 min (i.e. one on each CSF channel), whose downward ray bisector components are grouped together in a single column of matrix S, to match a single element of vector $\boldsymbol{u}$. It is important for the reflection points to be sufficiently diverse in their locations (ray directions) so that the pseudo-inverse is computable; in other words, the ionosphere must be reasonably disturbed, with enough off-angle returns to capture different projections of the drift vector.

Figure 7.27 shows a 3-hour early evening sample of CSF peak observables corresponding to the TID-like disturbance in the ionogram tilt plots of **Figure 7.24**. In terms of angular diversity, this is a good candidate for drift estimation. To provide a sufficient number of peaks for the least-squares inversion, all three CSF channels on different (time-varying) F2 frequencies have been combined, with the dependence on frequency captured in **Figure 7.28**. Although this combining causes some spread in delay between the three channels, the other observables are fairly uniform across all frequencies. The overwhelming pattern remains very much quasi-periodic in nature and largely monochromatic. Doppler velocities of up to +80 m/s were observed on this path from Cloncurry to Laverton, although this is still only a Doppler frequency shift of -2 Hz, well within the ±5 Hz unambiguous range.

After calculating reflection points for all six CSF paths into Laverton, the result is the geographic maps in **Figure 7.29**. As noted during the array calibration work, the distributions are highly elongated along the down-range axes, as marked by the grey dashed arrows. The estimated horizontal drift velocity components, along with their standard errors, are printed in the lower right corner of each panel, while the time-varying vertical drift velocity components are plotted separately in **Figure 7.30**.

At a glance, it is quite clear that the horizontal drift estimates, plotted as black arrows in **Figure 7.29**, are not consistent across paths, nor is there any meaningful geographical variation. It is possible that the geometry of each path makes it sensitive to a different
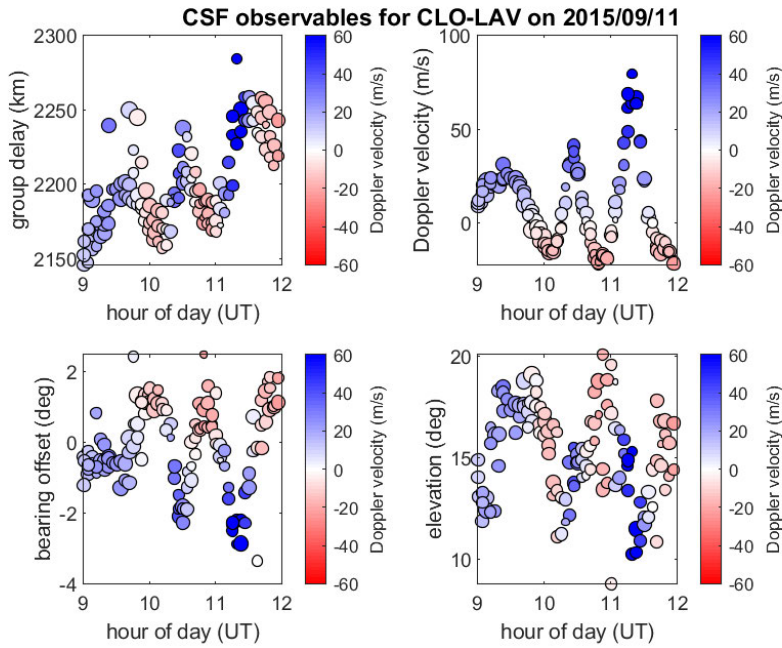
Figure 7.27. Peak observables for the Cloncurry to Laverton CSF path, combining all three channels from 9–12 UT on 11 September 2015. From top to bottom and left to right, the panels show group delay, Doppler velocity, bearing offset, and elevation. This period corresponds to the TID-like fringe pattern in **Figure 7.24**. Markers are colour-coded by Doppler velocity and scaled in size by SNR.
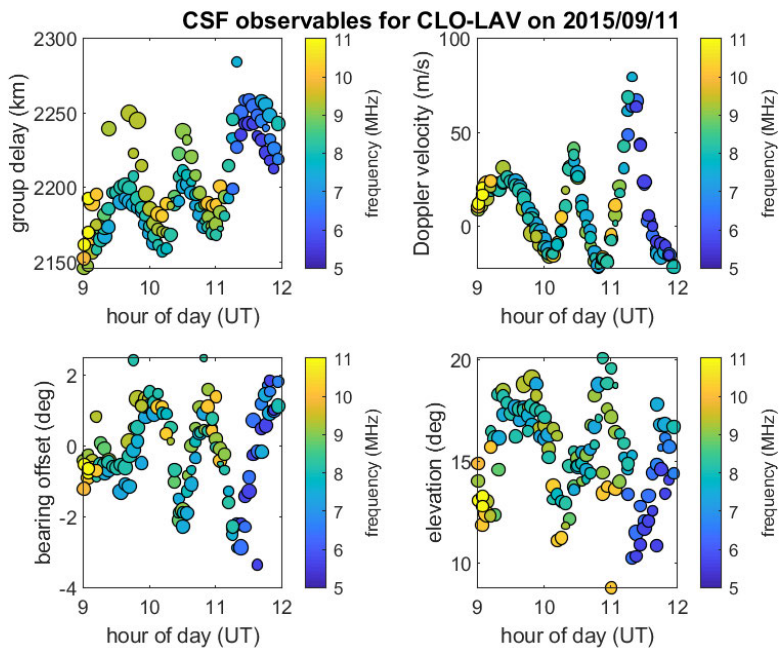
Figure 7.28. As above, but with peaks colour-coded by the CSF frequency (across all three channels).
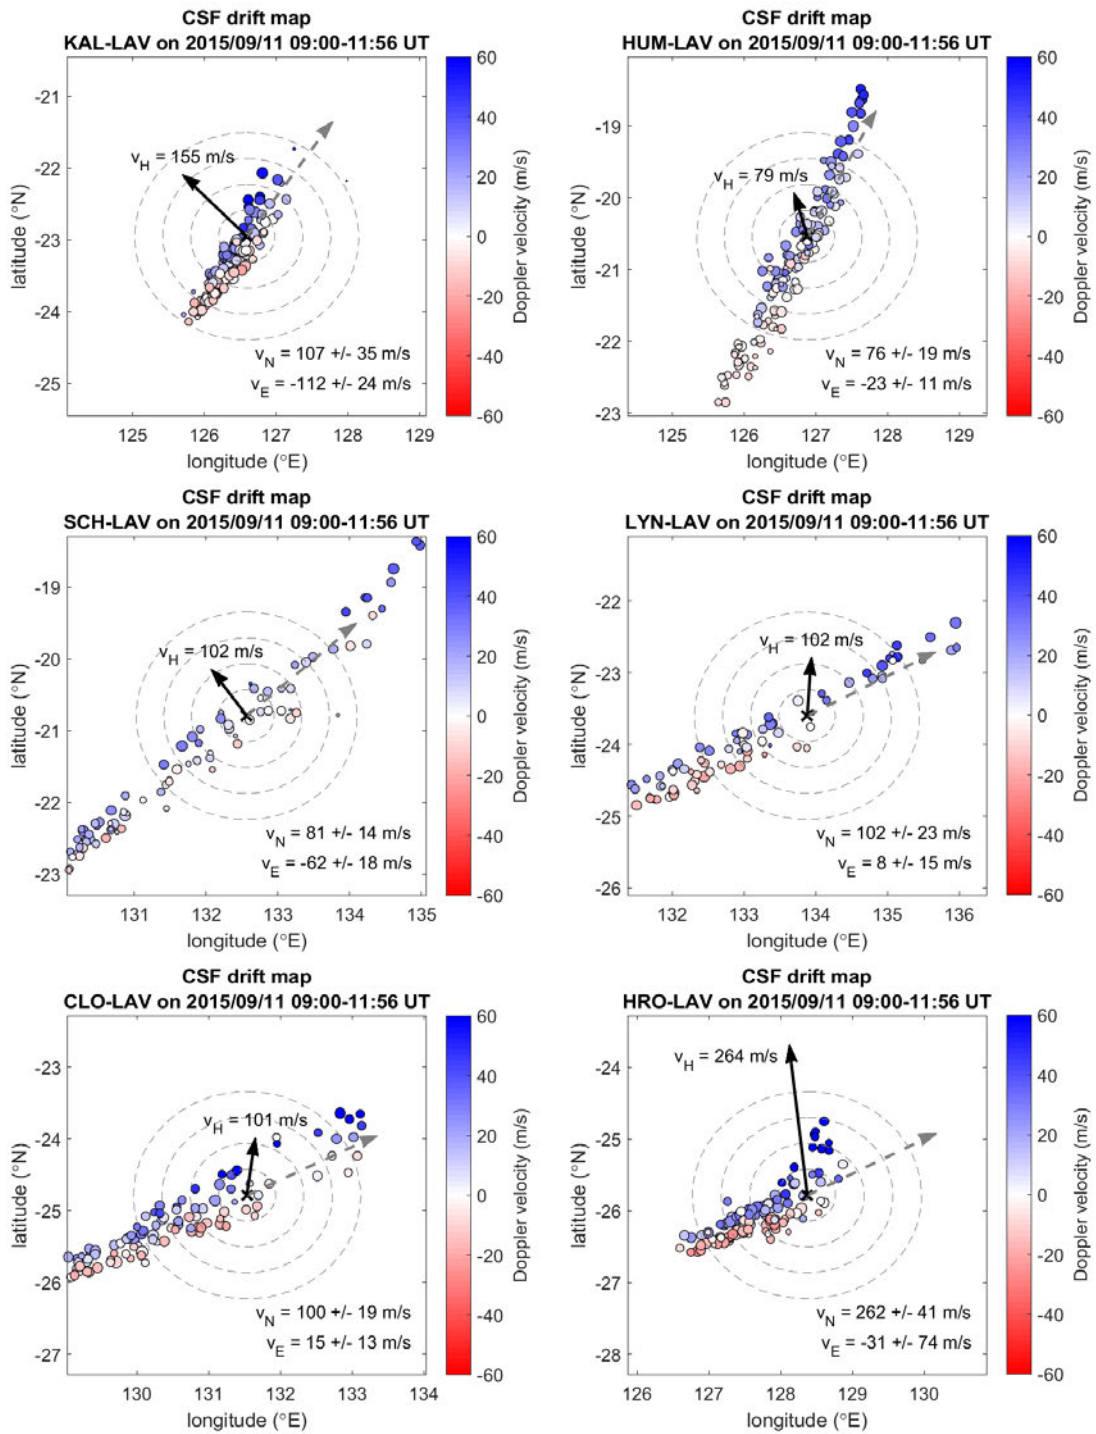
Figure 7.29. Geographic maps of CSF F2-low reflection points, each colour-coded by Doppler velocity and scaled by SNR, and the fitted horizontal (northward and eastward) drift velocity components. All three CSF channels from 9–12 UT on 11 September 2015 have been aggregated in these plots. The six panels show the six CSF paths into Laverton: Kalkarindji (top-left), Humpty Doo (top-right), Scherger (middle-left), Lynd River (middle-right), Cloncurry (bottom-left), and Harts Range (bottom-right). The concentric grey circles represent radii of 40, 80, 120 and 160 km about the great circle midpoint (black cross), while the grey dashed arrow indicates the midpoint bearing.
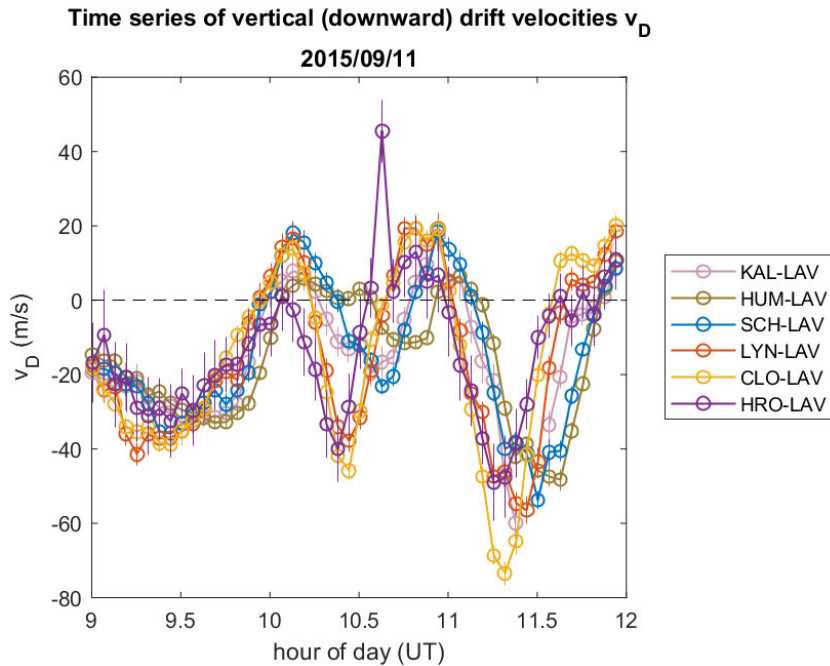
Figure 7.30. Time series of the vertical (downward) drift velocity component estimated for all six CSF paths into Laverton, corresponding to the drift maps in **Figure 7.29**. The vertical bars represent the standard error estimates.

component of the disturbance field, and that in effect the drift analysis is preferentially drawing out the velocity of just one structure ("ripple") that meets the condition for specular reflection, rather than the bulk motion of the ionosphere as a whole (e.g. see discussion in Rawer [1993, p. 207]). Suspicion is also once again cast on calibration issues and the significant assumptions of the analysis described above. Whatever the case, it can only be concluded that the horizontal drifts are unlikely to be physically meaningful, for example as a measure of the TID wave velocity.

On the other hand, the vertical (time-varying) drift estimates, shown in **Figure 7.30**, look considerably more sensible. All paths exhibit a similar quasi-periodic pattern, with amplitudes broadly in agreement once standard errors are factored in. The phase of the more southerly path midpoints, like Cloncurry to Laverton (CLO-LAV) and Harts Range to Laverton (HRO-LAV), tends to lead the more northerly midpoints, like Humpty Doo to Laverton (HUM-LAV) and Scherger to Laverton (SCH-LAV). This implies a predominantly north to north-westward propagating disturbance, in line with the earlier conclusion regarding **Figure 7.24**. Interestingly, the horizontal drift vectors

in **Figure 7.29** are indeed largely north to north-westward in direction, but with a wide range of spurious-looking magnitudes from 79 to 264 m/s. Without an accurate source of truth, these estimates are admittedly difficult to validate. In **Chapter 8**, a distinction is made between the TID propagation direction (i.e. normal to the TID phase-front) and the axis of reflection, as these are not generally the same.

7.6   Chapter summary

- This chapter has described the effects of medium- to large-scale spatial and temporal ionospheric variability on the ELOISE oblique group delay, Doppler and AoA observables. Both persistent patterns of diurnal variability, such as around the dawn/dusk terminators, and the day-to-day and hour-to-hour perturbations associated with specific disturbances, such as TIDs, have been identified and interpreted.

- Whereas the observables themselves are very much dependent on the oblique path geometry, techniques for horizontal gradient (tilt) and drift velocity estimation have been implemented as a means of inferring the physical characteristics of the ionosphere. However, these rely on a number of approximations, including a tilted spherical mirror model to represent each reflection, which may not be valid under certain circumstances. The struggle to relate horizontal drifts between different ELOISE CSF paths reflects the challenges of performing such an analysis on relatively long oblique paths.

- Results in both the observable and tilt domains have been compared against a data-assimilative ionospheric model, constructed from conventional OIS and VIS scaled parameters (i.e. without AoA or Doppler measurements). Overall there is a good level of agreement, particularly in the observable domain using 3D magneto-ionic numerical ray tracing, for which a statistical analysis revealed differences on the same scale as the measurement errors. In the tilt domain there are more discrepancies, likely caused by assumptions in the reflection point transformation and omission of the residual anomaly field from the model tilt estimates. Even prior to spatial mapping, the scaled quasi-parabolic profile parameters, sampled from ionograms every 3.75 min, are often sufficient to characterise much of the Doppler

variability in the ELOISE AoA observations; that is, using just a finite difference approximation.

- Although some notable examples of monochromatic TID-like disturbances were observed across multiple intersecting or collinear ELOISE AoA paths, in general it was difficult to make meaningful conclusions about inter-site patterns of variability from the ELOISE AoA data set. This was in part due to limitations in the tilt estimation technique, and also a lack of suitable night-time observations from Coondambo.

8 Identification and parameterisation of travelling ionospheric disturbance signatures

8.1 Overview and examples of TID signatures

Travelling ionospheric disturbances (TIDs) [Davies, 1990, ch. 7.10; Rawer, 1993, ch. 11] are one of the key contributors to spatial and temporal variability of HF radio-wave propagation through the ionosphere. First observed and explained by Munro in the 1950s (e.g. Munro [1950]; Munro & Heisler [1956a]), these wave-like perturbations in the background electron density are typically interpreted as the ionospheric manifestation of atmospheric gravity waves (AGWs) in the thermosphere [Hines, 1960; Yeh & Liu, 1974; Francis, 1975; Hunsucker, 1982; Hocke & Schlegel, 1996]. Although many candidate source mechanisms have been proposed, within the F2 layer, medium-scale MSTIDs (with periods of 15–60 min and horizontal wavelengths of several hundred km) are broadly thought to originate from regional tropospheric sources [Gossard, 1962], while large-scale LSTIDs (with periods of 30 min to 3 hr and horizontal wavelengths of ~1000 km or more) are often associated with geomagnetic disturbances in distant auroral regions [Richmond, 1978]. Despite not conforming to the traditional definition of a TID, solitary waves or wave packets may still be treated as a class of TID if they possess the same spatial phase properties [Afraimovich et al., 2003]. TIDs are frequently responsible for large group delay and AoA perturbations in HF propagation paths [Jones & Reynolds, 1975], along with additional discrete modes, and their occurrence is also closely linked to night-time spread-F in ionograms and radio scintillation effects caused by small-scale plasma bubbles [Abdu et al., 2009; Booker, 1979; Bowman & Monro, 1988; Kelley et al., 1981].

In recent years there has been increased interest in monitoring and forecasting these disturbances to support a diverse range of technologies sensitive to space weather, including Global Navigation Satellite Systems (GNSS), low-frequency radio astronomy, and HF sky-wave radar systems. Projects such as the Network for TID Exploration

(Net-TIDE) in Europe [Reinisch et al., 2018] have been established to develop regional TID diagnostic and warning capabilities to such users. Over-the-horizon radars (OTHRs) designed for long-range aircraft and ship detection, such as those comprising the Jindalee Operational Radar Network (JORN) in Australia [Colegrove, 2000; Johnson et al., 2008; Sinnott, 2015], may have their ability to detect and geolocate targets degraded in the presence of TIDs. This is usually due to ionospherically induced perturbations in the backscattered returns from targets, and inadequate representation of the disturbance fields in propagation models constructed from sparsely sampled ionospheric measurements. A more detailed discussion of the challenges posed by TIDs for OTHR and potential mitigation strategies can be found in Fridman & Nickisch [2001] and Nickisch et al. [2006; 2007].

Several examples of likely TIDs observed during the ELOISE experiment have been presented earlier in this thesis, and their appearance in the AoA ionograms varies quite considerably. Classifying TIDs based on how they manifest in the data is in many ways the most practical scheme, compared to more conventional definitions based on their geophysical sources and spatial/temporal scales. LSTIDs, in the absence of other irregularities and neglecting geomagnetic effects, often support only a single reflecting path from the transmitter on a given frequency and mode (e.g. 1-hop F2-low). The compression and expansion of the layer associated with the passage of the TID [Lynn et al., 2016] causes the ionogram trace to move up and down in apparent range, and in and out in frequency, as the maximum observed frequency (MOF) increases according to the secant law. This brings about the appearance of the entire trace "breathing" in and out in a quasi-periodic fashion. Usually such observations are accompanied by a similar signature in AoA, as the disturbance shifts the reflecting point away from the great circle midpoint, although this depends on the geometry of the path with respect to the TID direction. **Figure 7.27** from **Section 7.5** illustrates a likely candidate for such an LSTID.

As the amplitude-to-wavelength ratio and, thus, iso-ionic surface curvature of the LSTID increases, additional off-angle ("satellite") traces are more likely to be visible in the ionogram [Georges, 1967b, ch. 5.2.2; Chum et al., 2010], as in **Figure 5.2**, **Figure 5.3** and **Figure 5.4** of **Section 5.2**. Often these ascending and descending traces appear in a quasi-periodic sequence, embodying the wave-like nature of the underlying AGW, though sometimes they can occur as an isolated event, indicative of a more frontal

(bore-like) structure passing over the field of view. Similar disturbance patterns have been seen, for example, in the virtual height plots of Weber et al. [1978], albeit for the equatorial region, and were found to correspond to bands of field-aligned depletions in all-sky airglow images at 630 nm, representing atomic oxygen emissions at F2 altitudes. No such corresponding cases were found in the ELOISE airglow data from Alice Springs, in part because of the limited observations and high number of cloudy nights that prevented many meaningful comparisons. Examples of inverse correlations between ELOISE airglow intensity and the F2 layer height from QVIS ionograms, reported in Unewisse et al. [2016], are likely to be "breathing" modes associated with the aforementioned class of weaker-amplitude LSTIDs.

MSTIDs, on the other hand, frequently cause multiple off-angle reflections due to their shorter horizontal wavelengths, yet this tends to be localised on the F2 trace owing to the fact that their vertical wavelengths (on the scale of ~100 km) are usually only a fraction of the thickness of the F2 layer. The trace "kink" that folds back on itself is a widely reported feature of this class of disturbance [Munro, 1950; Georges, 1968; Lobb & Titheridge, 1977; Morgan et al., 1978; Cervera & Harris, 2014]; see for example **Figure 5.1** of **Section 5.2**. In other cases, where the MSTID has a lower amplitude-to-wavelength ratio, the kink is reduced to a more subtle discontinuity ("wiggle") in the trace curvature (e.g. **Figure 4.5** of **Section 4.2.1**). Over sequences of ionograms, such features are seen to progressively descend down the trace, with the O-mode usually leading the X-mode [Munro & Heisler, 1956b; Davies, 1990, p. 201], due to ray divergence and the forward tilt of the TID phase-front.

As noted in the previous chapter, not all quasi-periodic signatures in the ELOISE data will have been caused by TIDs. The following characteristics are usually prerequisites:

- Common periodicities (across multiple observables) that fall within the physical range of values for AGWs reaching ionospheric heights. The balance between viscous damping, which preferentially depletes energy from the smaller-scale waves, and exponential amplitude growth, which likewise increases with height due to decreasing atmospheric density, means that different AGW components attain their maximum amplitude at different heights. As such, a relatively narrow set of periods typically dominate at any given height [Hines, 1960; Pitteway & Hines, 1963; Georges, 1967a]. In the F2 region, most of the energy tends to be found in

wave periods of 20–30 min (e.g. Bowman [1992]), with a lower limit of about 10–15 min imposed by the acoustic cut-off frequency [Barry, 1963], and a more gradual roll-off above 1–2 hr due to the influence of ion drag (hydromagnetic damping) [Hines, 1960; Clark, 1970].

- A height-dependent phase pattern, whereby the perturbation is first seen at greater heights (delays) then descends, as a result of the forward-tilted TID phase-front [Bowman, 1968]. Height increases with frequency along the 1-hop F2-low trace, so in the case of the ELOISE CSF observations, the TID shows up as a lag in the Doppler signatures recorded on the different (concurrent) channels; see for example **Figure 8.1**, in which the larger perturbations (e.g. around 0930 and 1330 UT) tend to appear marginally earlier on the highest frequency channel and later on the lowest channel. The tilt angle, between the AGW or TID wave vector (real component) and the horizontal plane, is derived in **Appendix E** and varies from roughly 45° at F2 heights for a typical MSTID to 80° for an LSTID, depending on the wave period [Hines, 1960; Morgan et al., 1978]. The TID tilt angle is often measured from the slopes of iso-ionic contours, obtained via true height analysis of conventional ionograms (e.g. Clarke [1972, ch. 7.4.3]; Bowman [1992]; Pederick et al. [2017]), although this usually neglects off-angle reflections. The variation in the AGW tilt angle with height is illustrated in **Figure 8.2** for seven representative wave periods (15–180 min) and a model neutral atmosphere based on NRLMSISE-00 [Picone et al., 2002]. Across Australia, the TID tilt angle will generally be a little larger than the underlying AGW, due to the height-dependent AGW/TID phase offset which is a function of the geomagnetic field alignment (see **Appendix E**).

- Temporal lags between corresponding signatures detected on different sounder paths, and between 1-hop and multi-hop returns on a single path, consistent with horizontal motion between the respective ionospheric sample points. The spatial separation of ELOISE AoA midpoints was designed such that reasonable correlation of TID signatures between adjacent midpoints should be expected, although the differing oblique path geometries can make comparisons challenging. Some TID propagation directions will be favoured over others due to ion drag (i.e. selective damping of AGWs with fronts that are not field-aligned) [Hines, 1968], and the effects of wind filtering through the stratosphere and thermosphere [Cowling et al., 1971], where wind speeds can reach several hundreds of metres per second [Kohl &

King, 1967; Cho & Yeh, 1970]. Blocking diagrams, such as those in **Figure 8.3**, can be used to determine which AGW and, hence, TID horizontal velocities will be eliminated by background neutral winds before reaching F2 heights, according to critical layer theory [Taylor et al., 1993; Medeiros et al., 2003]. Despite significant hour-of-day variation, the observed preference for westward and equatorward propagation directions across mid-latitudes at night (e.g. Shiokawa et al. [2009]) is broadly compatible with the 12 UT panel. However, these predictions are only as good as the model for the neutral wind profile, which tends to suffer from lack of measurements in the thermosphere, and of course many AGWs/TIDs travel great distances, so no single height profile truly represents all interactions.
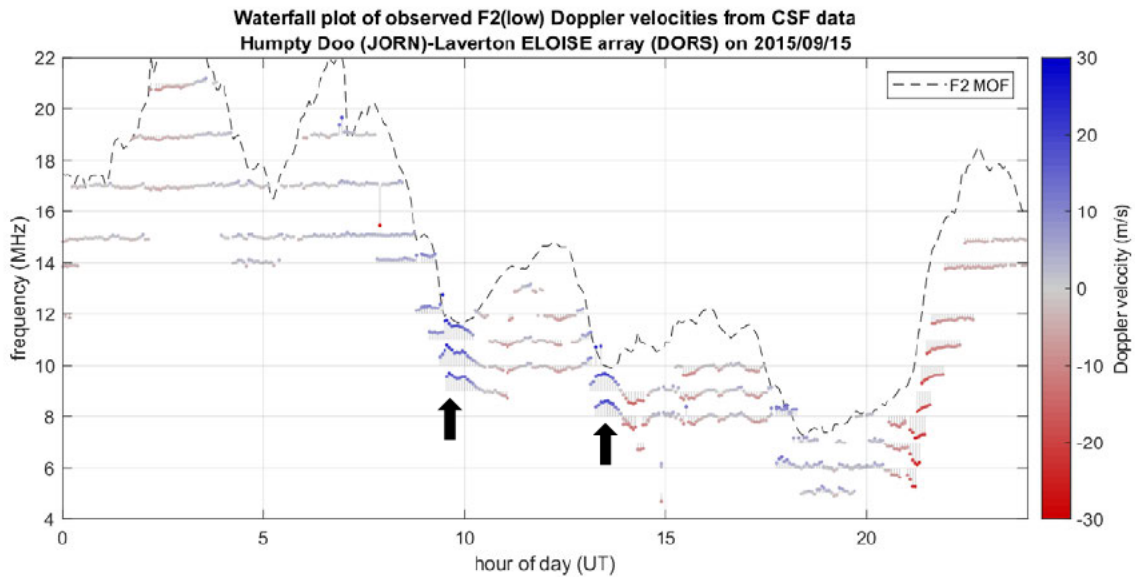
Figure 8.1. Waterfall plot of 1-hop F2-low Doppler velocities from a day of CSF observations on the Humpty Doo to Laverton ELOISE path (15 September 2015). At each 3.75 min epoch, there are three concurrent CSF channels in use, which are shown as nested time series centred on their respective operating frequencies, with markers displaced and coloured according to the Doppler velocity. For sake of clarity, only the dominant F2-low mode from each dwell is retained here. The dashed line represents the F2 MOF, extracted from the adjacent ionograms, while the large black arrows at ~0930 and ~1330 UT highlight two notable TID-like features, where the Doppler perturbations appear in descending frequency order. Note that the temporal lag from one CSF channel to the next is on the scale of one epoch (for a ~1 MHz channel separation), well in excess of the 12 s spacing between each of the three interleaved dwells within a single epoch.
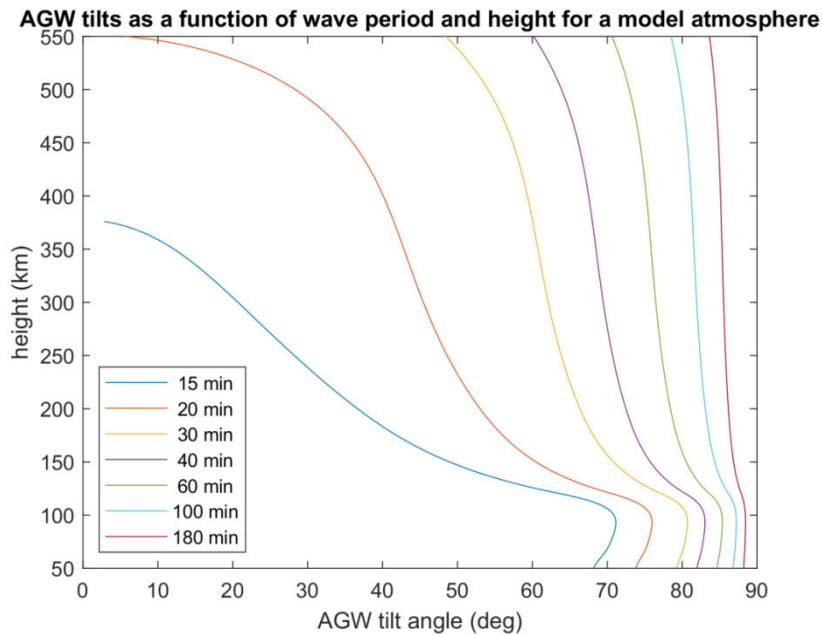
Figure 8.2. Tilt angle profiles for atmospheric gravity waves of various periods, evaluated using a model neutral atmosphere based on NRLMSISE-00 (night-time, Alice Springs). The longer period waves have larger (steeper) tilt angles, and in general the tilt angle decreases with height in the F2 region as the vertical wavelength increases, according to the solution of the AGW dispersion relation in equation (E.3) of **Appendix E**.

Other properties like focusing/defocusing effects [Jones & Lyon, 1974], preferred times of day (e.g. near dawn/dusk solar terminators), links with known meteorological sources (e.g. lightning storms) and space weather events (e.g. causing particle precipitation and thermospheric heating at high latitudes), and multiple simultaneous reflection paths are also common indicators of TIDs, yet not always conclusive. In reality, TIDs are ubiquitous and rarely is there only a single monochromatic component [Afraimovich et al., 2003]; examples like the time sequences of peaks in **Figure 7.1** and **Figure 7.13** of **Chapter 7** are typical of observations. Thus, even when a dominant periodicity is present, there is a potential for many more TID constituents to go undetected, and parameters inferred under a single-wave assumption are at serious risk of errors [Jacobson & Carlos, 1991]. The approach taken in this work was to treat any candidate quasi-periodic signature as a TID, and if the parameters required to synthetically replicate the signature were inconsistent with geophysical expectations, then the hypothesis was dismissed.

In the analysis that follows, two candidate TID models are proposed (as previously described in Heitmann et al. [2018]) to assist in classifying and parameterising TID-like features in the ELOISE AoA observations: one based on a simple mirror reflector, and the other a 3D electron density perturbation governed by physical interactions with AGWs in the neutral atmosphere. The merits of each model are discussed, in the context of their ability to synthetically reproduce the observables, and potential methods for the robust inversion of TID parameters in an automated manner are explored. In practice, this is a much more challenging problem at oblique incidence than vertical incidence. Finally, some comments are made about the ability to relate TIDs across multiple ELOISE AoA paths.
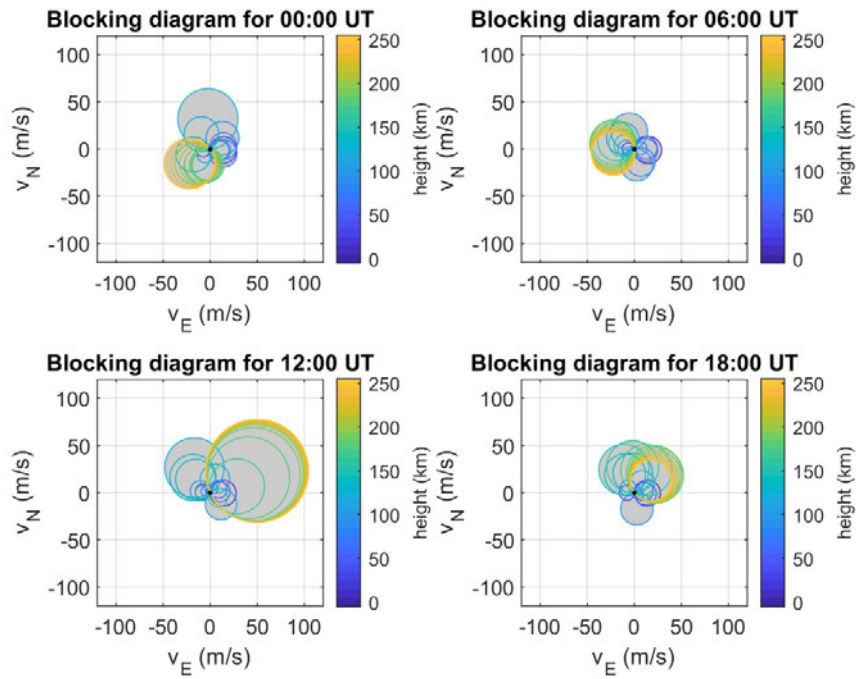
Figure 8.3. Sample blocking diagrams showing the proportion of AGW horizontal velocities that will be filtered by the prevailing neutral winds before reaching F2 heights (250 km). The four panels represent different hours of the day (0, 6, 12 and 18 UT) for which the Horizontal Wind Model HWM14 [Drob et al., 2008; Drob et al., 2015] has been evaluated at Alice Springs, where local solar time is approximately UT + 9 hr. Those AGW/TID wave velocities falling in the grey-shaded regions (i.e. with horizontal speeds $v_{TID} \leq u_z \sin \theta_{TID} + u_m \cos \theta_{TID}$, for zonal wind u_z, meridional wind u_m, and wave bearing θ_{TID}, measured clockwise from True North) have their intrinsic frequencies Doppler-shifted to zero and, thus, cease to exist at heights above this critical layer due to heavy absorption [Cowling et al., 1971].

8.2   Candidate TID models

To aid the interpretation of ELOISE results, realistic TID models and associated ray tracing/homing algorithms were developed to synthesise the observables from the AoA array (i.e. group delay, Doppler, bearing and elevation). The effects of TIDs from far-field sources on HF propagation have been widely synthesised in the past using such models, which seek to parameterise the TID (or its underlying AGW) in terms of amplitude, period, wave vector and phase. Perturbations may be applied to either the ionospheric height (i.e. hmF2) or the electron density (i.e. NmF2 or foF2). Commonly adopted TID model variants, in order of increasing complexity, include:

1   A corrugated (sinusoidal) mirror reflector (e.g. Lyon [1979]; Huang et al. [2016]).

2   A sinusoidal electron density perturbation, typically height-dependent, applied to a realistic background ionosphere (e.g. Ogawa et al. [2002]; Huang et al. [2016]).

3   A sequence of tilted Gaussian peak or trough electron density perturbations applied to a background ionosphere (e.g. Lobb & Titheridge [1977]).

4   An AGW-seeded physics-based ("Hooke") model for computing electron density perturbations, driven primarily by collisional interactions with the neutral atmospheric constituents (e.g. Hooke [1968]; Cervera & Harris [2014]). Similar coupled AGW/TID perturbation models have also been described by Clark [1970], Clark et al. [1971], Davis [1973], Francis [1974], Bottone [1993], and Kirchengast [1996].

In this work, only models 1 and 4 were considered, representing the simplest and most sophisticated variants, respectively. As models 2 and 3 still call for the same ray tracing engine as in model 4, they offer little computational cost saving, yet lack the stronger physical foundation in hydrodynamic theory and neutral/ion coupling. For sake of clarity, monochromatic TID perturbation fields were adopted as a starting point. It is reasonable to further assume that these are plane waves beyond the vicinity of the AGW source. Admittedly only a very limited fraction of ELOISE oblique observations, mostly night-time cases, actually fell into this monochromatic regime; for the remainder of the time, the observations were neither truly periodic nor was there just a single dominant component. Oblique propagation synthesis was carried out in 3D using basic specular ray homing for model 1 and full numerical ray tracing for model 4.

8.2.1 Corrugated mirror model

The corrugated spherical mirror reflector (model 1) is the simplest of the candidate TID models, consisting of a spherical shell ionosphere with a sinusoidal perturbation superposed. It has the advantage of being quick to compute (being specular reflection only) and directly invertible from observations under certain conditions; for example, using the Frequency and Angular Sounding (FAS) technique [Beley et al., 1995; Galushko et al., 2003; Paznukhov et al., 2012]. It can also be relatively easily extended to support multiple frequency components and large-scale (global) tilts. Such a model is therefore useful for classifying disturbances in the observed data and providing first-order estimates of the TID parameters.

Starting with a spherical shell (of fixed radius, but not necessarily Earth-concentric), one or more sinusoidal perturbations are superposed in 3D to build up the disturbance field. These perturbations each have a fixed amplitude and direction, with spatial phase defined in terms of a spherical distance from the great circle line that runs through the (unperturbed) midpoint orthogonal to the TID bearing; in other words, the horizontal wavelength is strictly an arc length along a bearing normal to the phase-fronts. As such, the radius of the reflecting ionospheric surface as a function of the local horizontal coordinates at the midpoint (with x-axis pointing in the down-range direction and y-axis in the cross-range direction) can be written as

$$r_i(x, y, t) = r_{i0} + \sum_n A_n \cos\left(\omega_n t - k_{h,n} g_n(x, y) + \varphi_n\right), \tag{8.1}$$

where r_{i0} is the unperturbed spherical radius, A_n are the TID amplitudes (with subscript n denoting a single harmonic component), $\omega_n = 2\pi/T_n$ are the TID angular frequencies (for periods T_n), $k_{h,n} = 2\pi/\lambda_{h,n}$ are the TID horizontal wave numbers (for horizontal wavelengths $\lambda_{h,n}$), $g_n(x, y)$ are the spatial phase functions (equal to zero at the unperturbed midpoint), and φ_n are the fixed phase offsets. For the n^{th} component, the TID azimuth χ_n, measured clockwise from the midpoint path bearing, is built into the function $g_n(x, y)$.

Given a disturbance field made up of an arbitrary number of TIDs, candidate ray paths satisfying the condition for specular reflection are found using a homing algorithm,

evaluated on a regular x-y grid of initial estimates (distributed according to the known TID wavelengths). Although the reflecting surface is specified as an analytic function, in general, for an arbitrary number of TID components and a spherical Earth/ ionosphere, the problem is believed to be analytically intractable. As such, an unconstrained numerical root finder was used to locate the x and y coordinates for which (1) the propagation path and unit normal vector to the reflecting surface at that point lie within a single plane, and (2) the unit normal bisects the upward and downward ray segments. These two criteria for specular reflection can be expressed by the following two equations, respectively:

$$\left(\frac{\boldsymbol{v}_R \times \boldsymbol{v}_T}{\|\boldsymbol{v}_R\|\|\boldsymbol{v}_T\|} \right) \cdot \widehat{\boldsymbol{n}} = 0, \tag{8.2}$$

and

$$\cos \gamma = \frac{\boldsymbol{v}_R \cdot \widehat{\boldsymbol{n}}}{\|\boldsymbol{v}_R\|} = \frac{\boldsymbol{v}_T \cdot \widehat{\boldsymbol{n}}}{\|\boldsymbol{v}_T\|}, \tag{8.3}$$

where $\boldsymbol{v}_R$ and $\boldsymbol{v}_T$ are the path vectors to the reflection point, as illustrated earlier in **Figure 4.17**, $\widehat{\boldsymbol{n}}$ is the unit normal to the model surface (directed outwards), and γ is the angle of bisection. The unit normal is simply the (normalised) surface gradient vector evaluated at the reflection point:

$$\widehat{\boldsymbol{n}} = \frac{1}{\sqrt{(\partial z_i/\partial x)^2 + (\partial z_i/\partial y)^2 + 1}} \begin{bmatrix} -\partial z_i/\partial x \\ -\partial z_i/\partial y \\ 1 \end{bmatrix}, \tag{8.4}$$

where the reflecting surface centred on $\boldsymbol{r}_0 \equiv [x_0, y_0, z_0]^T$ with spatially varying radius r_i is defined by

$$z_i(x, y, t) = \sqrt{r_i^2 - (x - x_0)^2 - (y - y_0)^2} + z_0. \tag{8.5}$$

Note that these equations form a core part of the FAS technique, described further in **Section 8.3.2**, and relate to the ray trajectory functions found later in equations (8.23)

and (8.24), under the approximation of a flat Earth/ionosphere and small-amplitude perturbation. Additional details can be found in Paznukhov et al. [2012].

While many nonlinear root-finding algorithms exist for numerical ray homing purposes, "mmfsolve" from the Mastering MATLAB® Toolbox [Hanselman & Littlefield, 2005] was applied in this case, based on Algorithm D6.1.3 from Dennis & Schnabel [1996]. A scalar cost function, based on the difference between the middle and right sides of equation (8.3), after satisfying equation (8.2), was chosen as the input objective function. Only those solutions that are not obscured (blanketed) by the mirror surface at any point along the ray path are retained; that is, partial transmission and reflection is not allowed. The Doppler shift is simply the time rate of change of geometric path length, calculated as a finite difference, under this mirror approximation.

Overhead and cross-sectional views of some of the possible reflection geometries, for monochromatic TIDs travelling both longitudinally and transversely[2] to an oblique propagation path, are illustrated in **Figure 8.4** and **Figure 8.5**. Apart from the direction, the TID parameters are identical. In each case, only the shortest ray path is shown in full, while other candidate reflection points are also marked by crosses in the top (overhead) panel. Of note is the fact that all reflection points lie along a single linear axis; to first order, this applies for any monochromatic TID propagation direction. For the transverse TID (**Figure 8.5**), there are a greater number of possible reflections for the same horizontal wavelength and, although the reflection points are perturbed by a similar distance in the x-y plane, the effect on the path length (i.e. delay and hence Doppler) is much larger than for the longitudinal TID (**Figure 8.4**). This suggests that a transverse TID should be more detectable in the ELOISE group delay, Doppler and bearing observables.

As an example of the utility of model 1, **Figure 8.6** shows a time series of CSF observables (fitted F2 peak data from Kalkarindji to Laverton, 1507 km, 39 °T) with synthesised values overlaid for a corrugated mirror moving transversely to the HF propagation plane at a mean height of 280 km (untilted). The modelled TID, constructed

---

[2] In this chapter, the terms "longitudinal" and "transverse" are used to refer to the relative alignment between the horizontal propagation direction of the AGW/TID and the midpoint bearing of the oblique HF path. This should not be confused with the oscillations of the atmospheric particles themselves, which are always transverse to the wave propagation direction (wave vector).
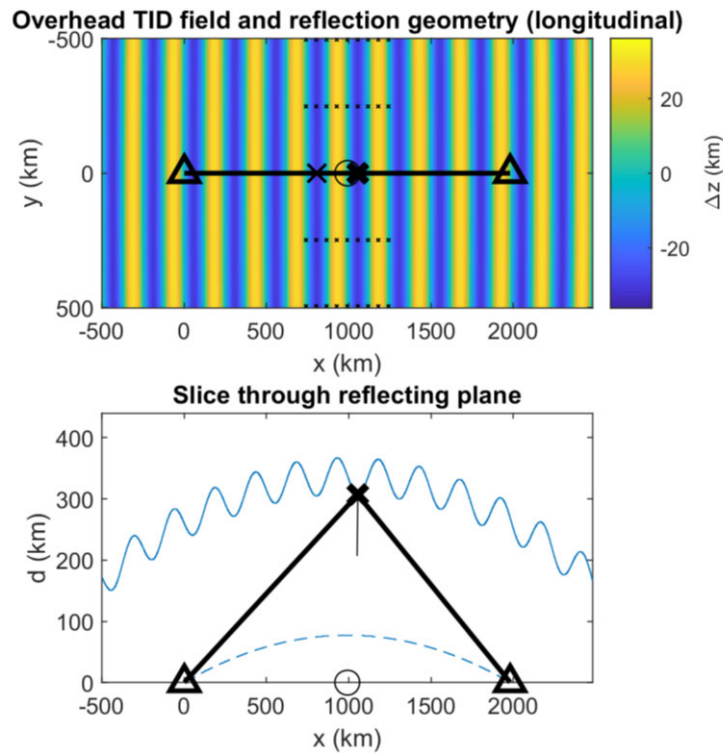
Figure 8.4. Illustrations of a constructed TID field and specular reflection geometry for the corrugated mirror model, for a single (monochromatic) TID propagating in a down-range (longitudinal) direction. The top and bottom panels show overhead and cross-sectional (slice) views of the shortest ray path, respectively. Candidate reflection points are indicated by the large 'x' markers in the top panel, which were converged upon from the initial estimates indicated by the grid of small markers; the reflection point with the shortest path length (delay) is shown in bold in both panels. Limits on the horizontal displacements of the initial estimates from the path midpoint (circle marker, projected on the x-y plane) eliminate some of the more off-angle solutions, which are less likely to be detected in practice. Whereas the coordinates are Cartesian, based on the left-handed system depicted later in **Figure 8.25** (i.e. with y-axis directed down the page), the Earth and unperturbed ionosphere are spherical, so for sake of clarity, only the perturbed component of the disturbance field Δz is shown in the top panel. In the bottom panel, the Earth's surface (dashed blue line) and ionosphere (solid blue line) are plotted as a distance $d = \sqrt{y^2 + z^2}$ through the reflecting plane. Note also, because of the highly oblique geometry (representative of the Humpty Doo to Laverton path), the aspect ratio is not one-to-one on either panel.
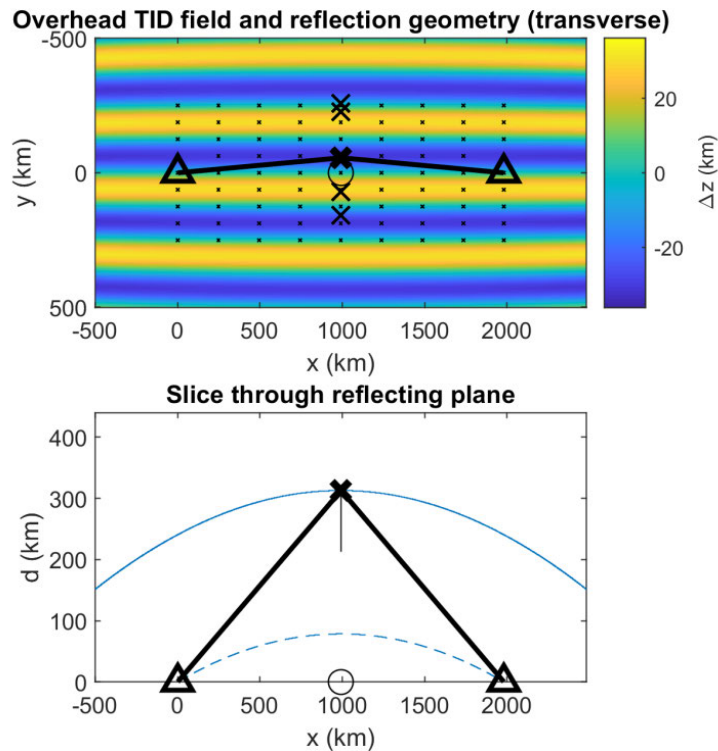
Figure 8.5. As in **Figure 8.4**, but for a TID propagating in the cross-range (transverse) direction.

by manually adjusting the parameters to match (by eye) the observables on this path, had a 15 km amplitude, 70 min period, 400 km horizontal wavelength, and midpoint bearing of -53 °T (north-westward). This period is characteristic of a large-scale TID, potentially of southern auroral origin, and similar to the disturbance signature shown previously in **Figure 5.4** of **Section 5.2**. While such disturbances were less common than the medium-scale perturbations shown in **Figure 7.1** of **Section 7.1**, similar "satellite" trace characteristics were seen on approximately a third of the nights during the month-long campaign. The example in **Figure 8.6** has been chosen for its single dominant periodicity, which makes it a less complicated modelling problem than in the aforementioned figure (although multiple TID components are supported in the model).

The apparent horizontal wavelength implicitly includes the effects of background neutral winds, and as such, it is credible to believe that the intrinsic AGW wavelength is in fact far greater than 400 km [Chum et al., 2012]; likewise for the intrinsic horizontal phase speed, which is currently an anomalously low ~95 m/s according to the model. Indeed, the empirical Horizontal Wind Model HWM14 [Drob et al., 2008; Drob et al.,

2015] predicts a strong prevailing east-north-eastward wind of over 100 m/s at this height, which largely opposes the apparent north-westward velocity of the TID model. Furthermore, if the modelled corrugation were to be extended to a "tilted" sinusoid (i.e. closer to a sawtooth wave), this could allow for a longer horizontal wavelength, more in line with the ≥1000 km LSTID wavelengths quoted in Hunsucker [1982]; as it stands, the pure sinusoid must be kept reasonably short in length to reproduce the sawtooth-like properties of the observables. Thus, it is fair to say that the precise TID model parameters should be interpreted with a degree of caution, as they are highly coupled to the model construction and an element of personal judgement.
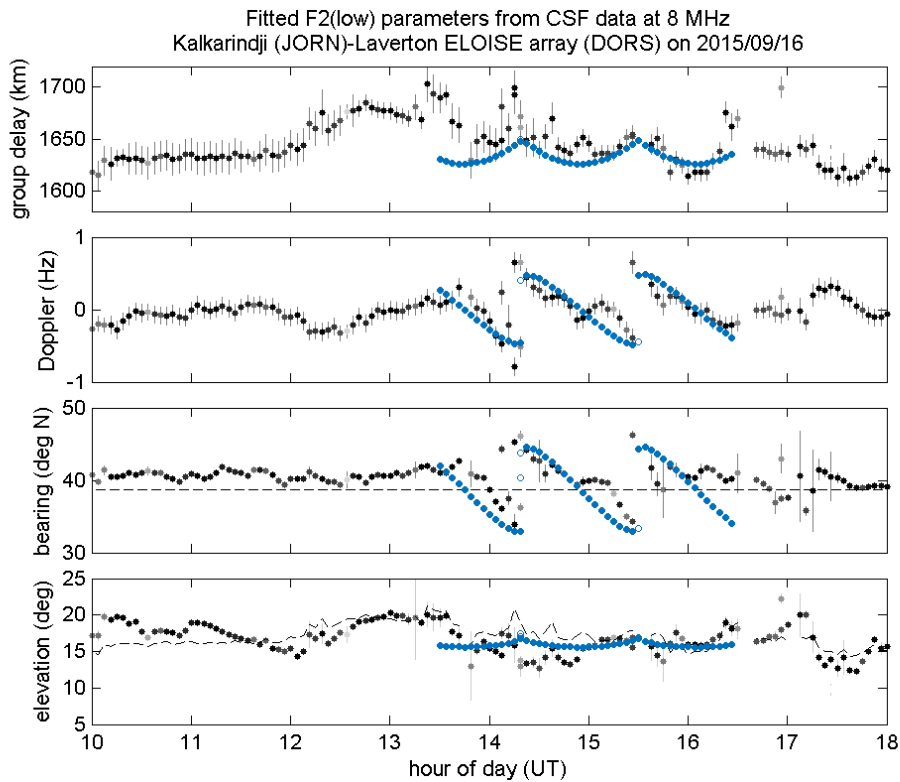
Figure 8.6. A night-time sequence of CSF peaks (grey-scale dots) for 1-hop F2-low propagation at an operating frequency of around 8 MHz. Synthesised values for the corrugated mirror TID model are overlaid on the most disturbed interval from 1330–1630 UT (blue dots, with solid markers showing the shortest path). From top to bottom, the panels are group delay, Doppler, bearing and elevation. The CSF markers are shaded according to their peak SNR, with dashed lines in the bottom two panels representing the equivalent great circle mirror geometry. Note that moderate delay spread was observed for much of this period, and at times the frequency was very close to the F2 leading edge, causing group delay estimates to be geophysically noisy.

While not a perfect match to the observations in **Figure 8.6**, the TID model is successful at reproducing the dominant features, including the magnitude of the variations in Doppler and bearing, and the sudden switchback in the propagation path from one wave period to the next. The fact that there is very little of the same quasi-periodic signature in elevation is due to the transverse geometry, as the reflection point is constrained to the transverse plane through the midpoint [Bramley, 1953]. However, with its large horizontal wavelength, the model is unable to reproduce the multiple simultaneous specular reflection points that were observed, except for one or two extra reflections at each switchback point. As noted in Davies & Baker [1966], the condition for these off-angle multi-reflections is that the radius of curvature of the reflecting surface must be less than the height of reflection. This implies a larger amplitude-to-wavelength ratio for the TID, although in this case neither a larger amplitude nor smaller wavelength could be found to acceptably reproduce the CSF observables in other respects.
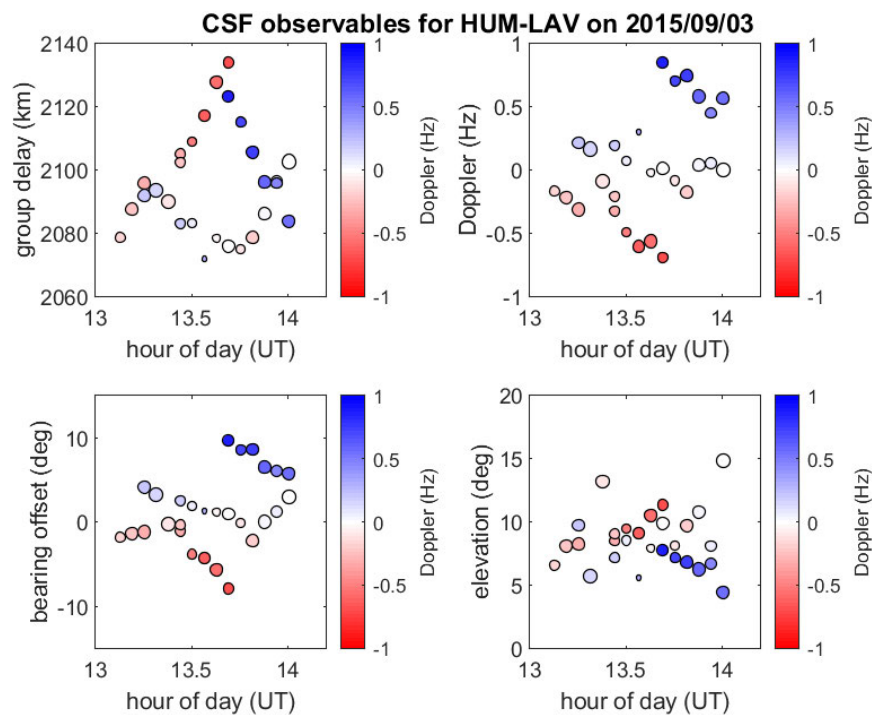
Figure 8.7. Peak observables for the Humpty Doo to Laverton CSF path, on a single ~5.7 MHz channel from around 13–14 UT on 3 September 2015. From top to bottom and left to right, the panels show group delay, Doppler, bearing offset, and elevation. This period corresponds to the evolution of satellite traces shown in **Figure 5.3** of **Section 5.2**. Markers are colour-coded by Doppler velocity and scaled in size by SNR.

One of the best examples of multiple off-angle reflections (satellite traces) was observed on 3 September 2015, and was shown previously in **Figure 5.2**, **Figure 5.3** and **Figure 5.4** of **Section 5.2**. The CSF observables for a ~1 hour subset, over which one of the frequency channels remained at a stable ~5.7 MHz, are plotted in **Figure 8.7**. Markers have been colour-coded by Doppler, to highlight the advancing (blue) and receding (red) parts of the signature. Although this only depicts a little over one period, the symmetrical rise and fall of off-angle traces in group delay is clearly captured, along with the associated offsets in Doppler and bearing. Elevation is the only dimension in which the extra traces cannot be easily resolved.

Figure 8.8 shows a synthetic reproduction of the observed signature from **Figure 8.7** using a transverse corrugated mirror model (as in **Figure 8.5**). Again, the TID parameters have been selected manually. It is arguable whether this represents an MSTID or LSTID, given the period (40 min) and horizontal wavelength (250 km), but
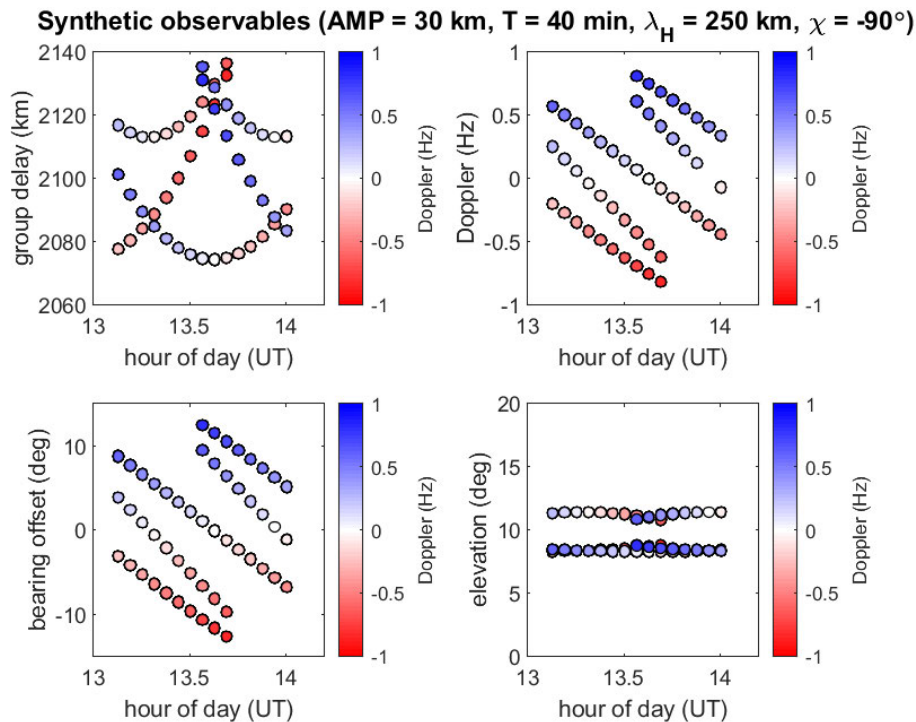
Figure 8.8. Synthetic observables for the corrugated mirror TID model, with amplitude 30 km, period 40 min, horizontal wavelength 250 km, and bearing -61° T (-90° with respect to the great circle midpoint bearing). The unperturbed F2 mirror height was 260 km. The panels are presented the same as in **Figure 8.7**, except for the fact that path loss has not been calculated, so all markers are a fixed size.

the key point is that the amplitude is sufficiently large (30 km) to support satellite traces and, thus, replicate the main characteristics of the observations. The shortest path, which is also likely to be the strongest in SNR, tracks one moving trough on the wave train, before jumping to the next when the intermediate crest passes overhead the midpoint. The extra "wings" between 2110 and 2120 km in the group delay (top left) panel correspond to reflections from wave crests and in practice are not observed, indicating the potential shortfalls of such a simple model.

It is not computationally difficult to synthesise more complicated examples, composed of multiple simultaneous TIDs, but quite quickly the synthetic observables deteriorate from having a periodic form to resembling noise. The results, shown in **Figure 8.9** for just two identical waves travelling orthogonally to each other, can begin to mimic mid-latitude spread-F conditions in terms of the multitude of reflections, reinforcing the
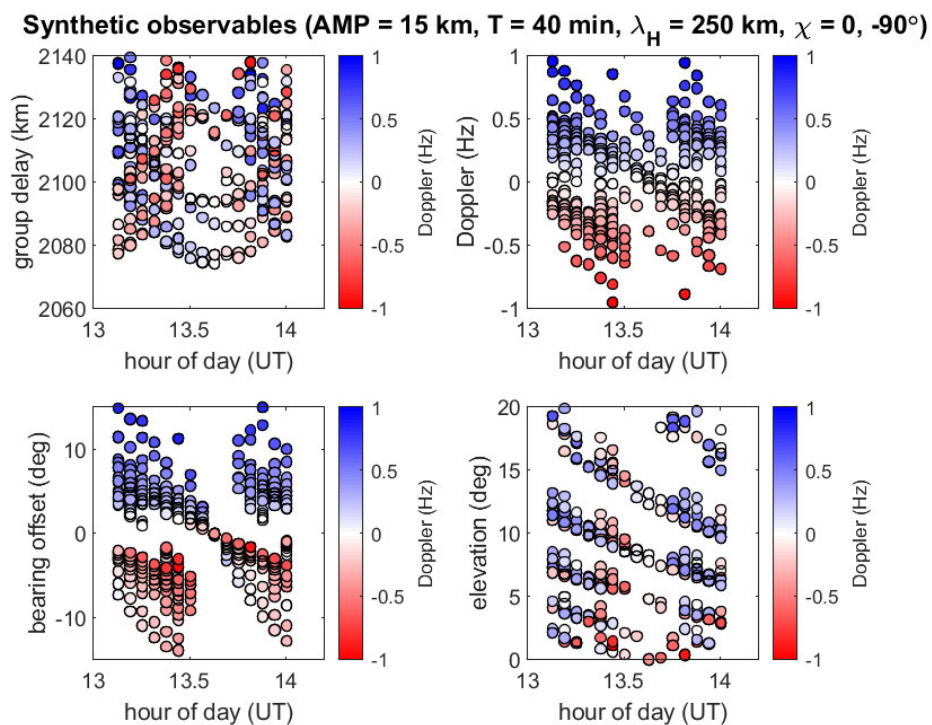
Figure 8.9. Synthetic observables for the corrugated mirror TID model, containing the superposition of two identical wave components with orthogonal directions: one longitudinal to the propagation path (as in **Figure 8.4**) and one transverse (as in **Figure 8.5**). Other TID parameters are the same as in **Figure 8.8**, except with half the amplitude (i.e. 15 km) assigned to each component.

view that this phenomenon may be closely related to the presence of TIDs and other tilted ionospheric structure [King, 1970; Bowman & Monro, 1988; Bowman, 1991a; Bowman, 1991b]. It is worth noting that while there is technically no limit as to the number of TID components in this model, the numerical solver may not converge as reliably and can be more sensitive to the initial estimates under very complicated disturbance fields.

More generally, the corrugated mirror model can be used to study the relationships between key TID parameters and the observables, and thus assist in broadly classifying disturbance signatures in the peak data. **Figure 8.10**, **Figure 8.11** and **Figure 8.12** show the impact of varying the TID amplitude, horizontal wavelength and bearing,
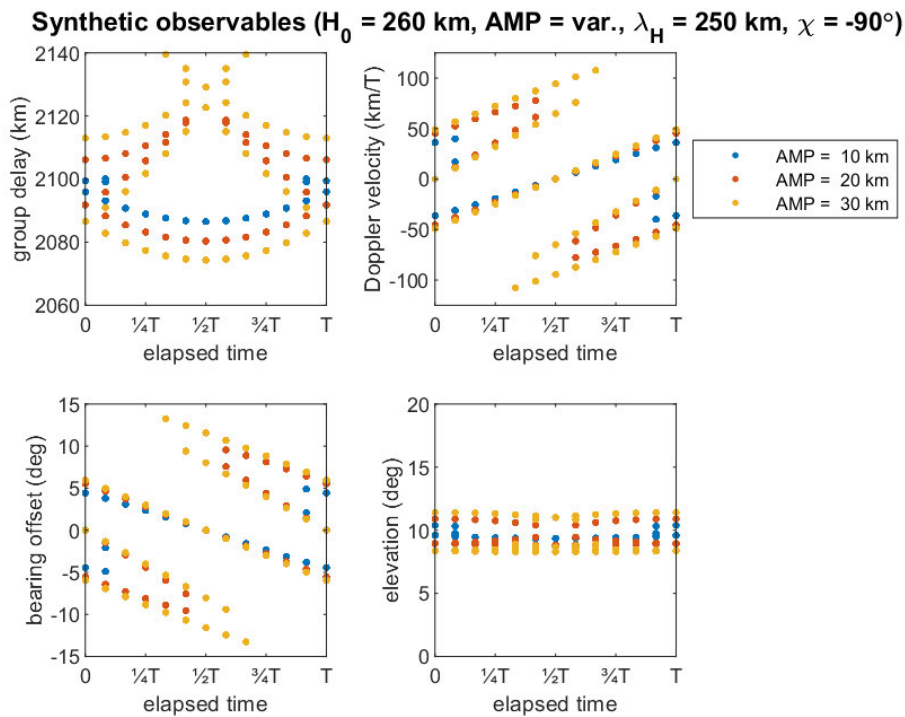
Figure 8.10. Synthetic observables for three generic TID model instantiations with different perturbation amplitudes: 10, 20 and 30 km, in blue, red and yellow, respectively. All other wave parameters, and the oblique ground range (1989 km), were fixed to the same values as the transverse TID in **Figure 8.8** (with period T scaled out of the time and Doppler axes). As the amplitude is increased, a greater number of simultaneous off-angle reflections become visible, and there is a near-proportional increase in the magnitude of the maximum bearing and elevation offsets (although the latter remains relatively small for the transverse geometry).
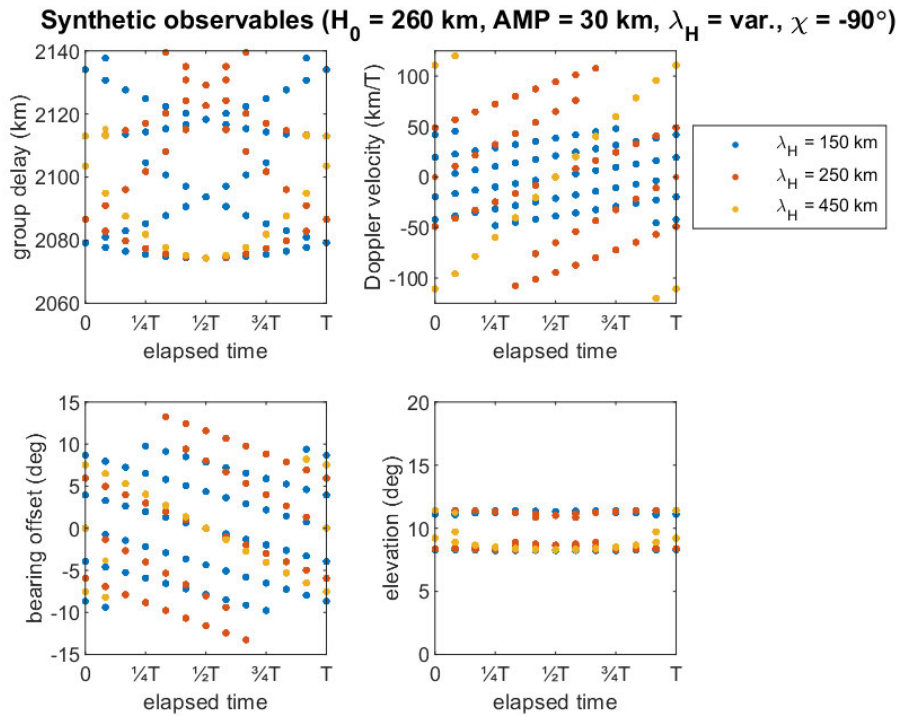
Figure 8.11. Synthetic observables for three generic TID model instantiations with different horizontal wavelengths: 150, 250 and 450 km, in blue, red and yellow, respectively. Again, all other wave parameters, and the oblique ground range, were fixed to the same values as the transverse TID in **Figure 8.8**. As the wavelength is increased, there is a reduction in the number of off-angle reflections, while the Doppler velocity naturally increases with the faster horizontal propagation speed.

respectively, starting with the same model as in **Figure 8.8**. The operating frequency and TID period have been deliberately scaled out of the results, with the time axis in units of the period and the Doppler velocity in units of km per period. As the wave amplitude increases, or horizontal wavelength decreases, the radius of curvature of the reflecting surface becomes smaller and the number of simultaneous specular reflections rises. The maximum magnitude of the perturbation in each observable also increases with increasing amplitude, while changing the wavelength mainly affects the magnitude of the Doppler velocity, as it is coupled to the propagation speed.

Of particular note is the fact that the scale of the perturbations depends strongly on the TID direction (**Figure 8.12**); this is the directional bias inherent in oblique incidence observations [Clarke, 1972, p. 268]. Wave bearings that are aligned longitudinally with

the oblique path bearing (i.e. relative azimuth $\chi \approx 0°$ or $180°$) are fundamentally more difficult to detect, as the perturbation signature is primarily in elevation (the least precise observable), whereas those that are transverse (i.e. $\chi \approx \pm90°$) have more pronounced signatures in all other observables. The difference becomes more accentuated the longer the oblique path is, reinforcing the idea that transverse TIDs are much more detrimental to OTHR coordinate registration, based on delay and azimuth measurements only [Barnes et al., 2000]. This finding emphasises the value of path orientation diversity in an experiment such as ELOISE, where the aim is to characterise the full directional spectrum of disturbances.
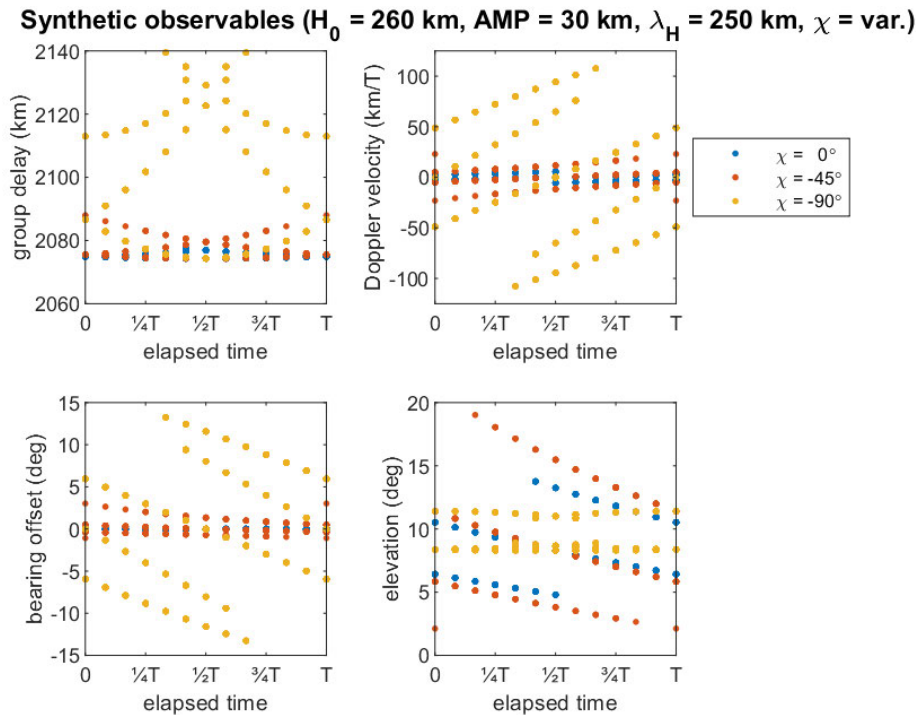
Figure 8.12. Synthetic observables for three generic TID model instantiations with different propagation directions: 0°, -45° and -90°, from longitudinal to transverse, in blue, red and yellow, respectively. Again, all other wave parameters, and the oblique ground range, were fixed to the same values as the TID in **Figure 8.8**. As the TID direction becomes more closely aligned with the oblique path bearing itself ($\chi \to 0$), the scale of the perturbations decreases in all panels except elevation. Among these three cases, the largest elevation perturbations are observed at -45°.

In all of these synthetic results, it is important to acknowledge that the passage of a sinusoidal reflecting surface does not, in general, produce sinusoidal variations in the observables; only for short (quasi-vertical) paths, where the horizontal wavelength of the TID is comparable to or larger than the ground range, and for small TID amplitudes can this occur. Although Huang et al. [2016] used a similar model implementation to produce examples of quasi-sinusoidal observables, their simulations were for a relatively short path (446 km) and extremely small TID amplitudes (0.3–2 km), most of which would be unlikely to be detected by any ionosonde, with or without an AoA measurement capability. Even in such cases, the functional form of the observables was often distinctly sawtooth in shape. This point will become particularly relevant later in the discussion about parameter inversion techniques (**Section 8.3**).

8.2.2 *Atmospheric gravity wave seeded (Hooke) model*

The AGW-seeded approach (model 4), originally developed by Hooke [1968; 1970], applies linearised perturbations to the continuity equations governing ion production, loss and transport, and is combined with 3D magneto-ionic ray tracing [Cervera & Harris, 2014] to produce a TID representation that is more physically realistic than the corrugated mirror model. Although more computationally demanding, it supports multiple ionospheric layers/modes (e.g. using the International Reference Ionosphere as the background electron density field) and the iso-ionic contours correctly capture the characteristic forward tilt of the AGW/TID phase-front, unlike the corrugated mirror model. Over time, this appears as a height-descending phase progression, an example of which is illustrated later in **Figure 8.15**. Both O- and X-mode components can be synthesised using this approach.

The derivation of the governing physical equations for the dynamics and propagation of TIDs under the Hooke model, starting from the fundamental dispersion relation for acoustic-gravity waves, is summarised in **Appendix D**. It is based on projecting the motion of the AGW along the magnetic field lines that constrain ion and electron motion and calculating the divergence to give the electron density perturbation $N_e'(\boldsymbol{r}, t)$ in space and time. Photochemical effects, caused by modulations to the neutral density and solar flux, have a less important role at F2 heights. This theory yields both the full lossless form of the TID perturbation in terms of the AGW frequency and wave vector

(**Appendix E**, eq. (E.14)) and the standard Hooke approximation (eq. (E.17)); the former is used in the simulations presented here.

Whereas an earlier DST Group study by Cervera & Harris [2014] focused on vertical-incidence geometry, the existing code (provided by Dr Manuel Cervera) for generating synthetic ionograms had to be generalised by the current author for application to the ELOISE oblique paths. This included an updated ionospheric profile specification (to align with the output of the ELOISE on-board processing), modified ionospheric and geomagnetic sampling grids, the use of height-dependent atmospheric parameters (including the Brunt-Väisälä frequency), and general speed improvements and bug fixes. Equations (E.3) and (E.14) from **Appendix E** are now also used in place of the less accurate equations (E.6) and (E.17) for calculating the vertical wave number and electron density perturbation.

Returning to the example in **Figure 8.6**, a set of synthetic observables generated by ray tracing through the Hooke TID model at 8 MHz is shown overlaid on the CSF peak data in **Figure 8.13**. Profiles of the neutral atmospheric parameters, computed from NRLMSISE-00, and the background ionosphere, scaled in terms of plasma frequency, are plotted in **Figure 8.14**, while the ionospheric perturbation field as a function of time is plotted in **Figure 8.15**. The background ionosphere is a fixed spherically-symmetric electron density profile based on the mean quasi-parabolic parameters fitted to the surrounding 3-hour ionogram sequence (i.e. interleaved with the CSF observations in **Figure 8.13**). It is assumed that the unperturbed neutral atmosphere is motionless, as the background horizontal winds at F2 heights are not known with sufficient certainty.

In the F2 region, the neutral scale height is around 40 km (**Figure 8.14**, first panel), meaning that in the absence of dissipation, the AGW amplitude increases by $\exp(1/2)$ over a ~40 km span in altitude; as such, relatively small-amplitude waves in the lower atmosphere can become significant at ionospheric heights. For internal gravity waves, the (intrinsic) angular frequency must satisfy $\omega < \omega_B < \omega_a$, where ω_B and ω_a are the Brunt-Väisälä and acoustic cut-off frequencies, respectively, which imposes a lower limit on the wave period of ~13 min at 250 km (**Figure 8.14**, third panel).

The AGW perturbation parameters were chosen to synthesise a broadly similar class of disturbance as the one in **Figure 8.6** (using the corrugated mirror model). The modelled large-scale TID is again moving transversely to the HF propagation plane (-53 °T) with

a period of 70 min and horizontal wavelength of 400 km, and with amplitude governed by the neutral wind perturbation and dissipative terms. The magnitude of the projection of the wind perturbation (~15 m/s) in the direction of the geomagnetic field vector at the F2 reference height (250 km) is equal to 7 m/s. Although not formally included in the Hooke model, wave dissipation effects due to molecular viscosity, thermal diffusion and ion drag are partially accounted for by reducing the imaginary component of the vertical wave number (as in **Appendix E**, eq. (E.5)).
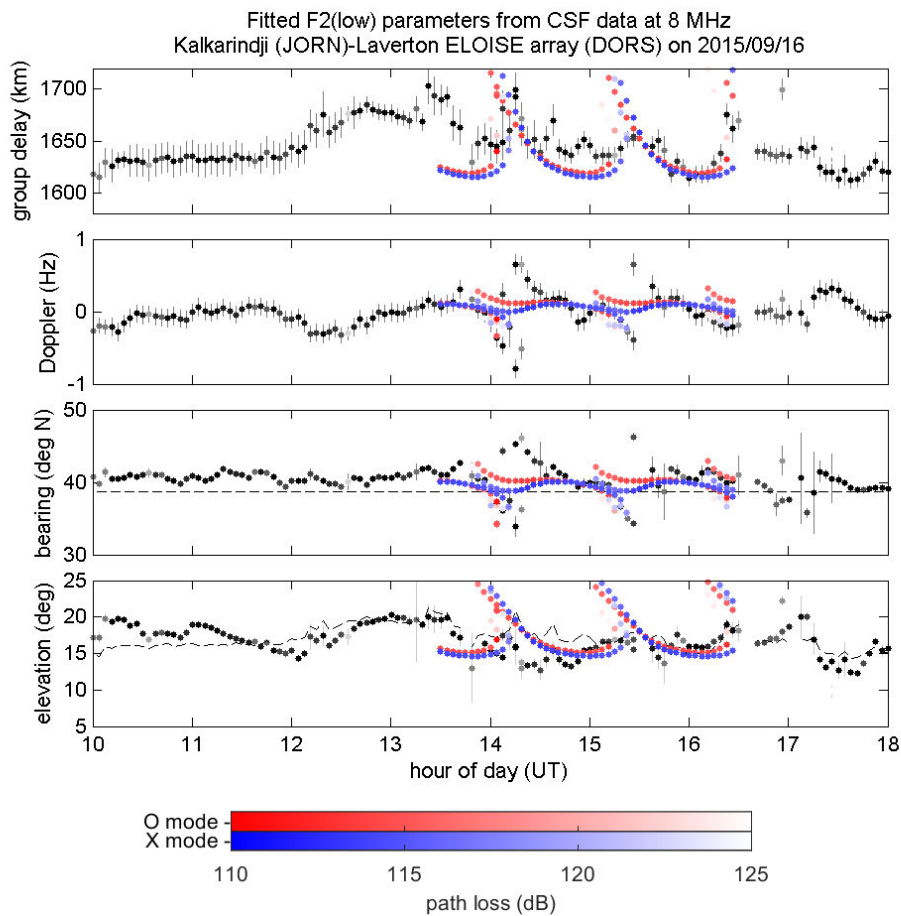
Figure 8.13. Observables synthesised from the AGW-seeded (Hooke) TID model in **Figure 8.15** (red/blue dots), overlaid on the CSF observations from **Figure 8.6** (grey-scale dots). Again, from top to bottom, the panels depict group delay, Doppler, bearing and elevation, for a fixed operating frequency of 8 MHz. Both O-mode (red) and X-mode (blue) components are calculated by 3D numerical ray tracing [Cervera & Harris, 2014], with colour saturation indicating the estimated path loss. Note that a very similar plot appeared as Figure 16 in Heitmann et al. [2018], however this was using the standard Hooke approximation and included an error in the call to NRLMSISE-00.
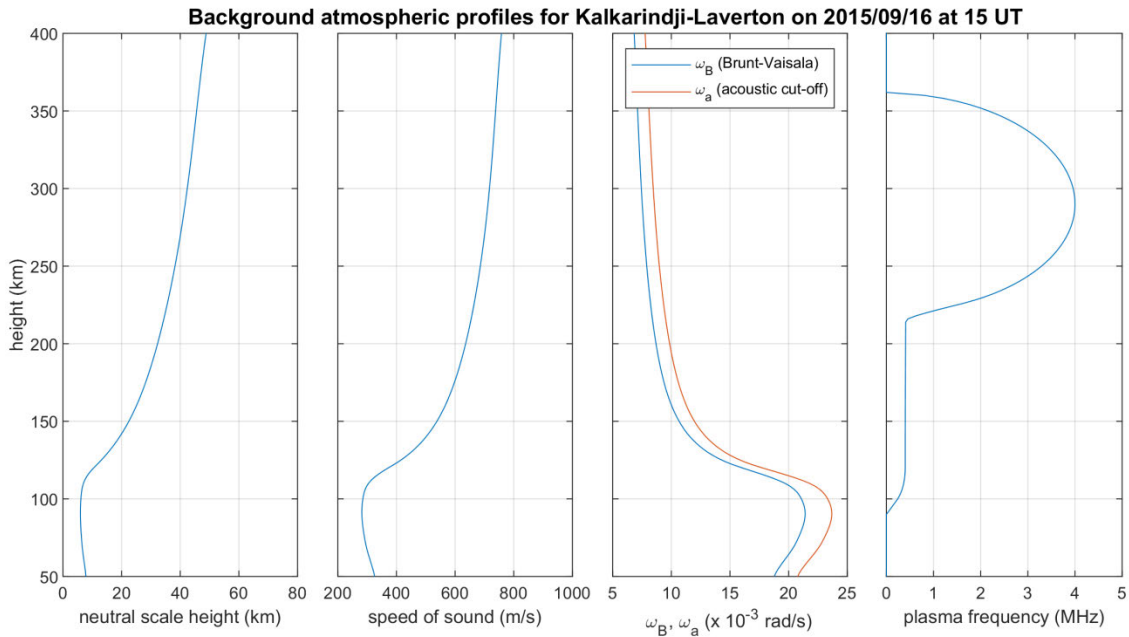
Figure 8.14. Profiles of neutral atmospheric parameters computed from NRLMSISE-00 for the unperturbed (great circle) midpoint at 15 UT, and the background (mean) plasma frequency profile fitted to ionograms from Kalkarindji to Laverton over the surrounding 3-hour period.
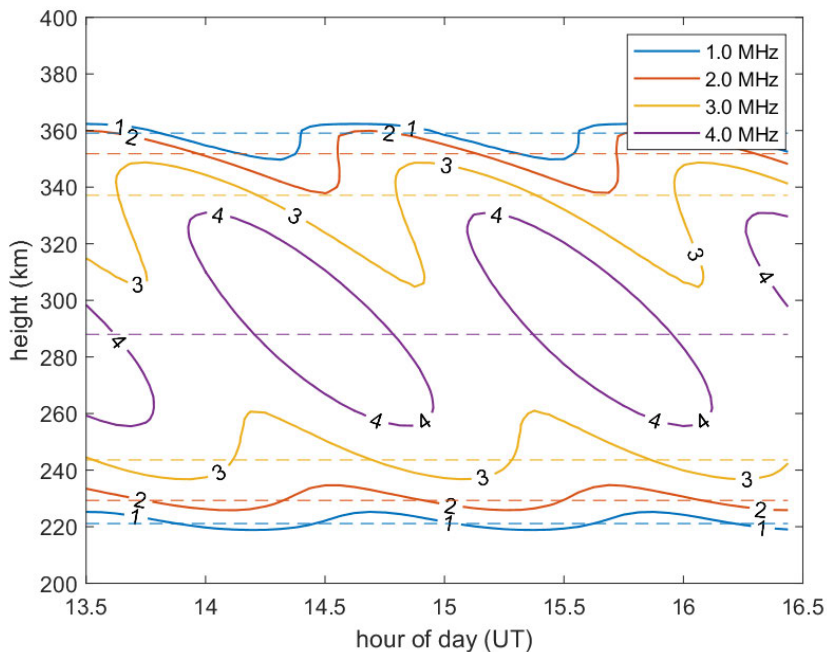
Figure 8.15. Illustration of perturbed iso-ionic contours (i.e. lines of constant plasma frequency in MHz), overhead a fixed point for the AGW-seeded (Hooke) model. The background (unperturbed) heights are indicated by the dashed horizontal lines.

Not all scales of AGWs are able to penetrate to F2 heights [Hines, 1960], and in this night-time example, it appears to be difficult to realistically reproduce the full amplitude of the observed TID signature with the AGW-seeded model, without resorting to physically unrealistic perturbation parameters. This in turn would trigger anomalously high/low electron density crests/troughs. The modelled field in **Figure 8.15** contains perturbations of up to 15–17% from the background plasma frequency profile at around 250 km; this is thought to be within the typical amplitude range of naturally occurring TIDs [Thome, 1964], and comparable to the deviations of ±15% that Nickisch et al. [2006] arrived at using the Bottone [1993] model for a simulated impulsive detonation. However, Doppler and bearing excursions in the peak data tend to be underestimated by the Hooke model in **Figure 8.13**, while group delay excursions are overestimated. Many of the higher elevation synthetic features, for rays approaching the F2 peak height and undergoing stronger refraction, are also not readily apparent in the observations, despite the additional estimated path loss being less than 15 dB (within the SNR margin of the instrument). Further work is required to better understand the source of these discrepancies and the limits of the Hooke model over a broader range of TIDs, some of which may not in fact be AGW-generated.

Note that for this 1507 km path, the ray apogee heights for an 8 MHz operating frequency, plotted in **Figure 8.16**, correspond roughly to the 3 MHz plasma frequency contour in **Figure 8.15** (yellow line centred about 240 km height). Even for the low ray, this puts it slightly beyond the sinusoidal contour regime, where the corrugated mirror model might ordinarily be expected to have most validity (at least, without modifying the functional form of the perturbation). Conversely, it means that at lower frequencies/ heights, or for lower amplitude perturbations, the mirror model may remain a good approximation. This is not true of the high ray, for which the apogee delves much deeper into the F2 region and almost tracks the peak height (i.e. within the purple 4 MHz contours of **Figure 8.15**). The fact that the operating frequency is quite close to the MOF (nose) of the ionogram accentuates these non-sinusoidal high-ray features.

It is also worth noting that the forward tilt angle, between the direction of phase propagation and the local horizontal, for such an LSTID is relatively steep (around 80° in the F2 region), and for the Australian region, will be greater than the underlying AGW, as discussed in **Appendix E**. The tilt angle as a function of height is plotted in **Figure 8.17** for both AGW and TID, alongside the plasma frequency profile and height-

dependent TID-AGW phase offset. For a 70 min wave period, the AGW tilt is very close to the asymptotic limit of ~79° at 250 km predicted by equation (E.6) in **Appendix E**: $\psi \approx \mathrm{atan}\left(\sqrt{\omega_B^2/\omega^2 - 1}\right)$ [Hines, 1960, eq. (4); Clarke, 1972, eq. (7.18)].
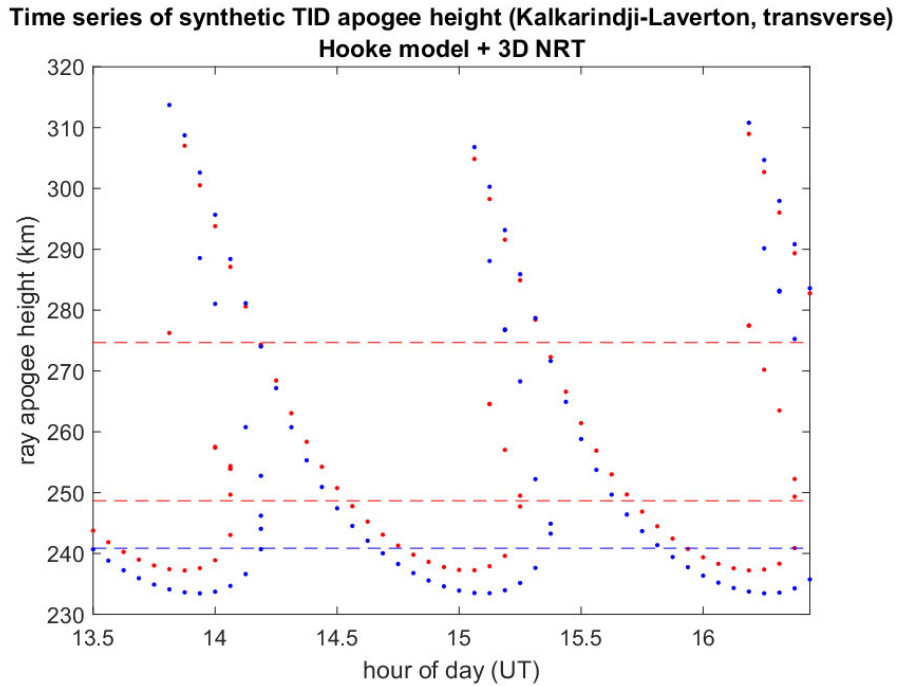
Figure 8.16. Time series of ray apogee heights (from 3D NRT) at a fixed operating frequency of 8 MHz, corresponding to the plasma frequency perturbation in **Figure 8.15** and the synthetic observables in **Figure 8.13**. The red and blue markers depict O- and X-mode returns, respectively. The unperturbed rays (for the background ionosphere only) are indicated by the dashed horizontal lines of the same colours, with the low rays appearing at 240–250 km and high rays at 270+ km (although only an O-mode solution was found for the latter case).

Three key areas in which the Hooke model could be further refined are: (1) allowance for prevailing neutral winds (i.e. characterising the AGW by its intrinsic frequency); (2) better compensation for wave dissipation effects (e.g. by incorporating a more realistic dispersion relation [Pitteway & Hines, 1963; Clark, 1970; Morgan & Tedd, 1983; Tedd et al., 1984; Vadas & Fritts, 2005; Godin, 2015; Emmons et al., 2020]); and (3) support for horizontal gradients and time variations in the background atmospheric

profiles. The challenges posed by items (1) and (2) in particular, as well as non-linearities in the wind system more generally, are discussed by Hines [1965]. The above refinements were deemed to be outside the scope of this initial demonstration, given that many authors have already considered this problem and developed substantially more sophisticated (non-linear) models over many years; for example, the 3D global-scale non-hydrostatic AGW ray tracing of Eckermann & Marks [1997], which couples in a full time-varying atmospheric model, and the AGW/TID interaction model of Kirchengast [1996], which retains additional non-linear terms to avoid artificial damping.
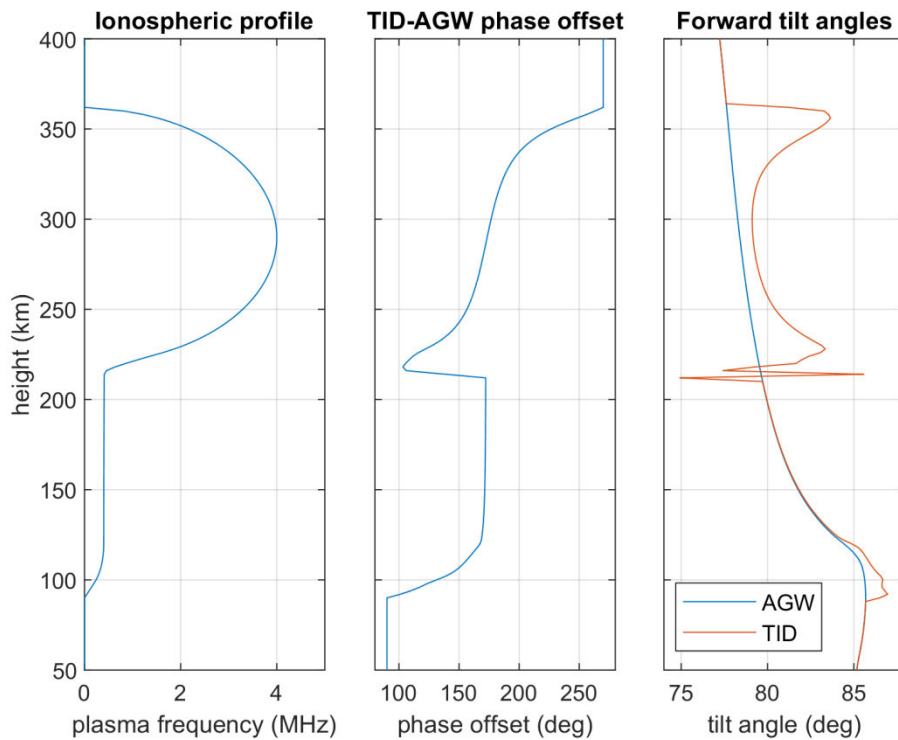
Figure 8.17. Phase offset and forward tilt angle profiles (middle and right panels, respectively) for AGW and TID under the Hooke model. The plasma frequency profile (left panel) is also included for reference. The standard Hooke approximation in equations (E.18) and (E.19) of **Appendix E** was used to compute these functions.

8.3   Parameter inversion techniques

While the parameterised TID models described above provide useful insights into the propagation of individual TIDs and their HF observational responses, any meaningful statistical analysis, along with the problem of real-time TID detection and correction in HF AoA data, clearly benefits from a means of automatic parameter inversion. Many authors have previously developed such techniques, predominantly for networks of vertical incidence sounders, including:

- True height analysis, based on the estimation of lags in the iso-ionic (or iso-height) contours via cross-correlation or cross-spectral functions (e.g. Klostermeyer [1969]; Morgan et al. [1978]; Ballard [1981]; Morgan & Tedd [1983]; Tedd et al. [1984]; Bowman [1992]; Emmons et al. [2020]). Usually this relies on a closely-spaced network of sounder sites, although Emmons et al. [2020] have described how the AGW dispersion relation can potentially provide for single-site sounder estimation of the horizontal TID wavelength, albeit without any directional information. Scanned incoherent scatter radar observations can also be used to similar effect (e.g. Thome [1968]; Hearn & Yeh [1977]; Ratovsky et al. [2008]), and with a realistic dispersion relation, this approach can provide parameters for both the TID and the causative AGW [Hearn & Yeh, 1977; Tedd et al., 1984].

- Virtual height or Doppler analysis, which is just a simpler form of the above; that is, estimating lags in the virtual height (group delay) or Doppler features directly (i.e. without performing profile inversion), often because only a single frequency is observed. Examples of using virtual height can be found in the early work of Munro [1948; 1950; 1958], Clarke [1972, ch. 6.4.2], and more recently, MacDougall et al. [2011] and Harris et al. [2012], while Crowley & Rodrigues [2012] and Chum et al. [2014] have made effective use of continuous-wave Doppler measurements. A similar approach applied to elevation estimates from a CSF beacon network has been reported by Knippling [2018]. It is also possible to cross-correlate the F2 critical frequency parameter from ionograms (e.g. Thome [1968]; Maeda & Handa [1980]), but like Doppler or elevation on a single frequency, this gives no information about the height structure of the TID, and is somewhat more difficult to interpret in terms of vertical wave displacements.

- Direct geometrical transformations, based on reflection from some class of corrugated mirror model, along with knowledge of AoA and Doppler. The Frequency and Angular Sounding (FAS) technique [Beley et al., 1995; Galushko et al., 2003; Paznukhov et al., 2012] is one such approach, which has been applied to the Net-TIDE project [Reinisch et al., 2018] in recent years. Simpler first-order geometrical arguments (e.g. Bramley [1953]) may also offer some utility for a more limited scope of TID characterisation.

- Joint profile inversion and horizontal gradient (tilt) estimation [Zabotin et al., 2006], from which TID parameters as a function of height can be inferred [Oliver et al., 1994; Negrea & Zabotin, 2016; Negrea et al., 2016]. This technique operates on the proviso that the horizontal TID wavelength is much greater than the spatial spread of reflection points (an issue discussed by Jacobson [1994] in the context of incoherent scatter radar, and by no means unique to this approach). Such an assumption is likely to be compromised at oblique incidence, although there is little doubt that TIDs do impart a signature on oblique tilt estimates (as shown in **Chapter 7**).

- Iterative procedures, which numerically fit a model of the disturbance field (usually an effective mirror surface composed of one or more waves) according to some optimisation criterion (e.g. Brownlie et al. [1973]; From et al. [1988]). Convergence may or may not be guaranteed. Since these algorithms have shown limited success beyond inverting synthetic data, they will not be considered any further here.

Regardless of the technique, a constituent TID wave is usually characterised by an amplitude, period (and phase), horizontal wavelength, and bearing. Where these approaches differ is in their ability to successfully decompose a complex (multi-component) wave field, and their underlying TID model and propagation assumptions. In order to be suitable for the ELOISE AoA data, the assumptions must remain valid for oblique paths of up to ~3000 km; this usually calls for a spherical Earth/ionosphere.

True height analysis falls short in that it neglects off-angle returns when constructing the ionospheric profile. Instead of relying on AoA, the TID parameters are deduced from spatial relationships between different iso-ionic contours (heights) and different sounders (horizontal sample points). For oblique incidence sounders, these horizontal sample points move around considerably and, without resorting to a more rigid form of

profile fitting [Heitmann & Gardiner-Garden, 2019], the fidelity of the true height profile inversion is likely to be poorer without polarimetric O-X separation.

As described in **Section 4.4**, virtual heights and equivalent vertical frequencies can be calculated for a tilted mirror model with regard to the AoA, and these exhibit similar descending phase signatures as in the true heights and plasma frequencies; see for example **Figure 8.18**, on which a TID-like trough feature is annotated with label (a). However, for a relatively constant forward tilt angle, the phase descent is no longer linear as a function of virtual height, owing to increased refraction at greater heights,
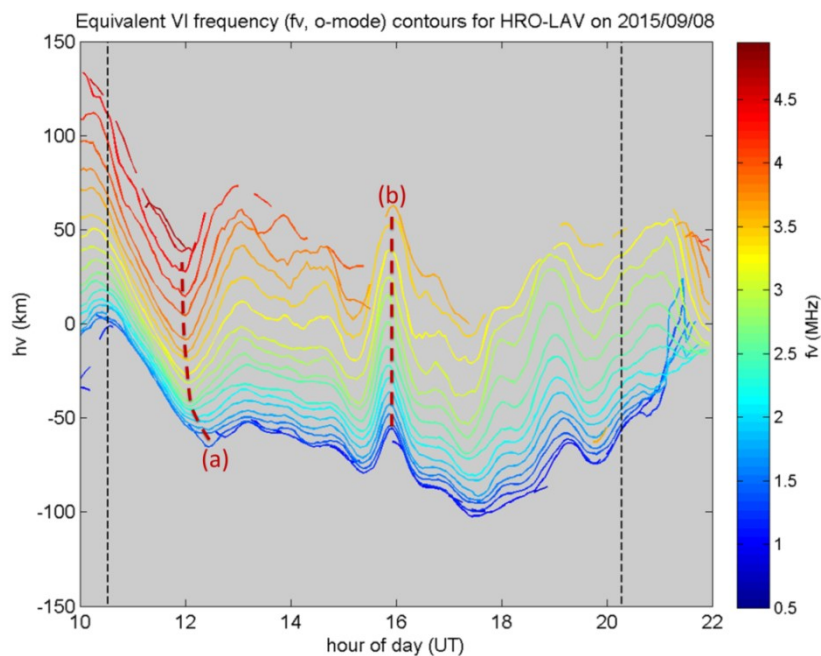
Figure 8.18. Equivalent vertical-incidence frequency (f_v) contours, as a function of hour-of-day and detrended virtual height (h_v), for 12 hours of ELOISE AoA ionogram peaks that have undergone reflection point mapping, along with some mild temporal smoothing. Under a mirror assumption, f_v can be thought of as a first-order proxy for the plasma frequency. A rudimentary O/X classification algorithm, based on identifying the cross-over point on the F2 low-ray trace, has been applied to isolate the F2 O-mode; however, it is noted that the oblique O-mode is still subject to geomagnetic effects. Detrending of h_v has been carried out using the International Reference Ionosphere's hmF2 model, leaving just the anomaly component. Two features have been annotated with red dashed lines: (a) a TID-like descending phase signature, and (b) a bulk movement of the F2 layer (e.g. associated with a change in the background neutral winds). Dusk/dawn terminators are indicated by the black dashed lines.

and the problem of horizontal displacement in the sampling point remains. As virtual height is a ray-integrated measure, an extra temporal phase shift will also be imposed on the signature [Pederick et al., 2017], although this is fairly small for most LSTIDs: <5% of the period for typical vertical wavelengths equal to or greater than the F2 layer semi-thickness.

In contrast, direct geometrical transformations simplify the propagation model to mirror reflection, but in doing so, make no assumptions about the ionospheric profile and therefore can be applied equally well to single-frequency observations, such as the ELOISE CSF peak data. As the Hooke TID model inherently relies on the electron density distribution, it cannot be inverted using geometry alone. With any mirror model approximation, delay and Doppler are in theory redundant quantities (with Doppler velocity interpreted as the rate of change in geometric path length), but in practice this can cause a contradiction due to the distinction between group and phase delay. As such, only one of either the group delay or Doppler observables is typically used; in the Lowell Digisonde implementation of the FAS technique, Doppler is favoured over delay for its greater precision. What this means in the ELOISE context is that the same inversion technique can also potentially be applied to ionogram peak data (i.e. without Doppler), although the use of group delay rather than Doppler may make the solution more susceptible to errors in the assumed mirror geometry [Obenberger et al., 2019].

Whether carried out manually or automatically, the procedure for determining TID parameters from the ELOISE ionogram or CSF peak data is much the same and faces similar challenges. The basic steps are:

1   Extract the TID period (and corresponding phase) from time sequences of the delay, Doppler and/or AoA observables. For a near-monochromatic TID field, this is simple and can be done by eye (i.e. based on the separation between periodic features in the observables); in other cases, spectral analysis is required. Although the fundamental periodicity in the observables should always be the same as that of the TID, in many cases the observed signature will not resemble a sinusoid, and may have multiple simultaneous reflections, so traditional Fourier analysis may not be suitable.

2   Once a single TID component has been isolated, the bearing of the TID can in theory be estimated from the axis along which the reflection points lie; that is,

using the measured delay and AoA, in conjunction with knowledge of the time evolution or the sign of the Doppler offset to resolve the inherent 180° directional ambiguity. In general, the axis of the reflection point perturbation is not the same as the horizontal phase propagation direction of the TID (i.e. the normal to the phase-front), but they are inherently related by the path geometry.

3   The amplitude and horizontal wavelength of the TID are determined last, being the most difficult parameters to separate in the observables. As noted earlier, it is the amplitude-to-wavelength ratio that determines the number of simultaneous reflections. For the given period and bearing, the horizontal wavelength of the TID is most discernible from the scale of variations in the Doppler observable, as it relates directly to the horizontal phase speed, with the amplitude of the TID largely responsible for the remainder of the variations in the other observables.

In the sub-sections that follow, two geometrical inversion approaches – namely, the Bramley ripple transformation and FAS technique – are evaluated for their suitability to the ELOISE data set, while drawing upon the synthetic corrugated mirror model results of **Section 8.2.1**.

8.3.1 Bramley ripple transformation

The brief appendix from Bramley [1953] analytically derives the geometry for specular reflection from an ionospheric "ripple" (sinusoidal corrugated mirror), assuming a flat Earth/ionosphere and a perturbation amplitude that is much smaller than the reflecting height and horizontal wavelength. One of the key results is a simple expression for the maximum amplitude and horizontal axis of the reflection point perturbations in terms of the TID amplitude and horizontal axis (or 180°-ambiguous bearing). This offers a possible method for extracting TID information from the mapped ELOISE reflection points (**Section 4.4**), even if only to yield a subset of first-order parameter estimates.

The so-called Bramley ripple transformation is summarised below, with a focus on its application to oblique reflection point data. The origin of the Cartesian coordinate system used in all the equations is taken to be the unperturbed (great circle) midpoint on the surface of the Earth, with the x-axis aligned in the direction of the unperturbed midpoint bearing.

For propagation over ground range D (km) via an effective reflection height h_0 (km), horizontal perturbations in the reflection points, relative to the (unperturbed) path midpoint, are approximately given by [Bramley, 1953, eq. (3) and (4)]

$$x_p = r_p \cos \alpha_p = kA \left(h_0 + \frac{D^2}{4h_0} \right) \cos \chi^* \sin \varphi_p, \tag{8.6}$$

and

$$y_p = r_p \sin \alpha_p = kAh_0 \sin \chi^* \sin \varphi_p, \tag{8.7}$$

where r_p and α_p are the time-varying displacement (km) and fixed azimuthal axis (degrees) of the perturbed reflection points in the horizontal plane, while A, k, χ^* and $\varphi_p = \varphi(x_p, y_p, t)$ are the amplitude (km), horizontal wave number (km^{-1}), azimuthal axis (degrees) and spatio-temporal phase (degrees) of the TID. For notational convenience, the subscript "h" has been dropped from the horizontal wave number of the TID, but this should not be confused with the wave number of the HF signal itself. Amplitude and wave number are taken to be strictly positive in this context, and although α_p may vary over a 360° interval, the geometry is indifferent to whether the reflecting surface is advancing or receding; indeed, for a simple corrugated reflector, α_p will flip by ±180° across different phases of each TID wave period.

At the reflection point, the phase of the TID is given by

$$\begin{aligned} \varphi_p &= k\left(x_p \cos \chi^* + y_p \sin \chi^* \right) - \omega t \\ &= kr_p \left(\cos \alpha_p \cos \chi^* + \sin \alpha_p \sin \chi^* \right) - \omega t. \end{aligned} \tag{8.8}$$

In this instance, to remain closer to Bramley's original notation, both α_p and χ^* are measured anti-clockwise relative to the midpoint path bearing; that is, in the opposite sense to conventional azimuth (as indicated by χ elsewhere in this chapter). If the midpoint bearing is denoted θ_m (°T), then the corresponding bearings for the perturbed reflection points and TID itself, measured clockwise from True North, but subject to a 180° directional ambiguity, are

$$\theta_p = \theta_m - \alpha_p, \text{ and} \tag{8.9}$$

$$\theta_{TID} = \theta_m - \chi^*. \tag{8.10}$$

Combining equations (8.6) and (8.7) gives the following expressions for the displacement r_p and axis α_p of the reflection point perturbations (cf. Bramley [1953], who uses the complement of α_p in his derivation):

$$r_p = \sqrt{x_p^2 + y_p^2} = kA |\sin \varphi_p| \sqrt{h_0^2 + \left(\frac{D^2}{2} + \frac{D^4}{16h_0^2} \right) \cos^2 \chi^*}, \tag{8.11}$$

and

$$\tan \alpha_p = \frac{y_p}{x_p} = \frac{\tan \chi^*}{1 + D^2/4h_0^2}. \tag{8.12}$$

For specular reflection, the TID phase φ_p must further satisfy equation (8.8) which, combined with equations (8.6) and (8.7), produces the transcendental equation [Bramley, 1953, eq. (5)]

$$\varphi_p + \omega t = k^2 A \left(h_0 + \frac{D^2}{4h_0} \cos^2 \chi^* \right) \sin \varphi_p, \tag{8.13}$$

whose solutions can be visualised as intersections between a straight line and a sinusoid.

At any time, there may be multiple solutions to equation (8.13); that is, supporting the condition for specular reflection. However, across each TID period, the maximum horizontal displacement (occurring when $|\sin \varphi_p| = 1$) will always be given by

$$r_{p,max} = kA \sqrt{h_0^2 + \left(\frac{D^2}{2} + \frac{D^4}{16h_0^2} \right) \cos^2 \chi^*}. \tag{8.14}$$

The square-root term in equation (8.14) governs the variation in maximum displacement for different TID bearings. For paths at or near vertical incidence (i.e. $D \ll h_0$), $r_{p,max} \approx kAh_0$ in all directions; that is, the response is isotropic. On the other hand, for longer oblique paths (i.e. $D \gg h_0$), $r_{p,max}$ varies from approximately $kA\frac{D^2}{4h_0}$ for TIDs propagating along the same bearing as the oblique path ($\chi^* \approx 0°$ or $180°$), to kAh_0 for TIDs propagating in a perpendicular direction ($\chi^* \approx \pm90°$).

This directional dependence ("bias") in Bramley's model was noted by Clarke [1972, p. 268], although it was interpreted in the wrong sense; in actual fact, the model predicts larger maximum displacements for the longitudinal TID compared to the transverse TID. However, the theoretical solutions contributing to this result are severely asymmetric paths, with reflection points displaced by many TID wavelengths and well outside the initial estimates used for synthesis in **Section 8.2.1**. Thus, the incidence angle to the ionosphere as a whole is likely to be so different that the virtual height of reflection at the midpoint is rendered meaningless; these paths may not even return to ground. Observational evidence suggests transverse TIDs remain more detectable in general, in part because of their stronger signatures in delay, Doppler and bearing.

To produce multiple simultaneous reflections, the TID amplitude and horizontal wave number must be such that (from equation (8.13))

$$k^2 A \left(h_0 + \frac{D^2}{4h_0} \cos^2 \chi^* \right) > 1. \tag{8.15}$$

In terms of the horizontal TID wavelength $\lambda_h \equiv 2\pi/k$, this requires

$$\lambda_h < 2\pi \sqrt{A \left(h_0 + \frac{D^2}{4h_0} \cos^2 \chi^* \right)}. \tag{8.16}$$

At vertical incidence, equation (8.16) simplifies to $\lambda_h < 2\pi\sqrt{Ah_0}$, while at oblique incidence the condition is again aspect dependent. For sufficiently long paths with $\chi^* \approx 0°$ or $180°$ (longitudinal TID), multiple reflections occur when $\lambda_h < \pi D\sqrt{A/h_0}$, whereas with $\chi^* \approx \pm90°$ (transverse TID), this becomes $\lambda_h < 2\pi\sqrt{Ah_0}$, the same as at

vertical incidence. The former case ($\chi^* \approx 0°$ or $180°$) supports a wider range of amplitude-to-wavelength ratios (given $D \gg h_0$), implying that multiple reflections should in theory be more common with longitudinal TIDs at oblique incidence. However, as noted above, this result is treated with a certain level of suspicion, owing to the extreme path asymmetries. As an aside, all the synthetic examples in **Section 8.2.1** easily satisfy the multi-reflection condition, except for the model in **Figure 8.6** which is at the critical threshold (i.e. both sides of equation (8.16) approximately equal).

Given a set of mapped reflection point data, the value of Bramley's model is that equations (8.12) and (8.14) together provide a direct analytic transformation to obtain the TID axis and amplitude-to-wavelength ratio, appearing here as kA. Furthermore, such a transformation can be computed for each reflection point independently, using equation (8.11) in place of (8.14) for a general TID phase φ_p, so multi-reflections need not be discarded. The so-called "instantaneous" amplitude-to-wavelength ratio kA^* (where $A^* \equiv A|\sin\varphi_p|$) is calculated in this way, using a single reflection point, yet it only has physical meaning when the maximum absolute value is taken over a full wave period. In theory, the instantaneous ratio should maximise to the true ratio kA given by equation (8.14). While not a complete parameter inversion, the Bramley ripple transformation nevertheless has the potential to be a fast and robust preliminary step in TID analysis.

It is important to note that all of the above assumes a flat Earth and an untilted mirror ionosphere, unlike the corrugated mirror model from **Section 8.2.1**. For longer paths in particular, this is a poor approximation. An incremental improvement, to accommodate a spherical Earth, may be attained by using the direct (straight line) transmitter-receiver separation for D, rather than the great circle ground range, and the height above this line for h_0, rather than the height above the ground. Tilts can potentially be accommodated by subtracting off the larger-scale, non-periodic variations in the reflection point latitude, longitude, height and bearing, prior to the above transform. However, as noted in **Section 7.4**, in practice it is difficult to separate this background from the TID variability of interest, which is typically evolving over time and quasi-periodic at best.

The next sub-section follows with an error analysis based on synthetic TID model data, to better understand the deficiencies of the Bramley ripple transformation, before subsequently applying the technique to actual ELOISE peak data.

8.3.1.1 Synthetic error analysis

Errors are introduced in the Bramley ripple transformation at both the reflection point mapping stage (i.e. under the assumption of a tilted mirror and no geomagnetic field) and during the estimation of TID bearing and amplitude-to-wavelength ratio (i.e. under the assumption of a single sinusoidal corrugation imposed on a flat mirror). In this sub-section, each of these contributions is analysed separately using the TID models from **Section 8.2** as the "truth", specifically focusing on the monochromatic examples from **Figure 8.6** and **Figure 8.13**. Again, ray tracing/homing algorithms are applied to produce synthetic observables for this purpose, but of course this is a considerable simplification of reality. Measurement errors in delay and AoA are also excluded.

Starting with the reflection point mapping stage, **Figure 8.19** shows errors in the estimated ground displacements (with respect to the great circle midpoint) for a Hooke TID model propagating both transversely (top row) and longitudinally (bottom row) to the oblique sounder path. The transverse case is the same as in **Figure 8.13**, while the longitudinal case has been rotated by 90°. The tilted mirror assumption does not discriminate between O and X modes generated by the 3D magneto-ionic ray tracing, so the errors are slightly different for each component; these are plotted in red and blue, respectively. Some rays penetrate deeper into the ionosphere, particularly around the troughs of the disturbances (e.g. see **Figure 8.16**), and these map to the largest displacement errors: 100 km or more on a 1507 km path. However, such points tend to be outliers, and the root-mean-squared error (RMSE) remains a more modest 7–15 km for the transverse case and 27–44 km for the longitudinal case. In comparison, the RMS values of the absolute displacements are substantially larger, at 26–28 km and 97–111 km for the transverse and longitudinal cases, respectively; in other words, the tilted mirror model (TMM) correctly accounts for roughly half to three-quarters of the displacement, which would be neglected altogether under a spherical mirror model (SMM).

Note that the error distributions are asymmetric and somewhat bimodal, due to multiple simultaneous reflections having different layer penetration depths and AoAs (as seen in **Figure 8.13**). The O-mode (red) for the longitudinal case (bottom row) is the most obvious example of this, with the mode of the histogram being positive (+12 km) and the mean being negative (-15 km). It is for this reason that the simple RMSE is presented here rather than other potentially misleading statistics.

The corresponding plot for the reflection point perturbation bearing (measured likewise with respect to the midpoint) is shown in **Figure 8.20**. For this variable, there is quite a substantial difference in the scale of errors between transverse and longitudinal TIDs, with a clear bias of around +20° in the O-mode time series under a transverse geometry. The reported RMSE statistics are somewhat misleading, as the bearing is poorly defined when the absolute displacement is near zero, resulting in occasional gross errors out to ±180°. These are not truly representative, and in practice, such weakly perturbed points are unlikely to contribute to any TID parameter inversion. Needless to say, in the special
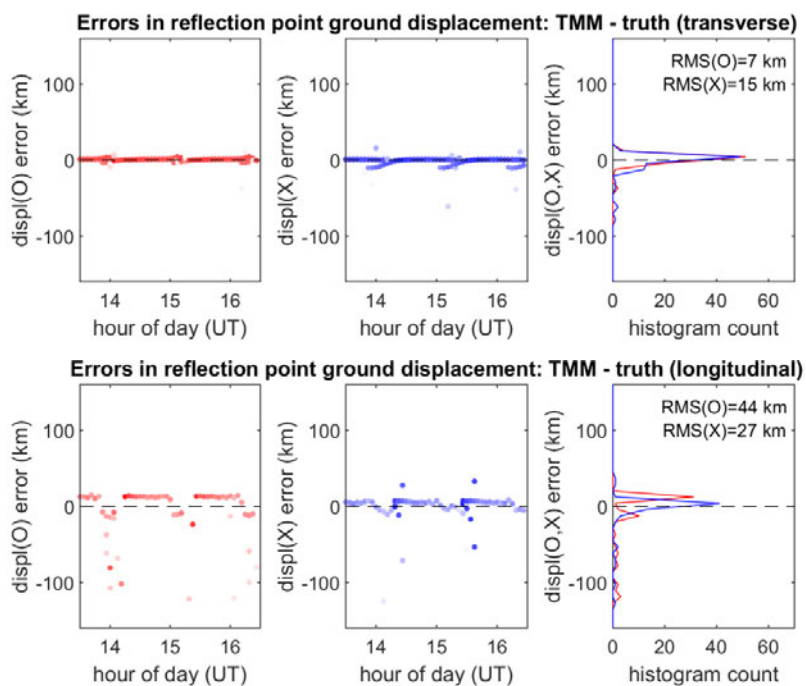
Figure 8.19. Errors in the reflection point ground displacements, estimated using group delay and AoA observables synthesised from two AGW-seeded (Hooke) TID models and mapped via a tilted mirror model (TMM). The transverse TID model in the top row of panels is the one previously depicted in **Figure 8.13**, for the Kalkarindji to Laverton path (1507 km) ray-traced at an operating frequency of 8 MHz, while the longitudinal model in the bottom row is the same, in terms of the neutral wind perturbation, but with the AGW/TID bearing rotated by 90°. As in **Figure 8.13**, the red and blue markers represent the O- and X-mode rays, respectively, which are each mapped separately with no knowledge of the geomagnetic field. The left and middle columns show O/X ray displacement errors as a function of time, where the "truth" is based on the position of the ray apogee, while the right column presents the same data as histograms, with RMSE statistics printed alongside each.
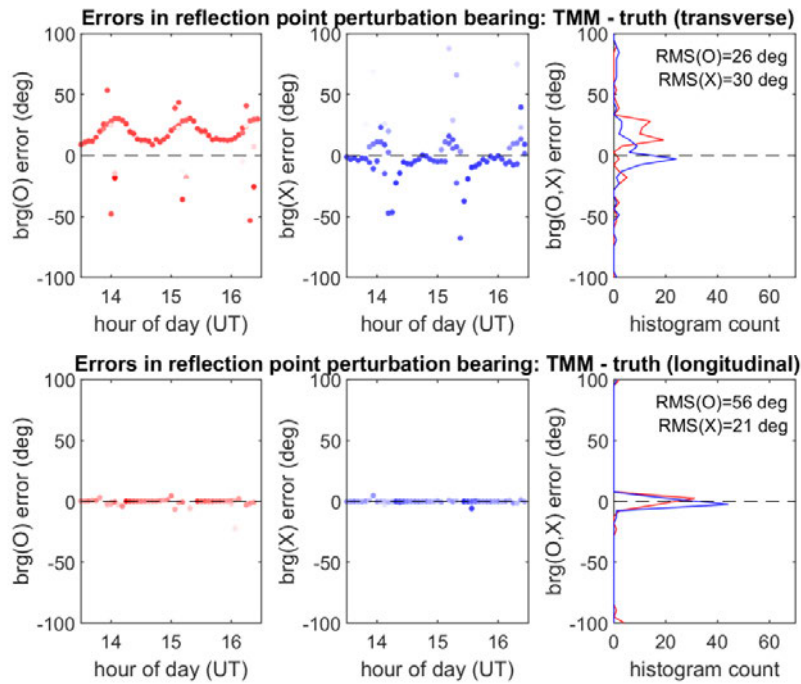
Figure 8.20. As in **Figure 8.19**, but for errors in the reflection point perturbation bearings.

no-field case where the ionosphere is completely unperturbed (spherically symmetric), all reflection points will inherently have zero displacement and an undefined bearing.

It is important to remember that this is just a single example, sampled discretely and containing a non-integer number of periods, so the RMSE values should be interpreted with some caution. Furthermore, the results will be strongly dependent on geomagnetic latitude, with alternative path alignments to the field leading to variations in the O/X differences. Dao et al. [2016] have demonstrated the importance of simulations in understanding the limitations of the tilted mirror model approximation in direction finding applications; for some paths, the use of O-mode observables may be more resilient to errors, while the X-mode (or even a combination of the two) may be a better choice for other paths. In the example shown here (Kalkarindji to Laverton), the X-mode tends to give smaller errors and less bias overall.

The second stage in which errors are introduced is the Bramley transformation itself, between ionospheric reflection points and TID parameters. While both sides of the transformation fall within a mirror reflection regime, the latter imposes additional

227

assumptions on the shape of the reflecting surface (i.e. flat with a single sinusoidal perturbation) and the scale of the disturbance (i.e. small amplitude-to-wavelength ratio). The resultant TID parameter estimates are evaluated in **Figure 8.21** using the corrugated spherical mirror model in **Figure 8.6** rotated through different angles with respect to the midpoint bearing. Note that although estimates are produced for every reflection point, it is the maximum absolute value of the instantaneous ratio, occurring at maximum displacement, that best defines the underlying disturbance, and it is at this point that the TID parameters (amplitude-to-wavelength ratio and bearing) are extracted. The true values of these parameters are indicated by the dashed horizontal lines on each panel.
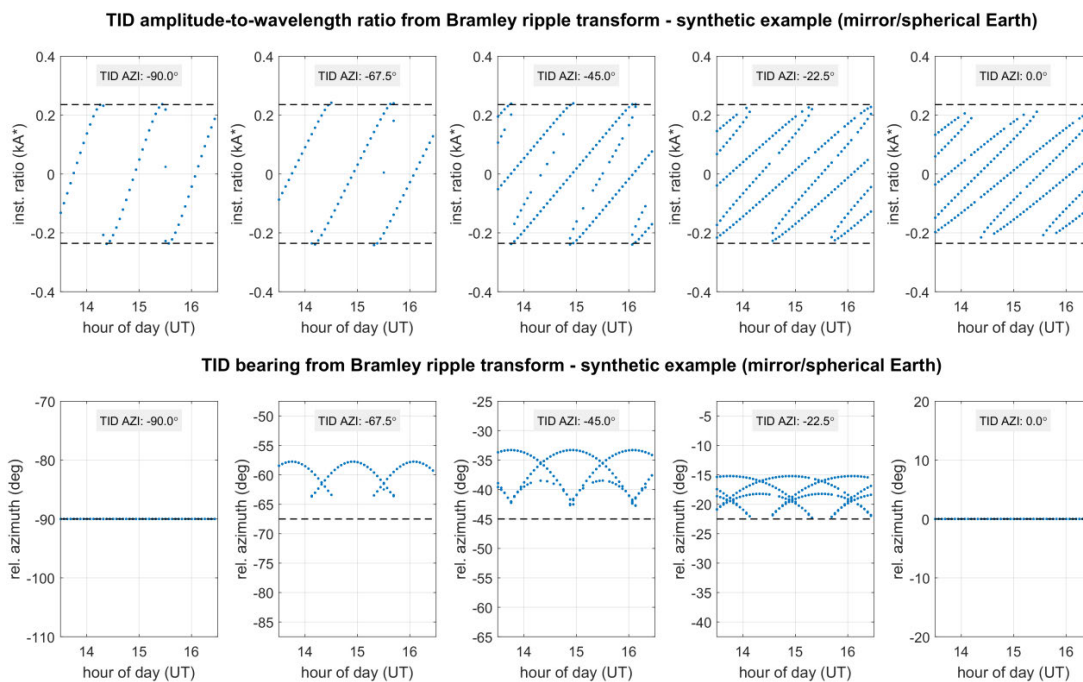
Figure 8.21. Time series of instantaneous amplitude-to-wavelength ratio (top row) and bearing (bottom row) estimates, derived via the Bramley ripple transform for TID observables synthesised with a corrugated mirror model on the Kalkarindji to Laverton path (1507 km). This is the same mirror model as in **Figure 8.6**, but with the TID direction rotated by various relative azimuth angles (clockwise with respect to the great circle midpoint bearing). The five columns represent different azimuths from -90° (transverse TID) to 0° (longitudinal TID), as labelled at the top of each panel, and to make things clearer, bearing estimates have been converted back to relative azimuth in the bottom row of panels. The true amplitude-to-wavelength ratio (kA = 0.24 rad) and relative azimuths are indicated by the dashed horizontal lines.

228

With an amplitude-to-wavelength ratio of 0.24 rad for the TID model in **Figure 8.21**, the small-angle approximation $\tan(kA) \approx kA$ is only in error by about 2%. As such, the transformed estimates of kA^* do indeed reach a maximum absolute value that is reasonably close to the truth in the top row of panels. The transformed estimates of TID bearing, scaled as relative azimuth angles in the bottom row of panels, are identical to the truth for azimuths of -90° (transverse) and 0° (longitudinal), but otherwise tend to be biased towards zero azimuth. While not shown here, this also applies for positive TID azimuths (i.e. for which the bias has the opposite sign) and is the result of curvature in the background (spherical) reflecting surface. However, if only the azimuth estimate at maximum displacement is used, the error is considerably reduced.

Figure 8.22 plots the TID parameter errors as a function of the relative azimuth angle. As in **Figure 8.21**, for the untilted case in blue, the amplitude-to-wavelength ratio error is largest in magnitude for azimuths near 0° (longitudinal TID), while the bearing (azimuth) error is largest for azimuths of ±60–70° (not accounting for the ±180° bearing ambiguity). Both errors are close to zero for azimuths of ±90° (transverse TID), which corresponds to the peak in instrumental detectability discussed earlier. Like the real ELOISE observations, the synthetic reflection points have been discretely sampled at 3.75 min, so there is a potential for amplitude underestimation due to a mismatch between the critical sample and the actual time of maximum displacement. Nevertheless, this is expected to be small compared to the other error sources identified.

A key factor neglected in the above error analysis is the influence of large-scale tilts, which act to displace the centre of the oscillating pattern in the reflection point positions. With even a modest tilt angle of 1°, the accuracy of the TID parameter inversion begins to degrade, as shown by the red (down-range) and yellow (cross-range) variants in **Figure 8.22**. The former has a larger impact on the bearing error, whereas the latter has a larger impact on the amplitude-to-wavelength ratio error, particularly for TID azimuths of ±90° where the untilted errors are essentially zero in both cases. In the ±60–70° azimuth range, the ratio error now reaches up to 10% of the true value, with the bearing error rising to as much as 8°. Any residual bias in the reflection point positions due to array calibration issues would have a similar influence on the parameter errors. For weaker TIDs, especially on longer paths and in the presence of stronger horizontal gradients, such issues may render the Bramley transformation unusable despite its appealing simplicity.
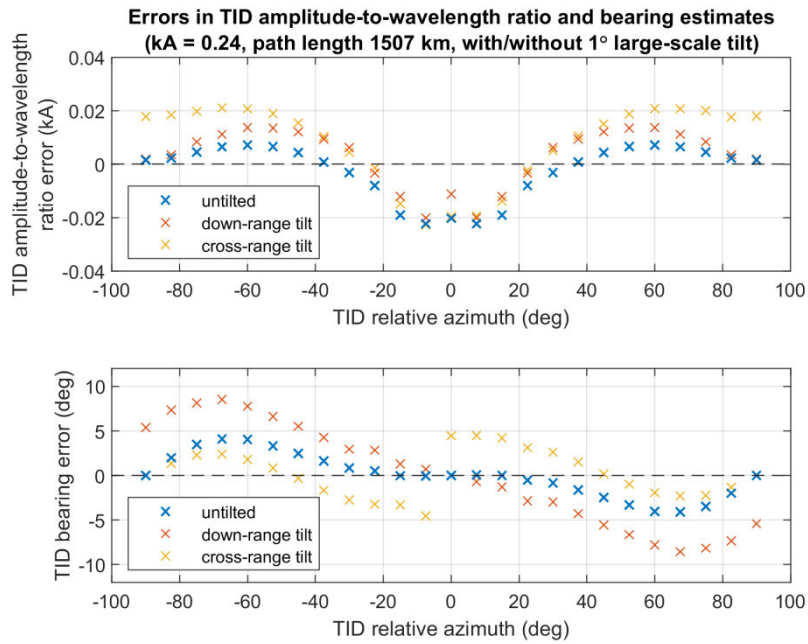
Figure 8.22. Errors in amplitude-to-wavelength ratio estimates (top panel) and bearing estimates (bottom panel) for the TID model in **Figure 8.21** (also evaluated at additional azimuths). At each azimuth, relative to the midpoint path bearing, the estimated TID parameters have been derived from the synthetic reflection point(s) with the maximum displacement over one or more wave periods. Along with the untilted case (blue markers), errors have also been calculated with a 1° large-scale tilt applied in both the down-range (red) and cross-range (yellow) directions about the midpoint. This is based on shifting the centre of the background spherical mirror, as described in **Section 8.2.1**.

8.3.1.2 Evaluation on ELOISE peak data

An 8-hour sample of mapped reflection point data from ELOISE AoA ionograms is shown in **Figure 8.23**. Instead of being expressed in terms of latitude and longitude, each point has been converted to a bearing and ground displacement from the great circle midpoint, corresponding to θ_p and r_p in the derivation above. Given that TIDs tend to induce a quasi-periodic perturbation that oscillates on either side of the midpoint, the plotted bearing has been wrapped over the interval between ±90 °T to minimise the apparent discontinuity, leaving the plotted displacement to take both positive and negative values (as per the colour bar). The fact that many of the points lie very close to the unperturbed path bearing (dashed lines) simply reflects the earlier

finding that the distribution is highly elongated along the path axis (cf. **Figure 4.18** in **Section 4.4**). The greater number of positive (blue) displacements above and negative (red) displacements below the dashed lines is indicative of a persistent large-scale tilt for much of this period, causing a systematic shift in the centre of the distribution to the north and/or east.

Figure 8.24 presents the output of the Bramley ripple transformation, with the location of the reflection point defining a TID bearing (θ_{TID}) and an instantaneous amplitude-to-wavelength ratio (kA^*) according to equations (8.12) and (8.11), respectively. Recall that the maximum absolute value of kA^* should be representative of the true ratio term kA. These results apply the spherical Earth correction described earlier, but are otherwise for a flat, untilted background ionosphere. Although a four-quadrant arctangent is used for solving equation (8.12), to preserve a consistent phase variation with time, the TID direction inferred from any single point is still inherently ambiguous by a factor of 180° (without resorting to the CSF-derived Doppler). As in **Figure 8.23**, to improve the presentation of **Figure 8.24**, the plotted TID bearing has been wrapped between ±90 °T, while allowing the instantaneous amplitude-to-wavelength ratio to take both positive and negative values.

Alternating blue/red horizontal bands in plots such as **Figure 8.24** are the conjectured appearance of simple (corrugated mirror) TIDs in isolation, however such features are rarely seen in practice. Large-scale horizontal gradients (tilts) act to shift the reflection points to one side of the path midpoint, so that the axis along which they lie no longer passes through this reference position and, hence, the Bramley geometry is broken. Instead, the TID will leave a signature that oscillates in bearing and may not incur a change of sign (colour) in the instantaneous amplitude-to-wavelength ratio as expected. Attempts at removing the background gradients (i.e. by subtracting off the larger-scale, slowly varying component of motion in the reflection points) were not overly successful, as there is not a clear delineator between the characteristics of tilts and TIDs.

The TID example from **Figure 8.7** can be seen at around 13–14 UT in **Figure 8.23** and **Figure 8.24**. The reflection points are by no means fixed in their bearing with respect to the midpoint, yet there is evidence of their displacement changing in sign (colour), and this banded signature becomes slightly clearer after the Bramley ripple transformation. At a virtual height of 300 km, the majority of points with a strong instantaneous

amplitude-to-wavelength ratio are clustered around a TID bearing of roughly -60 °T, which incidentally is much the same as was synthesised with the corrugated mirror model in **Figure 8.8** (i.e. -61° T). Although the transformed estimates do vary with
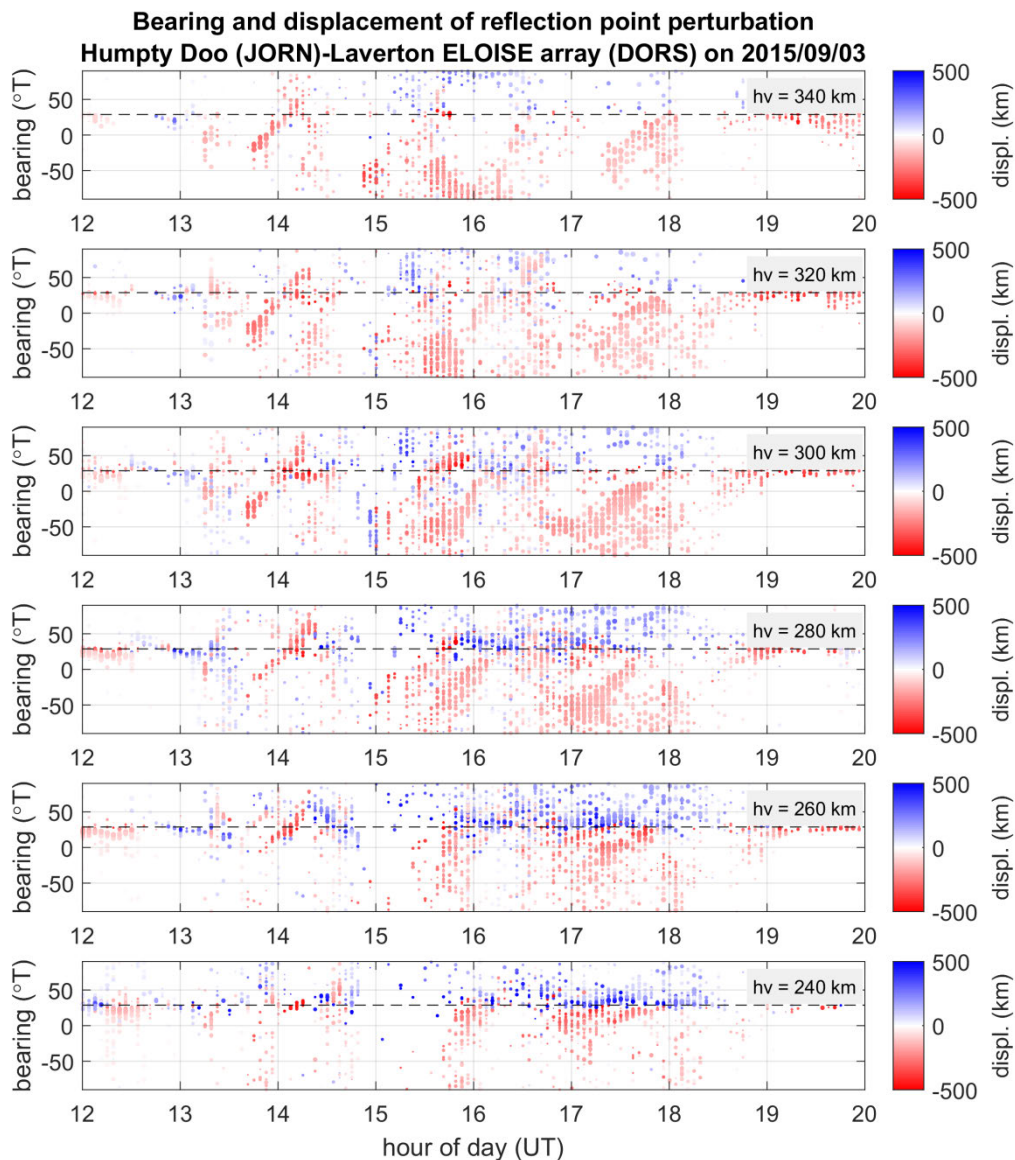
Figure 8.23. Time series of 1-hop F2-low reflection point perturbations, in terms of the wrapped bearing and ground displacement from the great circle midpoint, for six different virtual height windows spanning 240 to 340 km (bottom to top panel, as marked in the top-right corner). The ELOISE AoA path (Humpty Doo to Laverton) and day (3 September 2015) are the same as shown in **Figure 8.7**. The dashed horizontal lines on each panel show the unperturbed path bearing at the midpoint (29 °T), which differs slightly from the bearing at the receive site (31 °T). Each panel includes a ±5 km virtual height window to capture a sufficient quantity of ionogram peaks.
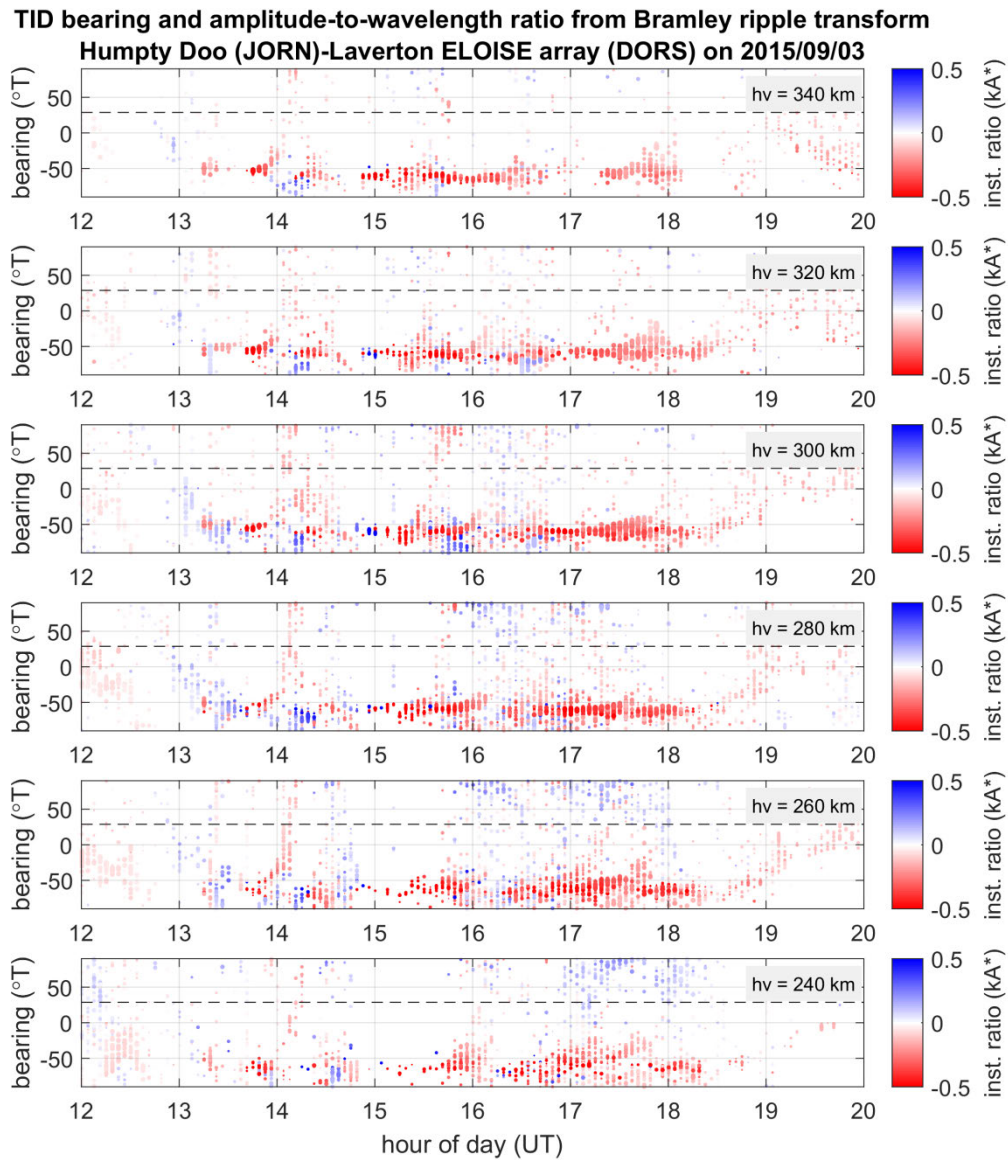
Figure 8.24. Time series of transformed TID bearing and instantaneous amplitude-to-wavelength ratio estimates (in radians), for the 1-hop F2-low reflection point data in **Figure 8.23**. Again, each panel represents a different ±5 km virtual height window, from 240 to 340 km, and the dashed horizontal lines show the unperturbed path bearing at the midpoint.

height, it is somewhat disappointing to see no obvious pattern of characteristic lags between the six panels in **Figure 8.24**, which casts serious doubt over the ability of this technique to be used for truly meaningful TID identification. **Section 8.4** will later return to evaluating this technique in the context of multiple paths.

8.3.2 Frequency and Angular Sounding technique

The FAS technique is a fairly mature method for automatically calculating TID parameters from AoA and Doppler ionosonde measurements, which has become notable recently in its Lowell Digisonde DPS-4D implementation as part of the Net-TIDE TID alert system [Reinisch et al., 2018]. Although originally developed for oblique reception of HF emitters by the Ukrainian UTR-2 radio telescope [Beley et al., 1995], it was adapted for the Digisonde in oblique mode by Galushko et al. [2003] and later generalised for both vertical and oblique incidence geometries by Paznukhov et al. [2012] (see also Paznukhov [2004]). A recent application of the technique to high-fidelity AoA and Doppler measurements from the new Long Wavelength Array (LWA) radio telescope at Sevilleta, New Mexico, has further demonstrated its utility for near vertical incidence geometries [Obenberger et al., 2019]. In addition, comparisons against simultaneous measurements from the Millstone Hill incoherent scatter radar have shown a reasonable level of agreement [Galushko et al., 2003; Galushko et al., 2008].

Like the simpler Bramley ripple transformation, the FAS technique operates on a single path and assumes reflection from a corrugated mirror surface. It does, however, allow for a more complex disturbance field composed of multiple periodicities, provided each is separable by spectral analysis and corresponds to a single plane wave only; this is a key assumption under the standard deterministic ("dynamic") implementation. An extension to support arbitrary background electron density distributions is possible, although this adds considerable extra complexity [Galushko et al., 2008].

The FAS technique starts with the following expression for the height of the reflecting surface, based on a flat-Earth system with horizontal coordinates x and y, as sketched in **Figure 8.25** (cf. Figure 1 of Paznukhov et al. [2012]):

$$h(x, y, t) = h_0(1 + \gamma(x, y, t)). \tag{8.17}$$

The scale of the height perturbation term $\delta h = h_0 \gamma$ is assumed to be small compared to the unperturbed height h_0.
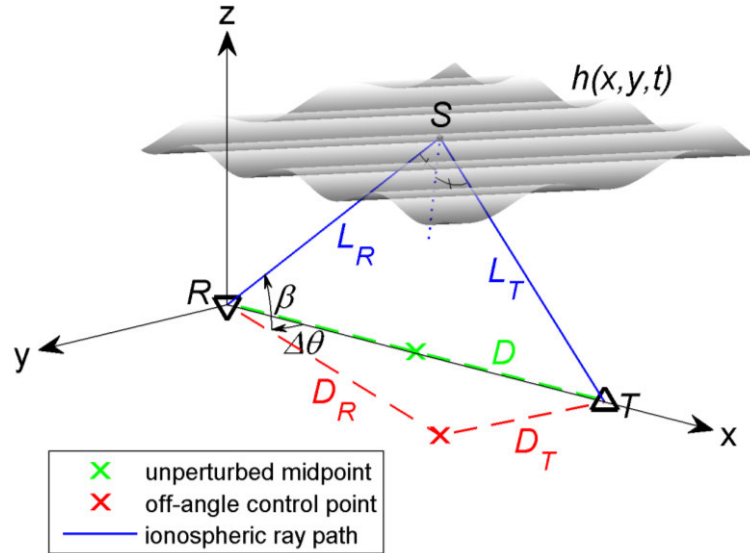
Figure 8.25. Diagram of receiver-centred Cartesian coordinate system used for the corrugated mirror TID model and FAS inversion (not to scale). For consistency with the coordinates found in Paznukhov et al. [2012], the same left-handed axes convention has been adopted. A perturbed (off-angle) path from transmitter "T" to receiver "R" via the ionospheric reflection point "S" on the sinusoidal surface $h(x, y, t)$ is shown in blue. This has the effect of shifting the control point on the ground from the green cross (unperturbed midpoint) to the red cross. Although the flat-Earth case is illustrated here, the geometry was extended to a spherical Earth for the ELOISE oblique paths by contracting the transmitter-receiver separation D and transforming the azimuth angle $\Delta\theta$ and elevation angle β with respect to the local horizontal plane at the receiver. The spherical-Earth case is thus related to the coordinate system used for reflection point mapping (**Figure 4.17**) by a simple rotation and x-y axis exchange.

The total geometric path length, as a function of the reflection height, is then given by

$$L(h) = L_R + L_T = \sqrt{D_R^2 + h^2} + \sqrt{D_T^2 + h^2}, \qquad (8.18)$$

where D_R and D_T denote the horizontal displacement of the reflection point from the receiver and transmitter, respectively. In the unperturbed case, $D_R = D_T = D/2$, where D is the transmitter-receiver baseline, and the path length simplifies to

$$L_0 \equiv L(h_0) = 2\sqrt{(D/2)^2 + h^2}. \tag{8.19}$$

The TID power spectrum, which defines the amplitude of the height perturbation for each spectral component, is derived by linearising the formula for the geometric path length about the midpoint. That is,

$$L(h) \approx L_0 + \delta h \left.\frac{dL}{dh}\right|_{h=h_0} = L_0 + \frac{4h_0^2}{L_0}\gamma. \tag{8.20}$$

Again, this assumes that the perturbation is small and, therefore, the changes in L are due to midpoint height variations only. For mirror reflection, the measured excess delay is simply given by

$$\tau = \frac{L - L_0}{c} \approx \frac{4h_0^2}{cL_0}\gamma. \tag{8.21}$$

Hence, the complex frequency spectrum of the surface perturbation at the midpoint can be estimated as [Paznukhov et al., 2012, eq. (16a)]

$$\tilde{\gamma}(\Omega) = \frac{cL_0}{4h_0^2}S_\tau(\Omega), \tag{8.22}$$

where $S_\tau(\Omega) = \int_{-\infty}^{\infty} dt\, e^{-i\Omega t}\tau(t)$ is the Fourier transform of the excess delay function, in terms of angular frequency Ω. In practice, a fast Fourier transform (FFT) is used, with a Hann taper applied to the sampled data presented here, and the power spectrum is computed as $|h_0\tilde{\gamma}(\Omega)|^2$. An equivalent expression for $\tilde{\gamma}(\Omega)$ in terms of the Fourier transform of Doppler data also exists [Beley et al., 1995], which may be useful when delay measurements are not available or are less precise.

The corresponding TID azimuth and wave number spectra are derived by invoking the condition for specular reflection [Paznukhov et al., 2012, eq. (2)], subject to the assumption of a small-amplitude and smoothly varying perturbation field; that is,

$|\gamma|, \left|\frac{\partial \gamma}{\partial x}\right|, \left|\frac{\partial \gamma}{\partial y}\right| \ll 1$. This gives the following ray trajectory functions, which relate to the surface slopes in the two horizontal directions [Paznukhov et al., 2012, eq. (5)]:

$$\alpha_s \equiv \frac{L \cos \Delta\theta \cos \beta - D}{\sqrt{D^2 + L^2 - 2DL \cos \Delta\theta \cos \beta}} \approx -h_0 \frac{\partial \gamma}{\partial x}, \qquad (8.23)$$

and

$$\beta_s \equiv \frac{L \sin \Delta\theta \cos \beta}{\sqrt{D^2 + L^2 - 2DL \cos \Delta\theta \cos \beta}} \approx -h_0 \frac{\partial \gamma}{\partial y}, \qquad (8.24)$$

where x is the down-range (transmitter-receiver) axis, $\Delta\theta$ is the received azimuth angle (or bearing offset, measured clockwise from the x-axis for a left-handed coordinate system), and β is the received elevation angle (not to be confused with β_s).

Under the further assumption that each temporal TID frequency Ω corresponds to a unique spatial frequency (horizontal wave number) $k_\Omega = k(\Omega)$ and direction (midpoint azimuth) $\chi_\Omega = \chi(\Omega)$, the surface perturbation field can be expressed in spectral form as

$$\gamma(x, y, t) = \int_{-\infty}^{\infty} d\Omega \, \tilde{\gamma}(\Omega) e^{i\Omega t - ik_\Omega(x \cos \chi_\Omega + y \sin \chi_\Omega)}, \qquad (8.25)$$

representing a continuous superposition of plane waves. Taking the partial derivatives of equation (8.25) with respect to each horizontal coordinate x and y, then applying a Fourier transform, yields

$$S_{\alpha_s}(\Omega) = ih_0 k_\Omega \cos \chi_\Omega \, \tilde{\gamma}(\Omega), \qquad (8.26)$$

and

$$S_{\beta_s}(\Omega) = ih_0 k_\Omega \sin \chi_\Omega \, \tilde{\gamma}(\Omega), \qquad (8.27)$$

where $S_{\alpha_s}(\Omega)$ and $S_{\beta_s}(\Omega)$ are the Fourier transforms of the ray trajectory functions in equations (8.23) and (8.24), respectively.

The inversion problem for the remaining TID parameters k and χ is thus solved by combining equations (8.22), (8.26) and (8.27) to give [Paznuknov et al., 2012, eq. (16b) and (16c)]

$$k(\Omega) = \frac{4h_0}{cL_0}\sqrt{\frac{\left|S_{\alpha_s}(\Omega)\right|^2 + \left|S_{\beta_s}(\Omega)\right|^2}{\left|S_\tau(\Omega)\right|^2}}, \tag{8.28}$$

and

$$\tan\chi(\Omega) = \frac{S_{\beta_s}(\Omega)}{S_{\alpha_s}(\Omega)}. \tag{8.29}$$

Together the frequency and angular spectra defined by equations (8.22), (8.28) and (8.29) encapsulate the generalised FAS technique. The AGW/TID dispersion relation can be used in conjunction with these to give the vertical wave number if required (e.g. **Appendix E**, eq. (E.3)). While some minor subtleties have been omitted for sake of brevity, such as normalising the TID power spectrum for taper losses and resolving the ±180° ambiguity in TID azimuth (noting that the four-quadrant arctangent function requires strictly real inputs), the technique nevertheless remains attractively simple and quick to compute. The primary challenges in applying this to actual AoA observations are effectively estimating and removing the unperturbed (mean) delay L_0, and selecting a suitable period of data for spectral analysis, over which the TID parameters are static. It is also important to note that multiple simultaneous reflections are not supported, so data must be reduced to one reflected path per time sample before computing the FAS spectra; for example, by selecting the shortest path (minimum delay) only.

Just as in the Bramley ripple transformation, a first-order correction to accommodate a spherical Earth, but still a flat background ionosphere, involves substituting the direct transmitter-receiver separation for D, instead of the great circle ground range, and extending the unperturbed height h_0 to be the distance from this direct line, instead of the height above the ground. All results presented herein include this correction when dealing with synthetic observables from a spherical model.

To illustrate the output of the FAS technique, **Figure 8.26** shows the TID power spectrum, azimuth spectrum, and wave number spectrum (in terms of horizontal wavelength) for a simple monochromatic corrugated mirror model propagating transversely to the oblique path. This is a rather contrived case, with a very small perturbation amplitude ($\delta h = 0.1$ km) and a flat Earth/ionosphere geometry to best satisfy the underlying assumptions; in practice, it would be almost impossible to detect this TID above the measurement noise. An oblique ground range of 1989 km was chosen, representing Humpty Doo to Laverton, and a total of eight hours of synthetic observables, sampled every 3.75 min, were used in constructing the spectra. The likelihood of observing a single "frozen-in" TID wave train for this length of time is again highly unlikely.

Such a small-amplitude TID yields observables that are extremely close to sinusoidal, so it is not surprising to see the harmonics many tens of dB below the fundamental period of 30 min in the top panel of **Figure 8.26**. In theory, this leaves significant scope for separating out other TID components, were they to be present. The true sinusoidal amplitude of 0.1 km translates to an RMS power of 0.005 km^2 or -23 dBkm^2 (marked with a black cross), which is in reasonably good agreement with the peak spectral estimate at -25 dBkm^2; the loss of ~1.8 dB, after accounting for the taper, is due to spectral leakage into neighbouring frequency bins. In the second and third panels, the wavelength and azimuth estimates likewise reproduce their true values of 200 km and -90° with negligible errors.

In **Figure 8.27** and **Figure 8.28**, two of the key assumptions of the FAS technique are tested in terms of their effect on the synthetic TID parameter estimates. Firstly, it is seen in **Figure 8.27** that increasing the perturbation amplitude of the TID model used in **Figure 8.26**, from 0.1–25.6 km in logarithmic steps, starts to cause considerable degradation in power and horizontal wavelength estimates above about 2 km (3 dBkm^2). With an integer number of samples per period, the ability to estimate period and azimuth is not compromised for the transverse TID model, although this is not true in general. At greater amplitudes, the observables and ray trajectory functions are no longer pure sinusoids (cf. **Figure 8.10**), which leads to increased spectral leakage, beyond the small-amplitude limit noted above, and a substantial underestimation of power (amplitude), by as much as 13–15 dB (or ~80% of the true amplitude) at the
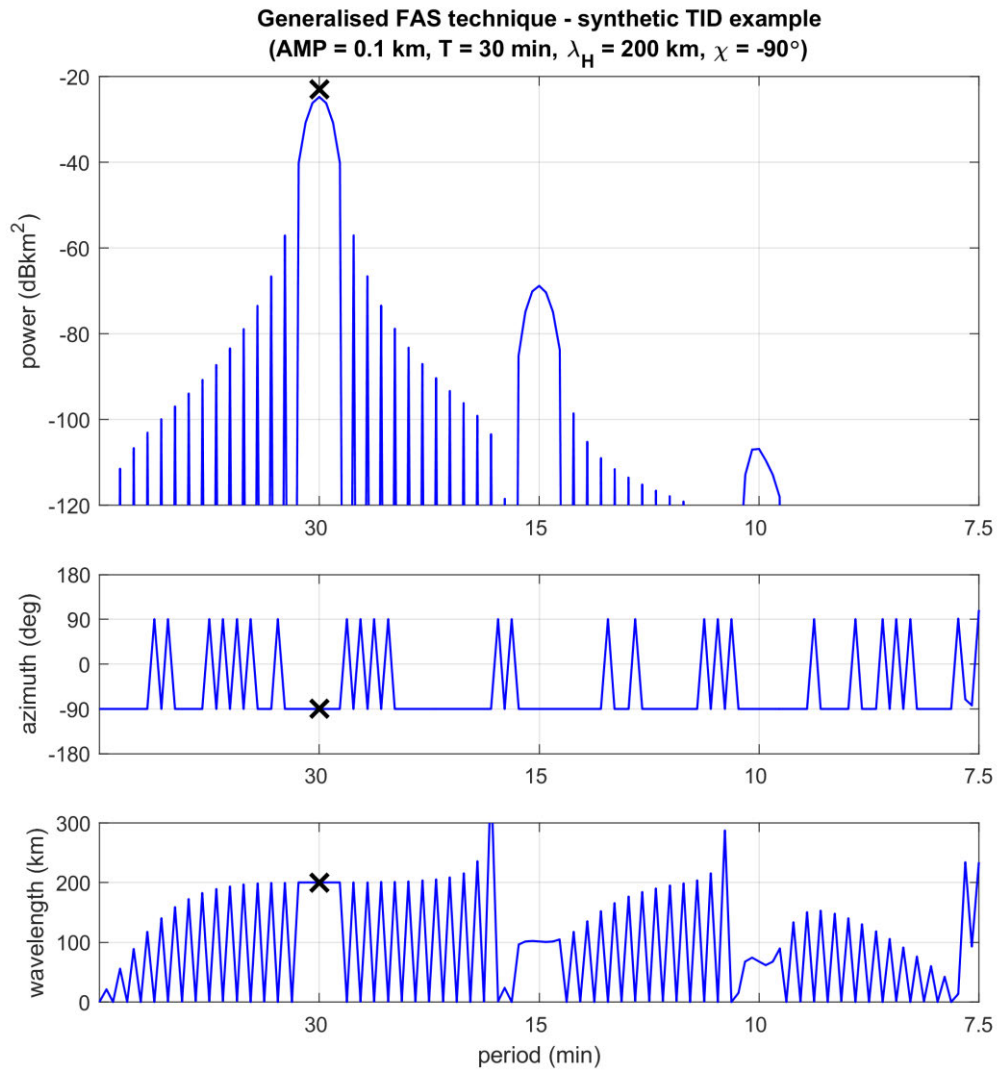
Figure 8.26. Generalised FAS spectra for a monochromatic corrugated mirror model with background height 250 km, perturbation amplitude 0.1 km, period 30 min, horizontal wavelength 200 km, and azimuth -90° (measured clockwise from the midpoint path bearing). Both the synthetic observables and parameter inversion have assumed a flat Earth/ionosphere in this case, so errors are minimal. The oblique ground range was taken to be 1989 km (Humpty Doo to Laverton). The panels, from top to bottom, show the TID power spectrum, azimuth spectrum, and horizontal wavelength spectrum, each with the frequency axis expressed in terms of wave period. The 128 data points were zero padded by a factor of two prior to calculating each FFT in order to more finely sample the discrete Fourier transform. A black cross on each spectrum marks the true TID model parameter.
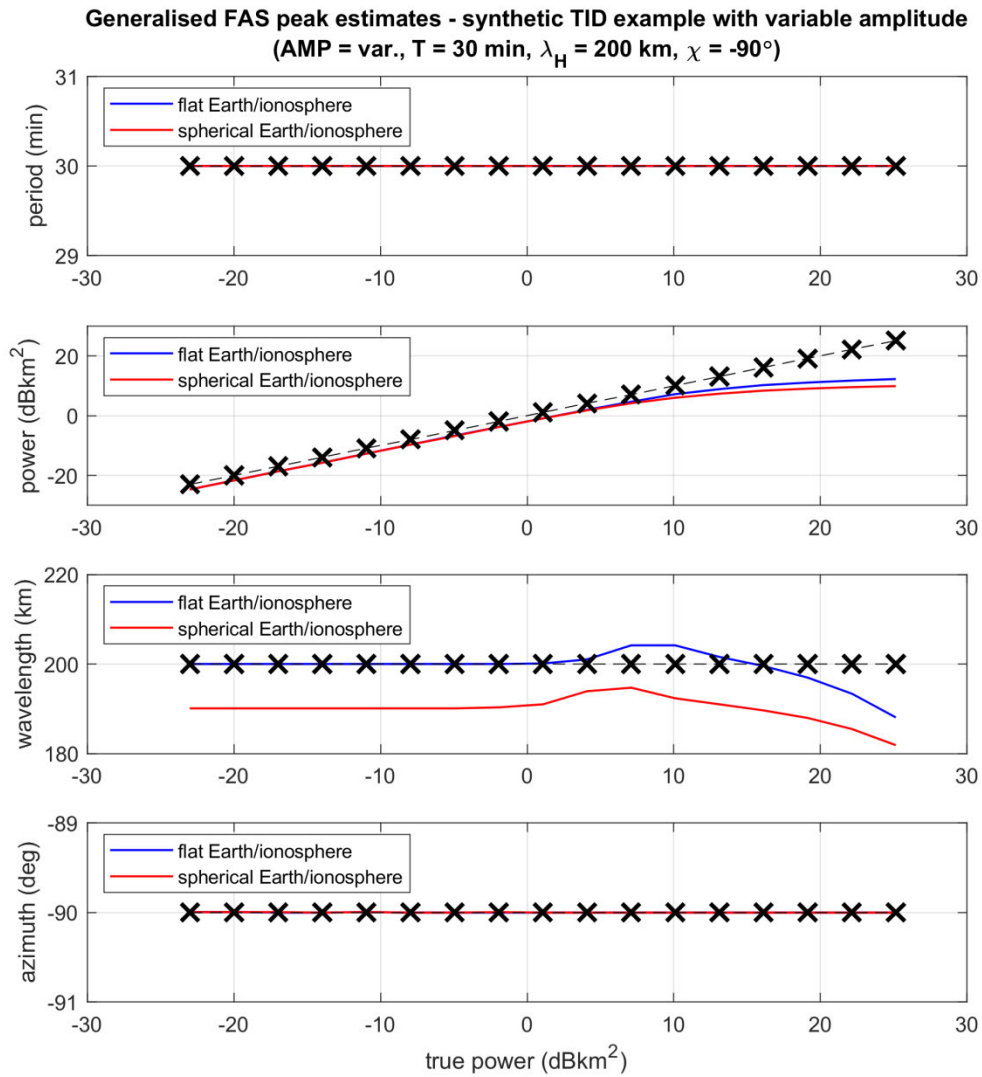
Figure 8.27. Inverted peak parameter estimates as a function of TID amplitude, determined by applying the generalised FAS technique to synthetic observables from a monochromatic TID model. This is based on the example analysed in **Figure 8.26**, with the same 1989 km oblique ground range. The panels, from top to bottom, show the TID period, power, horizontal wavelength and azimuth at the peak in the FAS power spectrum. True model parameters are indicated by the black crosses, while FAS estimates for flat and spherical Earth/ionosphere geometries are distinguished by the blue and red lines, respectively.
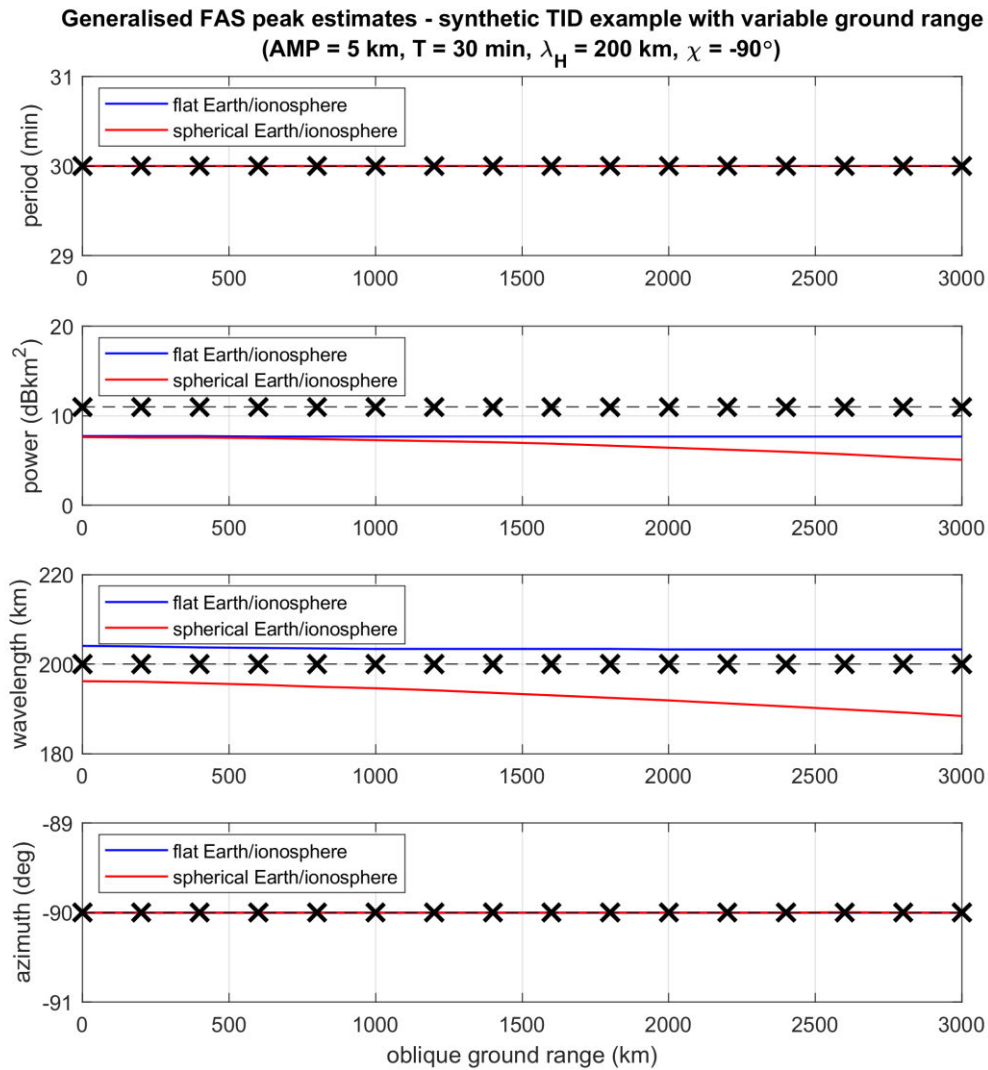
Figure 8.28. As in **Figure 8.27**, but for a variable oblique ground range, and with TID amplitude fixed to 5 km.

upper end shown. These errors, particularly in the wavelength estimates, are only made worse upon introducing a more realistic spherical Earth/ionosphere (red line). A similar degradation in power estimates occurs as the horizontal wavelength is reduced (i.e. increasing the amplitude-to-wavelength ratio), whereas the wavelength estimates themselves are prone to greater errors for larger-scale waves once spherical geometry is included in the TID model.

Secondly, **Figure 8.28** evaluates the dependence on oblique ground range (0–3000 km). With a fixed model amplitude of 5 km, likely near the lower limit of detectability, the

FAS estimates of power and wavelength are already in error at vertical incidence (zero ground range), and the difference progressively widens on longer paths for a spherical Earth/ionosphere; a flat geometry is merely fanciful at this point. At 3000 km, in the vicinity of the longest ELOISE AoA paths, power is underestimated by almost 6 dB (or ~50% of the true amplitude). The contrast between short and long paths becomes even more pronounced for larger TID amplitudes.

It is worth noting that while other authors have indeed applied the FAS technique to oblique observations, the published examples have path lengths ranging from 440 km [Galushko et al., 2007] and 450 km [Galushko et al., 2003; Paznukhov et al., 2012] to 700 km [Beley et al., 1995] and 500–1100 km [Reinisch et al., 2018]; that is, less than the typical ELOISE AoA path lengths. Furthermore, analysis of synthetic data has previously been restricted to fairly low-amplitude disturbances [Galushko et al., 2003; Galushko et al., 2008; Huang et al., 2016]; the latter paper in particular only considers model amplitudes below ~2 km, which are simply too small to be seen by the ELOISE AoA system given the measurement uncertainties in **Table 4.1** of **Section 4.4** (including median standard deviations of 1.2–1.6 km in group delay alone).

A further variant on these synthetic plots is shown in **Figure 8.29**, but this time with two TID components: the primary being the same as in **Figure 8.28** (i.e. amplitude 5 km, period 30 min, horizontal wavelength 200 km, and azimuth -90°), and the secondary being 3 dB weaker (i.e. amplitude 3.54 km), with variable period, wavelength 300 km, and azimuth -45°. In this case, the primary TID parameter estimates (in blue/red) continue to be reasonably good, besides the usual underestimation of amplitude, yet the secondary TID is poorly characterised. With spectral peaks extracted automatically, difficulties are naturally encountered when the secondary period is close to the primary period of 30 min. Results are also degraded for the secondary period at 15 min, being an exact harmonic of the primary period. Outside this unresolved region, the secondary azimuth estimates are generally quite acceptable, but the wavelength estimates are not; this remains the most challenging TID parameter to estimate reliably.
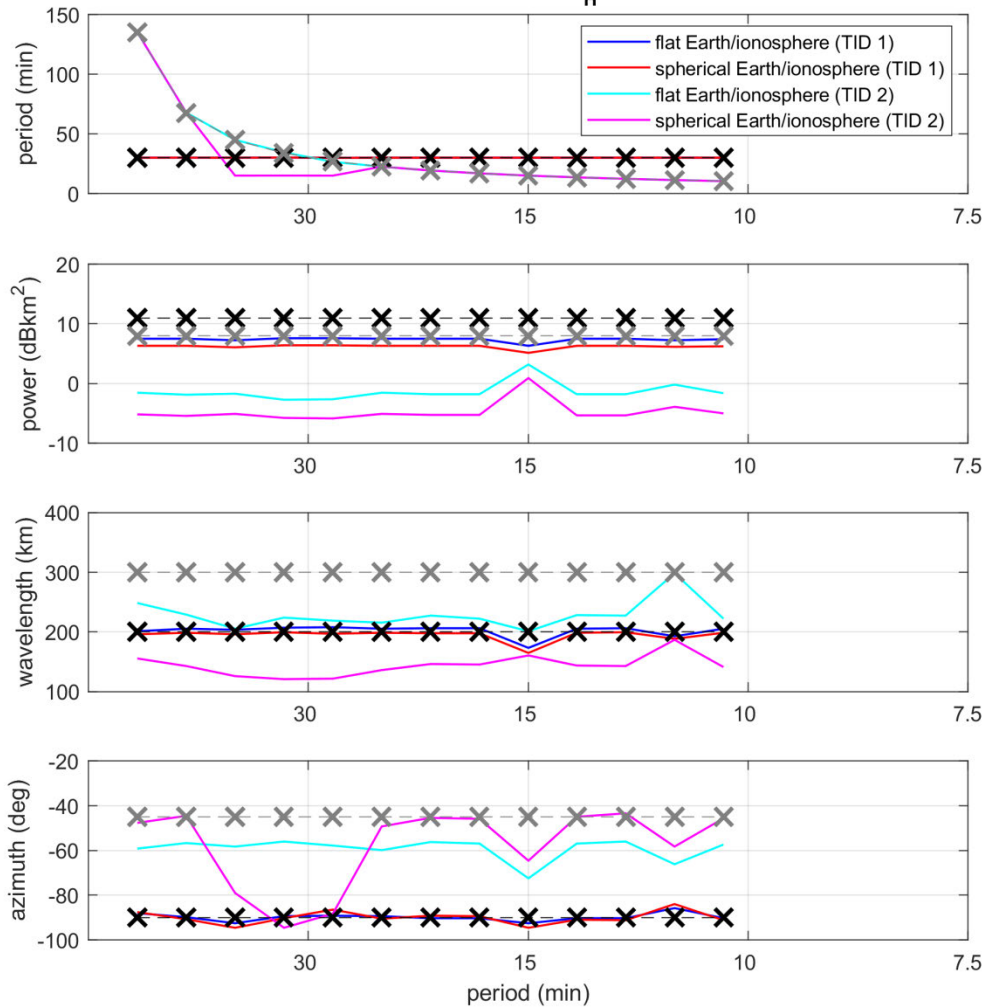
Figure 8.29. As in **Figure 8.27**, but adding in a secondary TID component with a variable frequency (period). The amplitude of the primary component ("TID 1") is fixed to 5 km, while the secondary component ("TID 2") is fixed to 3.54 km (or 3 dB lower in power). The secondary component also has a different horizontal wavelength (300 km) and azimuth (-45°). True model parameters for TID 1/TID 2 are indicated by the black/ grey crosses, while FAS estimates are shown in blue/cyan (for flat Earth/ionosphere) and red/magenta (for spherical Earth/ionosphere). Note that as the spectral peaks are extracted automatically, without knowledge of the true parameters, an incorrect solution is sometimes found for the secondary component, particularly when close to the primary period of 30 min.

Returning now to the representative ELOISE example in **Figure 8.6**, it is readily apparent that this TID model falls into the large-amplitude regime from **Figure 8.27**; that is, the observational response is in no way sinusoidal. It is expected that the FAS technique and its use of Fourier analysis will be compromised in such a case. To compare performance against the Bramley ripple transformation, TID parameters as a function of the relative azimuth angle were once again estimated for this model, using just synthetic delay and AoA observables. The results, shown in **Figure 8.30**, derive from an extended time series of 6 hours (96 samples), rather than 3 hours, to allow for the Hann taper. Unlike earlier FAS results, the primary period of 70 min is no longer an integer multiple of the sampling interval (3.75 min), so substantial zero padding has been applied to effectively interpolate the peak. While this addresses scalloping loss (i.e. the apparent spectral loss due to misalignment of the FFT bin from the centre of the peak), no attempt has been made to recover the loss due to spectral leakage more generally.

Figure 8.30 confirms that transverse TID parameters (i.e. azimuths of ±90°) are inverted most effectively by the FAS technique. Power is severely underestimated in other cases, by as much as ~17 dB (or ~86% of the true amplitude) when close to the longitudinal configuration, rendering the results largely useless. The horizontal wavelength and azimuth estimates are also substantially poorer away from the transverse geometry, although the latter is difficult to see on this scale. As such, the earlier transverse examples in this sub-section must strictly be viewed as best-case scenarios for FAS inversion.

Recasting these results in terms of parameter errors gives **Figure 8.31**, which has been presented in the same format as **Figure 8.22** (for the Bramley transformation). Power and wavelength have been combined as the amplitude-to-wavelength ratio kA (in radians) in the top panel, and azimuth error is equivalent to the bearing error in the bottom panel. Unlike the Bramley technique, knowledge of the time evolution resolves any 180° directional ambiguity in the FAS estimates, so the modelled azimuths span a full 360°. For a single TID component on an untilted spherical mirror (represented by the blue markers in both figures), the errors are larger for the FAS technique, by up to an order of magnitude for the amplitude-to-wavelength ratio and a factor of two for the bearing. This is believed to be a combination of amplitude loss due to spectral leakage, into both main-lobe and harmonics, and the fact that only the reflecting paths with

minimum delay are retained (whereas the Bramley transformation uses all reflection points).

Evidently there are a number of shortcomings imposed by the reflection geometry and conventional spectral analysis alone that make the FAS technique rather problematic for the ELOISE AoA paths and their typical measurement accuracy. This is without even considering factors such as large-scale tilts, realistic electron density profiles, geomagnetic (O/X) effects, and time variations in the background ionosphere and TID
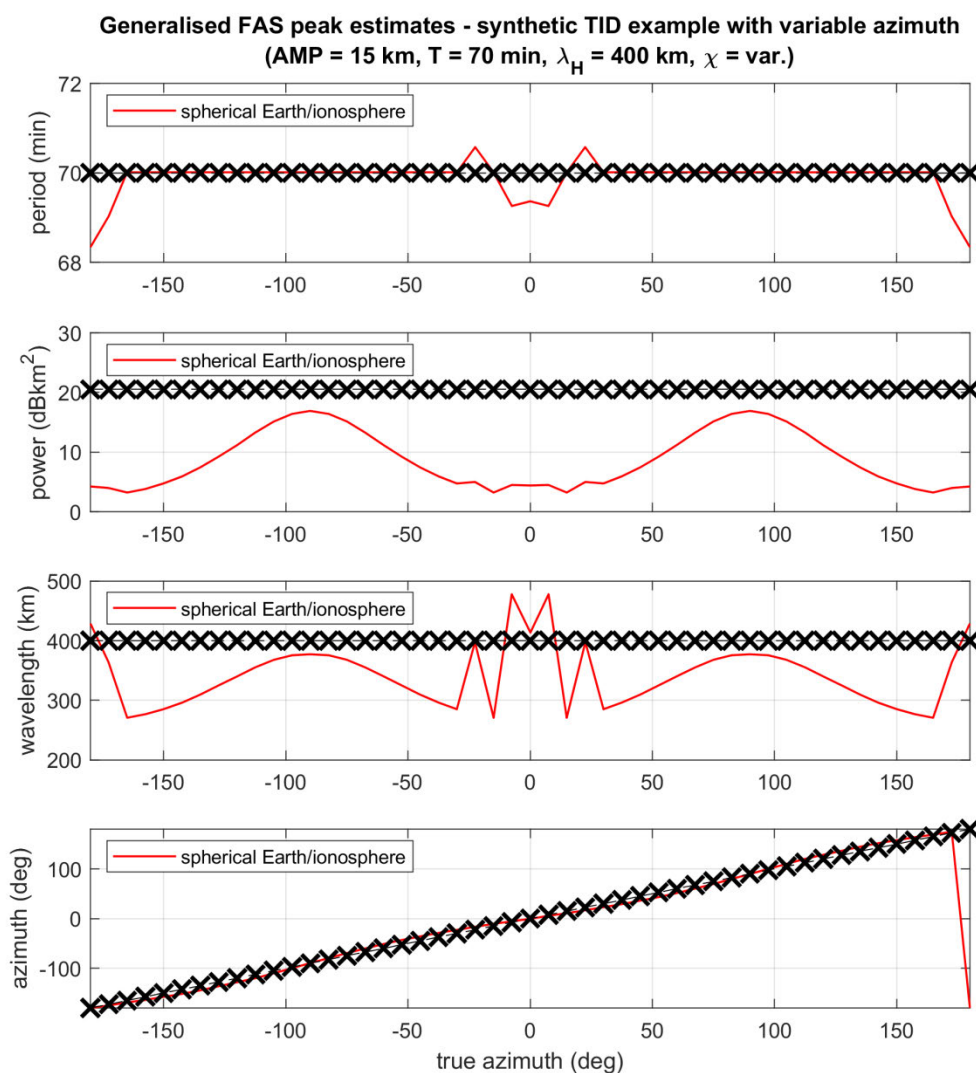
Figure 8.30. Inverted peak parameter estimates as a function of TID azimuth, based on the monochromatic TID model from **Figure 8.6**, with the same 1507 km oblique ground range. True model parameters are indicated by the black crosses, while FAS estimates for the spherical Earth/ionosphere geometry are plotted in red.

246
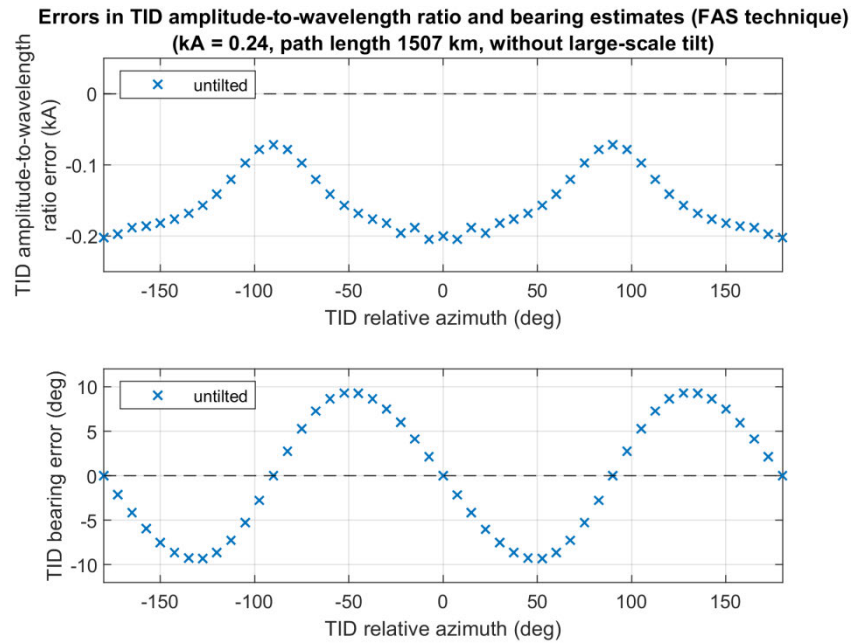
Figure 8.31. Errors in FAS-derived amplitude-to-wavelength ratio estimates (top panel) and bearing estimates (bottom panel) for the TID model in **Figure 8.6**, rotated across different azimuths (cf. **Figure 8.22** for the Bramley ripple transformation). Only the shortest ray path is used at each 3.75 min epoch, and the background spherical mirror is untilted.

components. The likelihood of observing 6–8 hours of a stable periodic component, particularly one with an amplitude of less than ~2 km, is almost non-existent in practice. Furthermore, actual HF observations will always contain measurement errors and uncertainty in the background ionospheric reference that act to the detriment of any TID parameter inversion.

In light of these synthetic results, and limited (failed) attempts to apply the algorithm to ELOISE AoA observations, the FAS technique was not ultimately pursued as a solution for the TID characterisation problem. Under certain conditions, it may remain suitable for shorter (quasi-vertical incidence) paths, where the geometry is more favourable against the FAS assumptions and AoA precision is generally better. Otherwise, it appears to be unsuitable for the class of TIDs detected in the ELOISE data set.

8.4 Relating TID signatures across multiple paths

The sections above have addressed the problem of extracting TID parameters from ELOISE AoA observables on a single path, which serves to remove the dependence on the oblique geometry and, thus, is an important precursor to enabling more direct comparisons across the ELOISE AoA network as a whole. However, the challenge of actually relating TID signatures and identifying spatial lags across multiple disparate paths remains outstanding. Recall that an earlier attempt, in **Section 7.5**, at using Doppler-based drift analysis to characterise the bulk motion of the ionosphere was inconclusive. Possible alternative approaches for TID analysis are: (1) the Bramley ripple transformation, (2) cross-spectral analysis of virtual height data, and (3) empirical mode decomposition. All three techniques use ionogram peak data only and do not require Doppler information. In this section, approaches (1) and (2) are explored.

As most of the ELOISE AoA midpoints are separated by more than a single TID wavelength, this is naturally a sparse spatial sampling problem, and one for which the reflection points are moving around on the same scale as the disturbance itself (routinely >100 km). Given the range of TID periods capable of reaching F2 heights, the observable response is also potentially a multi-harmonic signal, requiring wideband correlation processing, with the constantly evolving disturbance field creating additional obstacles for conventional spectral analysis. Suffice to say, it appears that the ELOISE data set alone is not sufficient to extract a meaningful inter-site relationship between most quasi-periodic signatures that may or may not be TID-related. The focus will therefore remain on the most promising transverse TID cases, which the oblique AoA system is most sensitive to.

Figure 8.32 shows a sequence of TID parameter estimates derived using the Bramley ripple transformation on F2-low peak data from five medium-length ELOISE AoA paths (1383–2055 km); the midpoint locations are mapped in **Figure 8.33**. From top to bottom, panels 1 and 2 represent similar path orientations (having midpoint bearings in the range 29–37 °T), and likewise for panels 4 and 5 (64–65 °T), while panel 3 is markedly different (-19 °T). In all but panel 3, the majority of the strongest amplitude points are clustered in the vicinity of -50 °T, suggestive of a consistent north-westward TID direction that persists for much of the night, although common features are not easily identifiable. However, the Kalkarindji to Coondambo results in panel 3 cast
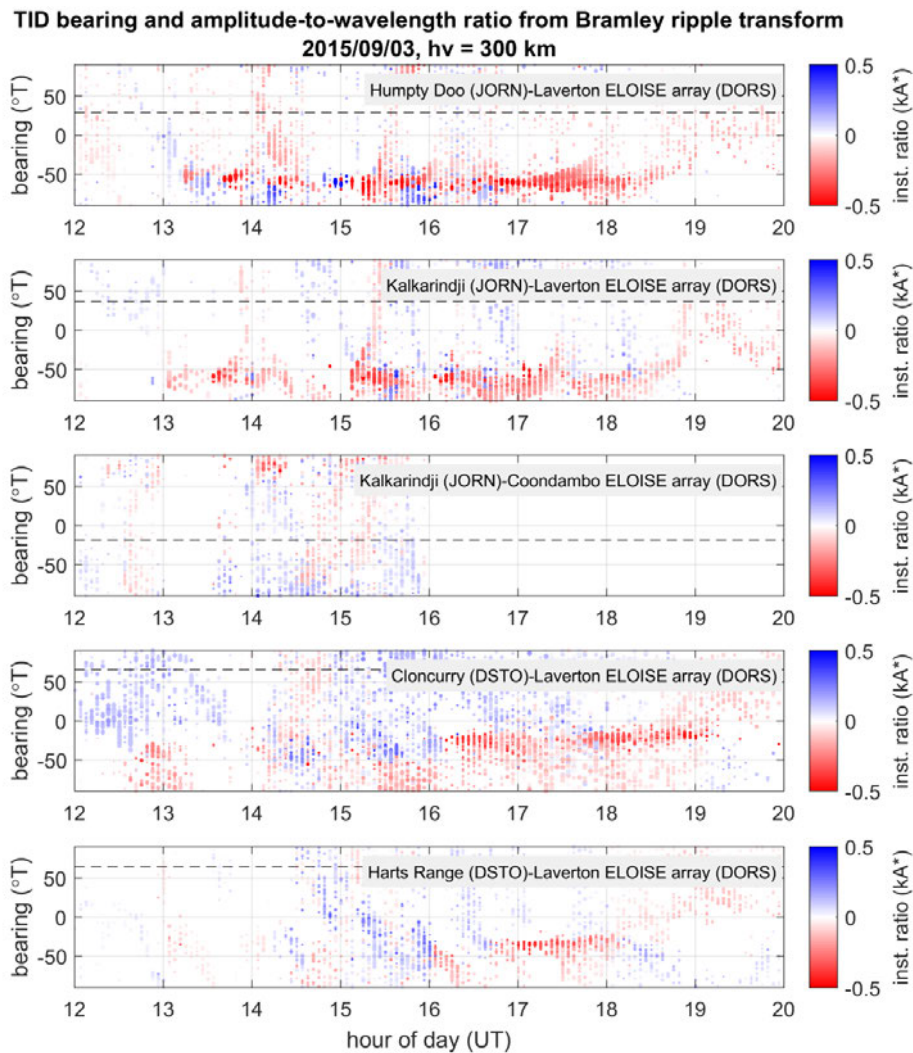
Figure 8.32. Time series of transformed TID bearing and instantaneous amplitude-to-wavelength ratio estimates (in radians), for 1-hop F2-low reflection points in a virtual height window of 300 ± 5 km. The five panels represent data from five medium-length ELOISE AoA paths (1383–2055 km) on 3 September 2015, including the earlier Humpty Doo to Laverton results from **Figure 8.24** in the top panel. The corresponding midpoint locations are shown in **Figure 8.33**, with panels ordered from the most northerly midpoint (top panel) to the most southerly (bottom panel). Again, the dashed horizontal lines show the unperturbed path bearing at the midpoint. Note that the Coondambo ELOISE array was not operating past 16 UT on this day, hence the data gap in the third panel.
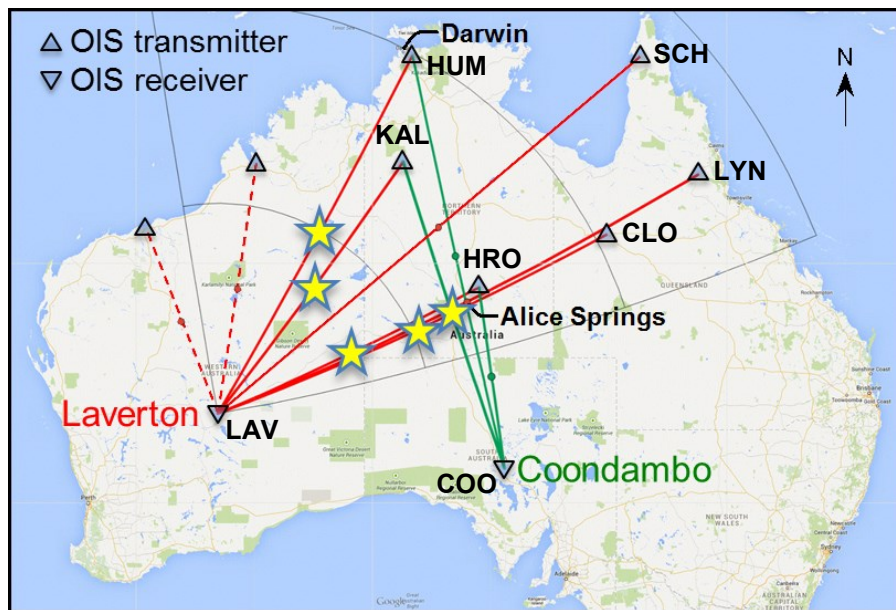
Figure 8.33. Midpoint locations (as yellow stars) for the five medium-length ELOISE AoA paths depicted in **Figure 8.32**. Four of those selected were received at Laverton (red), while the other was received at Coondambo (green).

serious doubt on this interpretation, with no evidence of a banded TID signature constrained in bearing, or indeed any obvious consistency with the other paths. Admittedly these TID directions are unfavourable (longitudinal) for the Coondambo paths in general, and moreover, the effects of ion drag and wind filtering are likely to preferentially deplete waves in those directions that are most favourable (transverse).

As an aside, although this chapter has concentrated on the 1-hop F2-low mode thus far, the same analysis can of course be carried out on the F2-high mode. An example of transformed F2-high peaks from five paths with near-collinear midpoints is shown in **Figure 8.34**, with the locations again mapped in **Figure 8.35**. Unfortunately, the difficulty with identifying TIDs in the high-angle data is that the underlying mirror assumption is less appropriate, with rays refracted through a greater volume, and the impact of large-scale tilts is more prevalent. In this case, no fixed-bearing TID signatures are obviously visible, yet there are quasi-periodic fluctuations in the transformed bearing that are reasonably similar on both Laverton and Coondambo paths (e.g. the negative excursion around 2–3 UT). This is despite a 60–80° difference in the midpoint path bearings, and much longer path lengths into Laverton (up to a factor of

three between the top Coondambo panel and bottom Laverton panel). Without a convincing north-to-south lag, from top to bottom panels, there is no reason to believe these patterns originate from TIDs, rather than more generic time-varying ionospheric gradients. Nevertheless, this figure does highlight the broader usefulness of such a transformation for relating observations on multiple oblique paths.
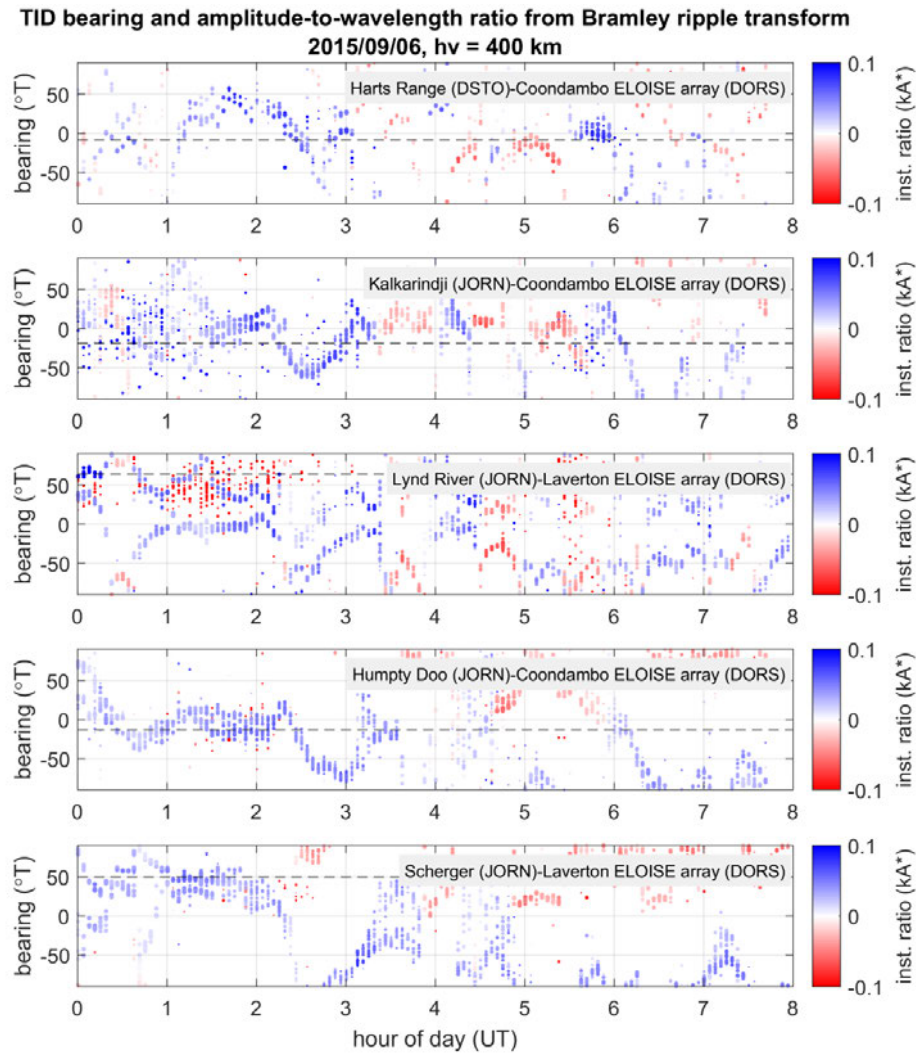
Figure 8.34. Time series of transformed TID bearing and instantaneous amplitude-to-wavelength ratio estimates (in radians), for 1-hop F2-high reflection points in a virtual height window of 400 ± 5 km. The five panels represent data from five intersecting ELOISE AoA paths (908–2712 km) with collinear sample points on 6 September 2015. The corresponding midpoint locations are shown in **Figure 8.35**, with panels ordered from the most northerly midpoint (top panel) to the most southerly (bottom panel). Note that Kalkarindji to Coondambo (panel 2) and Lynd River to Laverton (panel 3), in particular, have very similar 1-hop midpoints, separated by only 100 km.
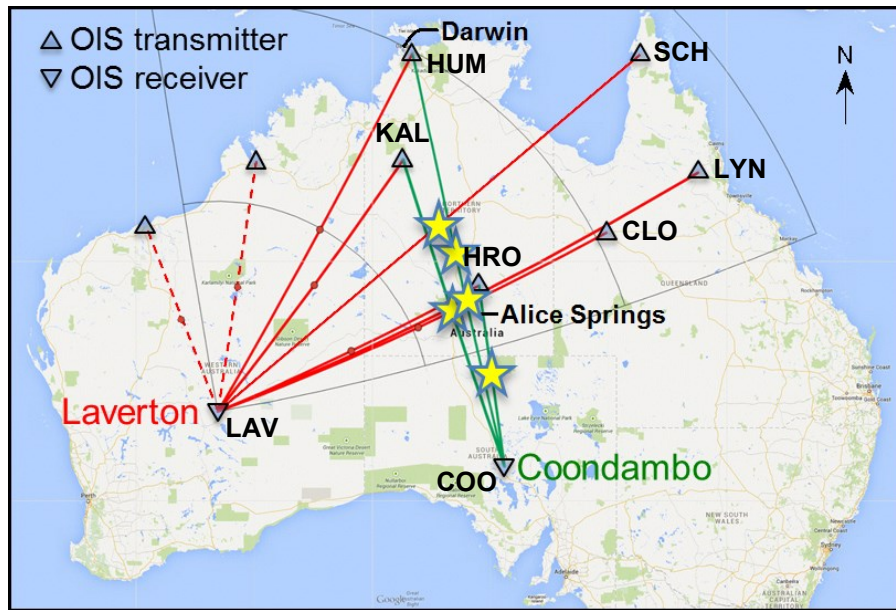
Figure 8.35. Midpoint locations (as yellow stars) for the five intersecting ELOISE AoA paths with collinear sample points depicted in **Figure 8.34**. Two of those selected were received at Laverton (red), while the other three were received at Coondambo (green).

An alternative approach that does not seek to explicitly parameterise the TID, but which still offers a means of relating signatures on different ELOISE AoA paths, is to make use of the virtual height data from reflection point mapping. Equivalent iso-ionic contours, as illustrated in **Figure 8.18**, should be largely agnostic to the oblique path geometry, although the point at which the overhead profile is sampled is strictly not fixed to the midpoint. Horizontal phase velocities are computed with a conventional spaced-sensor correlation technique (e.g. Harris et al. [2012]), based on interpreting spatial lags as apparent speeds between pairs of AoA paths. In theory, any set of three non-collinear paths yields a 2D horizontal velocity vector under the assumptions of fixed sample points and a simple frontal disturbance moving with constant velocity. To help satisfy the latter assumption, cross-spectra are used to isolate a narrowband subset of TID frequencies, in which it is plausible to have just one wave component.

A 12-hour collection of night-time virtual height data at an equivalent vertical-incidence frequency of 3 MHz is presented in **Figure 8.36**. This represents 1-hop F2-low returns on eight Laverton AoA paths (with Coondambo not operating during this time). While it is hard to identify a clear TID signature in the time series data (top and middle panels,

252

before and after detrending), the paths do exhibit many similar features and there is evidence of spatial lags. Calculating the power spectra over the entire interval (bottom panel) reveals candidate periodicities that may correspond to TID activity, at least in an aggregate sense, such as the broad peaks around 30–60 min.
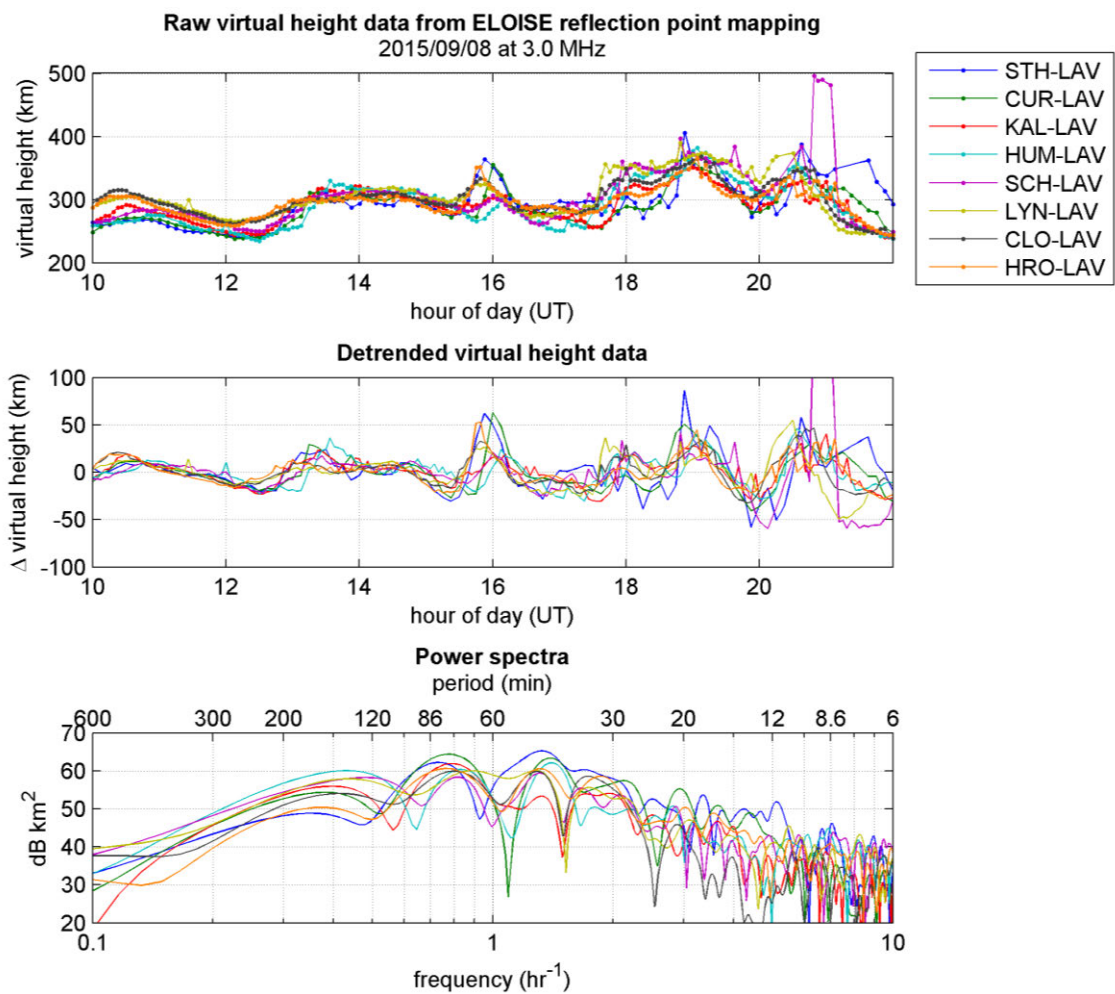
Figure 8.36. Time series plots showing 12 hours of night-time virtual height data and corresponding power spectra for all eight ELOISE AoA paths into Laverton on 8 September 2015. A fixed equivalent vertical-incidence frequency of 3 MHz, for 1-hop F2-low propagation, was sourced from the mapped ionogram peaks and corresponds to one of the iso-contours in **Figure 8.18**. This assumes a tilted mirror model, as described in **Section 4.4**. The raw virtual height data (top panel) is first interpolated and detrended (middle panel), using a ±1 hour (33 point) moving average to remove slowly varying changes in the background ionosphere. The power spectra (bottom panel), computed from the detrended data with a Hann taper applied, are then used to identify common periodicities within the range of typical TIDs (notionally 15 min to 3 hr).

Further processing of the cross-spectra between pairs of sounder paths then enables the determination of drift velocities corresponding to these peaks, as shown in **Figure 8.37** for the eight triangulated sectors bounded by the midpoints. It must be emphasised that these estimates do not represent a single TID spanning the full 12 hours but are merely indicative of the prevailing speed and direction of propagation for a ~45 min class of disturbance. On the whole, there is reasonable agreement in the horizontal direction of propagation (predominantly northward), although the speeds are not entirely consistent. The smaller sectors are likely to be more susceptible to horizontal displacements in the reflection points, contributing to a larger relative error in the baselines and thus the apparent speeds, while the more obtuse triangles (such as between South Hedland, Curtin and Humpty Doo paths in the top left) naturally provide a less ideal projection for certain directions. In contrast, the three larger and almost right-angled triangles in the top right produce fairly consistent speeds, between 237 and 275 m/s, and directions, between -28 and -18 °T.
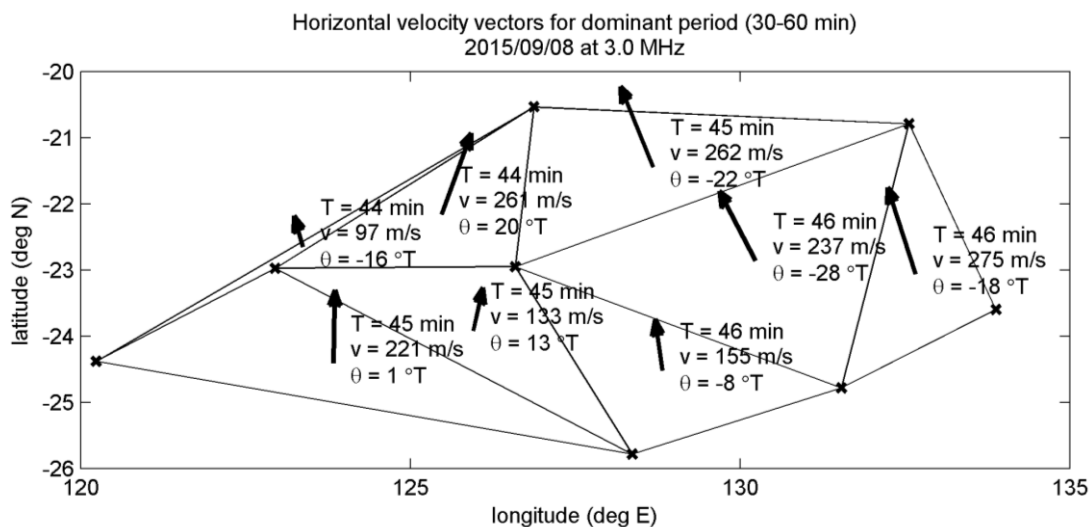
Figure 8.37. Horizontal phase velocity vectors calculated from the cross-spectral phases of detrended ELOISE virtual height data. For a given path triplet, the dominant 30–60 min periodicity has been selected, representing one of the main peak locations in the power spectra of **Figure 8.36**. The period, speed and direction (bearing) are printed alongside each vector arrow, which sits within a triangle bounded by three 1-hop midpoints (crosses).

The main limitation of this cross-spectral analysis approach is finding suitable data sets that are both sufficiently long and stable in frequency. Large-amplitude TIDs once again cause key assumptions to break down, leaving the velocity estimates very prone to errors for some of the better ELOISE TID candidates. Given that only the relative phase between paths is important here, the lag between virtual height and true height [Pederick et al., 2017] is believed to be a lesser issue, and that for oblique incidence soundings, errors involved in true height inversion (usually neglecting off-angle geometry) are likely to be far more significant.

There is clearly more work to be done in this area, to strengthen the understanding of relationships between the various ELOISE AoA paths. Wavelet analysis [Torrence & Compo, 1998] may be considered in future as a means of better capturing the non-stationary nature of TID observations, although care must be taken as results are highly dependent on the choice of mother wavelet, and the wavelet cross-spectrum for estimating lags may contain spurious peaks [Maraun & Kurths, 2004]. Multivariate empirical mode decomposition (EMD) [Huang et al., 1998; Rehman & Mandic, 2010], in conjunction with the Hilbert transform, has also been tested on the ELOISE data set, however the output was found to be very sensitive to noise and outliers; even a small amount of additive noise would drastically alter the intrinsic mode functions (IMFs), thereby making a physical interpretation difficult. As such, this approach was not pursued further.

8.5   Chapter summary

- Signatures of travelling ionospheric disturbances have been identified in the ELOISE AoA data. Comparisons with synthetic delay, Doppler and AoA observables from two candidate TID models show encouraging results, in terms of their ability to classify and parameterise the disturbance field. Despite its limitations, the simple corrugated mirror can still be effectively used to identify general classes of TIDs under certain conditions, although it is noted that some of the parameters matched by eye may be relatively imprecise. On the other hand, the AGW-seeded (Hooke) approach with 3D NRT can potentially offer a more realistic TID representation at higher altitudes (i.e. for which refractive effects are more

significant) but may be limited in its ability to fully capture the properties of some larger-scale disturbances.

- Using the corrugated mirror model, the relationship between the ionospheric height perturbation and HF observables has been explored as a function of key TID parameters. Many of the classic attributes of TID signatures can be reproduced by this simple model, including the transition from the sinusoidal to non-sinusoidal observable regime, presence of multiple simultaneous reflections, and spread-like behaviour arising from the superposition of even a modest number of TID components. In particular, the response is shown to become increasingly non-linear as the TID amplitude-to-wavelength ratio and/or oblique ground range is increased. For medium to long paths, the curvature of the Earth/ionosphere introduces additional errors. Perhaps most importantly, the model predicts a strong directional bias in favour of transversely propagating disturbances on representative ELOISE AoA paths.

- Without sufficient diversity in the path orientations (noting that the Coondambo site was not operating continuously), it is difficult to separate true anisotropy in the distribution of TIDs from this directional filtering effect. In hindsight, the alignment of the three AoA paths into Coondambo was not favourable for detecting TIDs with directions in the preferred north-westward quadrant, as this represents a less sensitive longitudinal geometry. Conversely, as an aside, such path bearings are likely to be strategically good look-directions for OTHR operations; that is, to avoid the detrimental influence of TIDs.

- Possible techniques for analytic TID parameter inversion, based on an assumed mirror reflector and direct geometrical transformations, have been considered and tested. With the use of synthetic observables, for various corrugated mirror models, it was confirmed that the TID period is in general the most robust and accurate parameter to extract, followed by bearing (azimuth), then amplitude-to-wavelength ratio; it is harder again to separate these last two parameters reliably. The Bramley ripple transformation is extremely elegant, but only offers a partial TID characterisation (with an azimuthal ambiguity) and is fundamentally based on a flat Earth/ionosphere assumption (although this can be corrected to some extent).

However, applying the transformation to actual ELOISE AoA observations left a rather unconvincing impression of its efficacy.

- The Frequency and Angular Sounding (FAS) technique extends this further to support separation of multiple TID components, but multiple reflections must be ignored. Unfortunately, for a typical oblique path from the ELOISE trial, only very shallow and likely undetectable TID amplitudes and/or long wavelengths yield observables that are even close to sinusoidal and, thus, appropriate for conventional spectral analysis within the FAS technique. The implications of larger amplitude disturbances and longer oblique paths on the FAS technique have not been previously published to the author's knowledge.

- While broad features can be reproduced by the TID models, at least qualitatively and for a single plasma frequency (iso-ionic) surface, more work needs to be carried out to interpret the many examples of AoA perturbations that do not have just one dominant and time-invariant periodicity. With TIDs being so ubiquitous, the majority of the ELOISE observations in fact represent the superposition of many components, from different sources, which may evolve over time or exist as solitary waves. Efforts to model each spectral component in a spatially consistent way across the network of ELOISE AoA paths are ongoing. A key enabler for permitting cross-spectral analysis between oblique paths is a method of transforming the observables to a common domain that is agnostic to the widely different path lengths: both the Bramley transformation and virtual heights from (off-angle) reflection point mapping have been proposed as potential options.

9   Summary and conclusions

9.1   Thesis summary

This thesis has described the design, calibration, analysis and interpretation of results from a new oblique angle-of-arrival (AoA) and Doppler ionosonde, aimed at attaining a better understanding about the nature of mid-latitude ionospheric disturbances over Australia and their impact on high frequency (HF) sky-wave propagation. The outcomes seek to inform and validate the ionospheric modelling and coordinate registration strategies for over-the-horizon radar systems such as the Jindalee Operational Radar Network (JORN). Over the course of the ELOISE experiment in September 2015, the AoA ionosonde measured the received power (path loss), group delay, bearing and elevation across eleven intersecting paths of between 900 and 2700 km, nine of which were also equipped with a channel scattering function (CSF) mode to generate delay-Doppler spectra, interleaved with the ionograms on three frequency-agile channels. All manner of off-angle perturbation signatures were observed in the soundings, including systematic bearing/elevation shifts due to horizontal gradients (particularly around the solar terminator), quasi-periodic trace distortions and discrete "satellite" traces (associated with the passage of travelling ionospheric disturbances or TIDs), and spread-F conditions at night.

The ELOISE AoA array was a 19-element configuration of vertical monopoles in two orthogonal arms, with a high-fidelity direct-digital HF receiver per element capable of simultaneously collecting soundings on multiple oblique paths. Two separate receiving systems were installed, near Laverton, Western Australia and Coondambo, South Australia, to offer some diversity in the path orientations from a set of eight down-range transmitters. Despite only using low-power (~20 W) chirp signals, the ionosonde offers excellent sensitivity to weaker off-angle propagation modes through its directive gain, and good bearing and elevation precision through its significant aperture size (albeit

with limited mode resolvability). Median AoA uncertainties of 0.2–0.5° in bearing and 0.4–0.9° in elevation, depending on the path geometry, were typical of the ionogram peaks, whereas the CSF peak uncertainties were roughly half that, likely due to their higher signal-to-noise ratios.

Array calibration efforts drew from an assortment of known line-of-sight, surface-wave and aggregated sky-wave sources in order to estimate the overall array alignment and inter-element phase delay offsets. Results were also successfully validated against the co-located JORN radar at Laverton. A low-elevation bias was identified on the longer paths, requiring an additional empirical correction to combat the overestimation of elevation below 10°; this is suspected to be partly a consequence of the 0° null in the antenna gain pattern. Simplified models of tropospheric refraction and antenna mutual coupling effects were unable to conclusively account for the dependence in both bearing and elevation dimensions, although in theory it was noted that mutual coupling can lead to errors on a similar scale as the observed biases.

Owing to the sheer volume of soundings (i.e. eleven AoA paths with a typical revisit interval of 3.75 min) and the remoteness of sites, automatic on-board signal and image processing was considered essential to the experiment. A complete suite of software was therefore assembled by the author, building on DST Group's existing sounder code base, to provide the fundamental ionogram and delay-Doppler map construction, radio frequency interference rejection, ionogram feature extraction and midpoint profile fitting, clear channel evaluation and adaptive CSF scheduling, and 2D AoA (bearing/ elevation) estimation. The standard output was a compressed binary file, suitable for real-time transfer back to a central data server. Peaks were later fitted to the image data and classified by propagation mode to provide a suitable form for further analysis, including reflection point mapping and tilt estimation. The F2-low mode served as the main focus because of its importance to radar operations and the likelihood of larger disturbance amplitudes at these heights.

The interaction between unresolved modes and small-scale temporal variability was investigated in some detail by breaking apart the CSF dwells into 64 individual sweeps and looking at how their peak power, AoA and Doppler fluctuated over the 6.4 s standard processing interval. Rapid signal fading of up to 20–30 dB and AoA jumps of up to 20° were associated with intra-dwell mode beating, with night-time F2-low being

the most susceptible, probably due to both the reduced CSF bandwidth allocations (i.e. poorer delay resolution) and likelihood of unseparated O/X components. A combination of metrics was proposed to help detect the presence of non-planarities and increased temporal AoA variance among the CSF peaks.

Medium- to large-scale variability in the peak data has been analysed in terms of diurnal patterns in the observables, zonal and meridional tilt estimates, drift velocity estimates, and case studies of TID signatures. Geophysical features, such as the sharp Doppler excursions and zonal gradients around the dawn/dusk terminators, appear largely in line with expectations, although the challenge is relating observables over vastly different oblique path geometries. A tilted spherical mirror model was used throughout as the basis for transformations to the path-agnostic domain of the iso-ionic reflecting surface. For oblique incidence, a flat Earth assumption is of course not appropriate, making the problem more difficult to solve analytically than at vertical incidence.

As a validation exercise, statistical comparisons were carried out between the ELOISE AoA observables and ray synthesis through a conventional data-assimilative model, the so-called Gardiner-Garden Ionospheric Model (GGIM), based on temporal filtering and spatial mapping of midpoint profile data. The results revealed that, for the most part, the model is able to reproduce the medium- to large-scale variability without direct AoA or Doppler inputs. In fact, fitted midpoint profiles from a single path alone can explain the majority of Doppler excursions, just from the time rate of change of the synthesised phase path. What remains are components at the medium-scale end of the disturbance spectrum that have been attenuated by the limited spatial and temporal sampling of the sounder network and subsequent model filtering steps. This typically means that, while the same perturbations are registered in the synthetic predictions, the AoA offsets tend to be underestimated. An additional comparison, between the AoA-derived tilt estimates and horizontal gradients in the GGIM's large-scale anomaly field, produced a similar conclusion, although the discrepancies were somewhat greater, most likely due to the mirror approximation used in determining the reflection points and omission of the residual anomaly field from the GGIM.

Finally, examples of quasi-periodic propagating disturbances from across the ELOISE AoA network have been interpreted as TIDs, the ionospheric manifestation of atmospheric gravity waves (AGWs) in the thermosphere, with attempts made to classify

and parameterise their wave-like properties from sequences of peak data. Two candidate TID models (the corrugated mirror and the Hooke AGW-seeded representation) have been used to synthesise realistic observables via ray tracing and homing, as a means of identifying the broad class of TID, while two direct geometric inversion methods (the Bramley ripple transformation and generalised Frequency and Angular Sounding technique) have been implemented to better quantify the wave parameters. Equatorward and westward propagation was a defining characteristic of the night-time TID cases presented. Among daytime cases, it was rare to find a single dominant periodicity.

The distinctive transition from sinusoidal to non-sinusoidal observables and emergence of multiple simultaneous reflections as the amplitude-to-wavelength ratio and/or ground range increases has been demonstrated with the aid of the above TID models, noting that beyond this point the simple inversion techniques become seriously compromised. For a typical ELOISE AoA path (1989 km), the transition can occur at TID amplitudes of just a couple of kilometres, placing the sinusoidal regime outside the limits of detectability. Another key result was establishing the strong directional bias in the amplitude of the HF response to TIDs, in favour of disturbances propagating transversely to the oblique path. Combined with the fact that elevation uncertainties translate to much greater down-range errors on longer paths, this has general implications on the design of AoA sounder networks for TID monitoring purposes; that is, shorter paths with more diverse orientations are preferable.

The part of the analysis that ultimately proved to be the most elusive was finding coherent propagating signatures across the ELOISE AoA network. Among the examples shown, paths with similar lengths and orientations were reasonably well-correlated, even with midpoints spaced hundreds of kilometres apart, but once the geometries diverged, this no longer held true. This may be partly a statement on instrumental limitations (particularly sensitivity and accuracy on longer paths), as well as the underlying assumptions of mirror reflection and a static ("frozen-in") disturbance field that were used to relate observables from the different paths. The overall lack of night-time observations from Coondambo also hindered this effort.

9.2   Future work

Being a new and novel ionosonde design, the ELOISE experiment represented the first time the AoA system had been put through its paces, and as would be expected, a number of lessons were learnt that will be incorporated into the next such experiment. A key element of any future work should be to collect and process a broader set of measurements (with more path diversity, particularly at shorter ranges) and investigate ways to relax some of the assumptions involved in the analysis. By transitioning more of the processing to run on-board the sounder, a longer-term synoptic-style analysis, incorporating greater data volumes becomes more feasible. Valuable insights would be gained from looking at the AoA and Doppler variability across multiple seasons, during both quiet and disturbed geomagnetic conditions, and characterising the spatial decorrelation scales for these observables.

In terms of the instrument itself, while there was reasonable confidence in the relative precision of bearing and elevation estimates, there still remain outstanding questions over the absolute AoA accuracy at low elevations. Ahead of the next deployment, an improved on-site calibration procedure (e.g. using a mobile emitter or direct signal injection) is required to alleviate mutual coupling concerns, and alternative array layouts should be considered to increase the average separation between antennas (e.g. with a minimum redundancy configuration) and/or deliver greater symmetry between the embedded elements (e.g. with a uniform circular array). Other calibration sources such as line-of-sight satellite/stellar objects, HF emitters of opportunity, and meteor echoes could also be explored. This will help ensure that ionospherically-induced wavefront distortions remain the limiting factor, not array performance.

Despite the ELOISE collection being focused on medium- to large-scale disturbances, it was not surprising to find evidence of smaller-scale irregularities contributing to the overall variance, as quantified by the intra-dwell CSF statistics. The inclusion of a faster revisit rate or longer duration ("stare mode") CSF dwells within the timeline of a future experiment would be highly desirable, subject to spectral licensing. Alternatively, passive reception of known HF emitters could be exploited to similar effect. Mode resolvability in the CSF dwells also has the potential to be improved by implementing high-resolution spectral analysis techniques.

The ability to scale zonal and meridional tilts from the AoA ionosonde presents an interesting opportunity for future data assimilation into ionospheric models. Although more validation is needed, the level of consistency with spatial maps of midpoint profile parameters suggests that the tilt estimates could add value, particularly in regions where conventional sounder sites are more sparsely distributed (e.g. over oceans). A better sense of the horizontal distance over which the tilts apply must first be investigated, however some preliminary comparisons against JORN transponder peaks (not shown here) indicate that coordinate registration can benefit out to a midpoint separation of at least 30 km.

With regard to TID identification and parameterisation, there is still much work to do in order to make the process more automated and robust to errors, particularly when faced with disturbance fields consisting of multiple components. Possible enhancements to the TID models include generalising for a background ionosphere that varies in space and time, allowing for the prevailing neutral winds, and modifying the functional form of the perturbation to take into account wave evolution over long distances. Relaxing the assumption of specular reflection as part of the TID parameter inversion should be a high priority, although the extra complexity of a non-analytic/iterative procedure may be insurmountable given the number of degrees of freedom. Joint profile inversion and horizontal tilt/TID estimation is also worthy of consideration. In addition, there have so far been no attempts to use modulations in the signal amplitude to evaluate the ray focusing/defocusing effects of TIDs, or multi-hop returns as a means of tracking spatial lags (even if their ionospheric sample points are ambiguous). Likewise, the differences between O and X modes, though difficult to separate without polarisation information, and between E and F height regions are important facets of the TID characterisation problem yet to be examined using the ELOISE AoA peak data.

An opportunity for another experimental collection would of course help in building up a climatological picture of TIDs in the region, including patterns of occurrence and preferred directions/velocities of travel. Continuous (24/7) data collections are essential to this. While many papers have looked at the TID response to isolated events, their quiet-time behaviour is comparatively less studied, and ionosonde observations are well-placed to capitalise on this. Both the Digisonde DPS-4D and VIPIR/Dynasonde have already taken steps towards addressing this, at least at vertical incidence. To truly demonstrate the worth of oblique incidence observations to this research area, the

ELOISE AoA ionosonde ideally needs to be run alongside a direction-finding vertical incidence sounder at the midpoint. It is then a matter of using a set of staggered path lengths to determine how the HF response transitions between these two geometries.

Beyond the scope of this project, but certainly within the objectives of ELOISE as a whole, is a more systematic comparison of observations by different instruments; the building blocks for this are described by Gardiner-Garden et al. [2019]. The AoA ionosonde is but one piece of a bigger puzzle, and the software and insights developed here are hoped to contribute in many ways to DST Group's experimental HF radar program in future.

Appendix A   Fundamentals of chirp sounding

The ELOISE AoA sounding system makes oblique measurements of the ionosphere by receiving a linear frequency-modulated continuous-wave (LFMCW) waveform from a network of down-range transmitter sites. Each array element, with its own direct-digital receiver, is analysed using standard chirp-sounder principles, which are outlined in this appendix for the wideband (ionogram) sweep. Analysis of the channel scattering function waveform is similar, except that the stretch processing described below is replaced by a range correlator (matched filter) and Doppler processing. The on-board code was largely adapted from DST Group's existing DORS oblique incidence sounder (OIS) ionosonde, originally designed for omni-directional reception. Refer also to Ivanov et al. [2003] for the detailed history, principles and applications of chirp soundings to ionospheric research.

The transmitted OIS signal (with quadratic phase) can be expressed as

$$x_T(t) = \exp[i\Phi(t)] = \exp\left[i\left(2\pi\left(f_0 t + \frac{Kt^2}{2}\right) + \varphi_T\right)\right], \tag{A.1}$$

where f_0 is the start frequency, K is the sweep rate, and φ_T is a fixed but potentially unknown phase offset.

Note that instantaneous sweep frequency $f(t)$ is related to the time-varying phase of the complex exponential $\Phi(t)$ by

$$f(t) = \frac{1}{2\pi}\frac{d\Phi}{dt} = f_0 + Kt. \tag{A.2}$$

This is the standard equation for a linear sweep.

The received OIS signal on antenna element n for a single propagation mode (without regard for signal amplitude, phase distortions or noise) is

$$x_{R,n}(t) = \exp\left[i\left(2\pi\left(f_0(t - T_{g,n}) + \frac{K(t - T_{g,n})^2}{2} + f_{id}t\right) + \varphi_{R,n}\right)\right], \qquad (A.3)$$

where $T_{g,n}$ is the group delay of the one-way ionospheric path to element n, f_{id} is the ionospheric Doppler shift, and $\varphi_{R,n}$ is a fixed unknown phase offset. Both $T_{g,n}$ and f_{id} are assumed to be constant for a suitably short coherent integration time (CIT) and a correspondingly narrow swept frequency interval.

The local oscillator reference for receiver n is

$$x_{LO,n}(t) = \exp\left[i\left(2\pi\left(f_0(t - T_{LO,n}) + \frac{K(t - T_{LO,n})^2}{2}\right) + \varphi_{LO,n}\right)\right], \qquad (A.4)$$

where $T_{LO,n}$ is the preset local oscillator time offset (typically set to the ground range plus 1.6 ms for DORS), and $\varphi_{LO,n}$ is a fixed unknown phase offset.

Applying stretch processing, the deramped and down-converted ("raw data") signal for receiver element n is therefore given by

$$\begin{aligned} y_{R,n}(t) &= x_{LO,n}^*(t)x_{R,n}(t) \\ &= \exp\left[i\left(2\pi\left(-f_0 T'_{g,n} - KT'_{g,n}t + \frac{K(T_{g,n}^2 - T_{LO,n}^2)}{2} + f_{id}t\right) + \vartheta_{R,n}\right)\right], \end{aligned} \qquad (A.5)$$

where relative group delay $T'_{g,n} = T_{g,n} - T_{LO,n}$, and fixed phase $\vartheta_{R,n} = \varphi_{R,n} - \varphi_{LO,n}$.

The signal ("beat") frequency is again the derivative of the time-varying phase of the complex exponential. From equation (A.5),

$$f_B(t) = -KT'_{g,n} + f_{id}. \qquad (A.6)$$

Note that the beat frequency is negative for positive $T'_{g,n}$ and includes a Doppler coupling component; a negative f_{id} increases the apparent group delay for a positive sweep rate. For now assume that $f_{id} = 0$.

What is actually measured is a sampled (digitised) version of the down-converted signal in equation (A.5); i.e. $y_{R,n}[t]$. Technically on a direct digital receiver such as DORS, the down-converted signal is calculated directly from sampled signals $x_{R,n}[t]$ and $x_{LO,n}[t]$. Sampling noise is not modelled in the formula above.

The OIS ionogram is formed by taking the discrete Fourier transform (DFT) of $y_{R,n}[t]$ over a sequence of finite time intervals (CITs) corresponding to different segments of the sweep. Typically this is evaluated using the fast Fourier transform (FFT) algorithm, at frequencies $f_m = \frac{m-1}{M}\frac{1}{T_s}$, where $m = 1, \dots, M$, given M time samples with sampling interval T_s. Note that in this context the frequencies of interest are on the scale of the beat frequency, not the wideband sweep frequency (i.e. Hz not MHz).

Consider a particular ionogram bin for which the CIT includes the samples $t_m = 0, \dots, (M-1)T_s$. The centre sweep frequency is therefore $f_c = f_0 + \frac{f_{bw}}{2}$, where $f_{bw} = K(M-1)T_s = 200\,\text{kHz}$ for typical DORS processing. Excluding windowing and normalisation, for sake of simplicity, the DFT of $y_{R,n}[t]$ is given by the following:

$$
\begin{aligned}
Y_{R,n}[f] &= \sum_{m=1}^{M} y_{R,n}[(m-1)T_s]\exp[-i(2\pi f(m-1)T_s)] \\
&= \exp\left[i\left(2\pi\left(-f_0 T'_{g,n} + \frac{K(T^2_{g,n}-T^2_{LO,n})}{2}\right) + \vartheta_{R,n}\right)\right] \\
&\quad \times \sum_{m=1}^{M} \exp\left[-i\left(2\pi(f + KT'_{g,n})(m-1)T_s\right)\right] \\
&= \exp\left[i\left(2\pi\left(-f_0 T'_{g,n} + \frac{K(T^2_{g,n}-T^2_{LO,n})}{2}\right) + \vartheta_{R,n}\right)\right] \\
&\quad \times \frac{1-\exp\left[-i\left(2\pi(f + KT'_{g,n})MT_s\right)\right]}{1-\exp\left[-i\left(2\pi(f + KT'_{g,n})T_s\right)\right]} \\
&= \exp\left[i\left(2\pi\left(-\left(f_0 + \frac{K(M-1)T_s}{2}\right)T'_{g,n} + \frac{K(T^2_{g,n}-T^2_{LO,n})}{2}\right) + \vartheta_{R,n}\right)\right] \\
&\quad \times \exp[-i\pi f(M-1)T_s]\frac{\sin\left(\pi(f + KT'_{g,n})MT_s\right)}{\sin\left(\pi(f + KT'_{g,n})T_s\right)} \\
&= M\exp\left[i\left(2\pi\left(-f_c T'_{g,n} + \frac{K(T^2_{g,n}-T^2_{LO,n})}{2}\right) + \vartheta_{R,n}\right)\right] \\
&\quad \times \exp[-i\pi f(M-1)T_s]\,\text{psinc}_M\left(2\pi(f + KT'_{g,n})T_s\right),
\end{aligned}
\tag{A.7}
$$

after applying the formula for the sum of a finite geometric series and simplifying in terms of the periodic sinc function $\mathrm{psinc}_M(x) = \frac{\sin\left(\frac{Mx}{2}\right)}{M\sin\left(\frac{x}{2}\right)}$.

It is clear the amplitude of $Y_{R,n}[f]$ (ignoring normalisation) is maximised when $f = f_B = -KT'_{g,n}$. In practice, because the FFT only samples this function at discrete f_m, there may be some scalloping loss.

Now consider such measurements from across an array of N elements, with a separate receiver per element. Express the group delay for element n in terms of a reference group delay for the array as a whole T_g and a spatially varying delay offset τ_n:

$$T_{g,n} = T_g + \tau_n. \tag{A.8}$$

Similarly, express the local oscillator time offset for receiver n in terms of a reference time offset T_{LO} and a receiver-dependent time offset v_n:

$$T_{LO,n} = T_{LO} + v_n. \tag{A.9}$$

The relative group delay for element n is given by the difference between the two:

$$T'_{g,n} = T'_g + \tau_n - v_n, \text{ where } T'_g = T_g - T_{LO}. \tag{A.10}$$

Hence, the ionogram bin at frequency $f = -KT'_g$ can be expressed as

$$
\begin{aligned}
Y_{R,n}\big[-KT'_g\big] = \exp\Bigg[& i\left(2\pi\left(-f_c(\tau_n - v_n) + \frac{K(2T_g\tau_n + \tau_n^2 - 2T_{LO}v_n - v_n^2)}{2}\right)\right. \\
& \left. + \vartheta_{R,n}\right)\Bigg] \times \mathrm{psinc}_M\big(2\pi K(\tau_n - v_n)T_s\big)\, A_R\big[-KT'_g\big],
\end{aligned}
\tag{A.11}
$$

where $A_R[f]$ incorporates all of the terms that are not spatially varying.

The geometry of a planar wavefront suggests measured time offsets of $\tau_n = \frac{\boldsymbol{k}^T \boldsymbol{u}_n}{2\pi f_c}$, for signal wave vector $\boldsymbol{k} = -\frac{2\pi}{\lambda_c}[\sin\theta\cos\beta, \cos\theta\cos\beta, \sin\beta]^T$ and element position vector $\boldsymbol{u}_n = [x_n, y_n, z_n]^T$. Note that θ and β are the azimuth and elevation of the signal with

respect to the receiver array (i.e. in the opposite direction to the signal wave vector itself). For a uniform linear array (ULA) along the x-axis with element spacing d, the position vector takes the simpler form $\boldsymbol{u}_n = [nd, 0, 0]^T$ and, hence, $\tau_n = -\frac{nd\sin\theta\cos\beta}{c}$.

The receiver-dependent components v_n of the local oscillator time offsets are zero by default, but can also be configured to have a similar spatial variation as for the measured time offsets τ_n; that is, in the form of $v_n = \frac{\boldsymbol{k}_0^T \boldsymbol{u}_n}{2\pi f_c}$, where wave vector $\boldsymbol{k}_0$ defines some arbitrary reference direction. In practice, non-zero values of v_n are most easily implemented by applying phase shifts after the fact to the down-converted signal (i.e. immediately prior to calculating the ionogram). This avoids having a different sounding schedule per element, and overcomes restrictions posed by the digital sampling limit on the amount by which the local oscillator can be shifted.

Classical beamforming combines the N receiver elements through multiplying and summing by a steering vector equal to (without any array tapering)

$$
\boldsymbol{v} = \begin{bmatrix} \exp(-i\boldsymbol{k}'^T \boldsymbol{u}_1) \\ \vdots \\ \exp(-i\boldsymbol{k}'^T \boldsymbol{u}_N) \end{bmatrix}.
\tag{A.12}
$$

For maximum beamformer output in the ULA case, where the array has been calibrated to remove the $\vartheta_{R,n}$ terms (or phase-locked if running off a common timing source), the beamformer wave vector should be set to $\boldsymbol{k}' = \boldsymbol{k} - \boldsymbol{k}_0$ (where $\boldsymbol{k}_0$ is ordinarily $\boldsymbol{0}$). After assembling the (complex) ionogram bins from each receiving element into vector $\boldsymbol{Y}_R$, this gives, for bin frequency $f = -KT_g'$, a beamformer output of

$$
\begin{aligned}
\tilde{Y}_R\left[-KT_g'\right] &= \boldsymbol{v}^H \boldsymbol{Y}_R\left[-KT_g'\right], \text{ where } \boldsymbol{v} = \begin{bmatrix} \exp(-i2\pi f_c(\tau_1 - v_1)) \\ \vdots \\ \exp(-i2\pi f_c(\tau_N - v_N)) \end{bmatrix} \\
&= A_R\left[-KT_g'\right] \sum_{n=1}^{N} \exp\left[i\left(2\pi\left(\frac{K(2T_g\tau_n + \tau_n^2 - 2T_{LO}v_n - v_n^2)}{2}\right)\right)\right] \\
&\quad \times \mathrm{psinc}_M(2\pi K(\tau_n - v_n)T_s).
\end{aligned}
\tag{A.13}
$$

For good coherence across the phased array, the narrowband assumption requires that (1) $\frac{K(2T_g\tau_n + \tau_n^2 - 2T_{LO}v_n - v_n^2)}{2} \ll 1$ (i.e. phase coherence), and (2) $K(\tau_n - v_n)MT_s \ll 1$ (i.e. amplitude coherence).

Note that with typical ELOISE parameters, namely $K = 250$ kHz/s, T_g and $T_{LO} \approx 10$ ms (worst case for a 3000 km path), $\tau_n \approx 1$ μs and $v_n = 0$ (worst case for a 300 m array at endfire), and $MT_s \approx 0.8$ ms (typical DORS CIT for a 250 kHz/s sweep rate), the left-hand side of requirement (1) should be no larger than ~0.025, but the left-hand side of requirement (2), the time-bandwidth product, could be as large as ~0.2. This is not ideal for phased array beamforming.

To reduce the time-bandwidth product, it is possible to judiciously choose a reference wave vector of, for example, $\boldsymbol{k}_0 = -\frac{2\pi}{\lambda_c}[\sin\theta_0, \cos\theta_0, 0]^T$ for the receiver-dependent time shifts v_n. For the ULA this becomes $v_n = -\frac{nd\sin\theta_0}{c}$, and as such,

$$\tau_n - v_n = \frac{nd}{c}(\sin\theta_0 - \sin\theta\cos\beta). \tag{A.14}$$

These time shifts "refocus" the array so that the effective boresight is now directed at the great circle OIS path azimuth θ_0, improving the narrowband assumption that underpins all phased array processing. Then in the worst case, where θ_0 and θ differ by say 10°, $\tau_n - v_n \approx 0.2$ μs and the time-bandwidth product in requirement (2) is ~0.04. This will improve the performance of the beamformer, by keeping the processing in the narrowband regime. The left-hand side of requirement (1) will also reduce to some extent.

In contrast, for an interferometric approach like the ELOISE AoA algorithm, using elemental phase angles only and ignoring amplitude, only requirement (1) applies (except where SNR becomes critically low). As such, the need to minimise the time-bandwidth product is relaxed. However, in the interests of achieving the best phase precision possible, the use of a unique local oscillator time shift for each element and each path remains desirable.

Appendix B   Imaging on the ELOISE array

The ELOISE twin-arm receiving array was designed primarily to function as an interferometer for AoA estimation. However, as a sparse 2D array, the receiver-per-element data can also be applied to imaging; that is, through the use of phased-array beamforming over the hemispheric field of view. A key advantage of such techniques is their ability to resolve multiple oblique propagation modes in the spatial domain. The purpose of this appendix is to illustrate some results from imaging on the ELOISE array, and highlight the implications of the array design on directional ambiguities and side-lobe performance. A basic understanding of phased-array beamforming techniques is assumed (see e.g. Pillai [1989], Krim & Viberg [1996] or Van Trees [2002]).

For demonstration purposes, consider again the sample CSF dwell from **Chapter 4**; this is reproduced as a delay-Doppler map in **Figure B.1**, scaled in terms of the median power over all array elements. The left panel clearly shows three resolved peaks: the peak at 1490 km is the 1-hop F2-low mode, while the cluster of two peaks at 1630 km and 1660 km are the F2-high O- and X-modes, respectively. To better understand the spatial imaging characteristics of the array, the F2-low mode is initially isolated from the rest of the dwell, as shown in the right panel. All imaging results that follow use the complex CSF data on each antenna element after delay and Doppler processing, and are prior to the bearing and elevation corrections described in **Chapter 5**.

Recall the layout of the ELOISE array in **Figure B.2**. Although each 10-element arm is a uniform linear array, the 2D array is by definition sparse and not symmetric; this was due to the availability of cleared land and other nearby infrastructure which constrained the footprint at Laverton. As such, there is fairly uneven and incomplete sampling of the different inter-element baselines, compared to the equivalent 10×10 filled array. This can be visualised in terms of the difference co-array [Hoctor & Kassam, 1990], shown in **Figure B.3**, which contains the horizontal displacements of all $19^2 = 361$ different
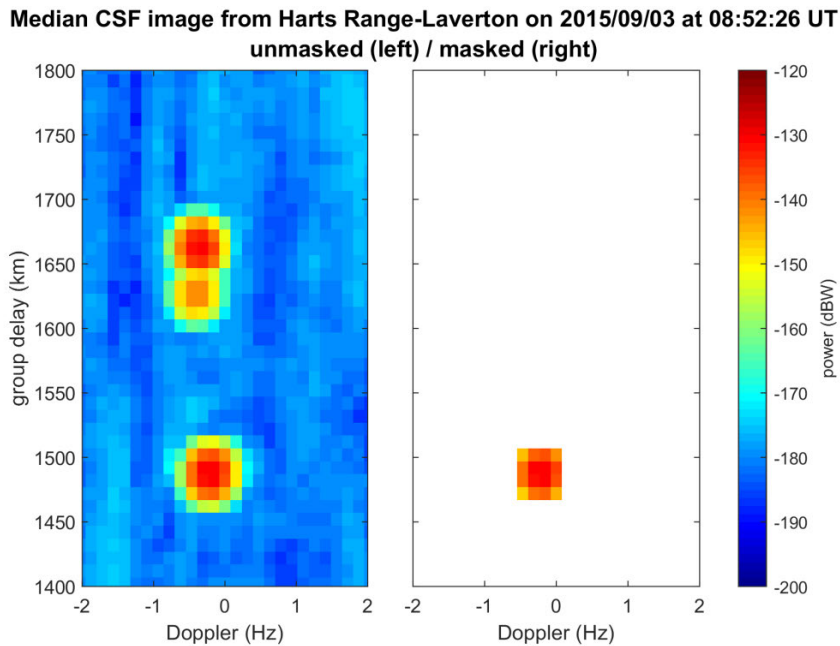
Figure B.1. Sample CSF delay-Doppler image, used to evaluate imaging on the ELOISE AoA array. This same CSF dwell appears throughout **Chapter 4** (e.g. **Figure 4.5**) and was collected at 13.18 MHz on the 1383 km OIS path from Harts Range to Laverton. The left panel shows the median power over all 19 array elements, while the right panel has had masking applied to isolate just the 1-hop F2-low mode pixels.

pair-wise combinations of elements (including with themselves). Since there are only 191 unique baselines, some combinations occur more frequently than others, leading to an effective set of co-array weights for the untapered array, indicated by the coloured markers.

The result of applying conventional (classical) phased-array beamforming over all bearing and elevation angles, without any taper, is the polar diagrams in the top row of **Figure B.5**. In the absence of a taper, each element has equal weighting going into the beamformer, and for a single primary mode with a high signal-to-noise ratio and near planar wavefront, this closely resembles the fundamental beam pattern. Not surprisingly the beams are quite broad, given that the aperture is only about 4 to 8 wavelengths long across each of the two arms, and the side-lobes are very high, only about 5 dB below the main-lobe, which will greatly compromise the ability to detect multiple sources. The latter issue comes down to the effective weights for the difference co-array, which in this case are rather harsh (not smoothly varying) and do little to control side-lobe levels.
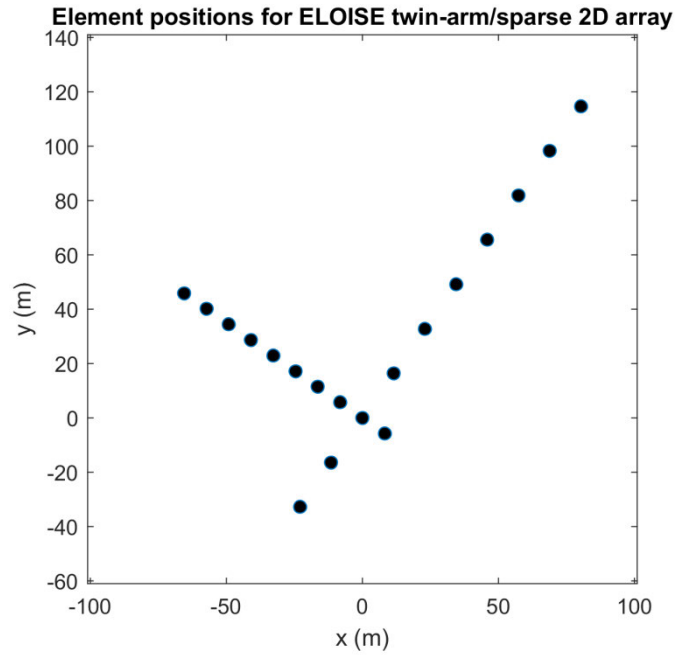
Figure B.2. Layout of antenna elements in the local horizontal plane for the Laverton ELOISE array. In this coordinate system, the x-axis is eastward and the y-axis is northward.
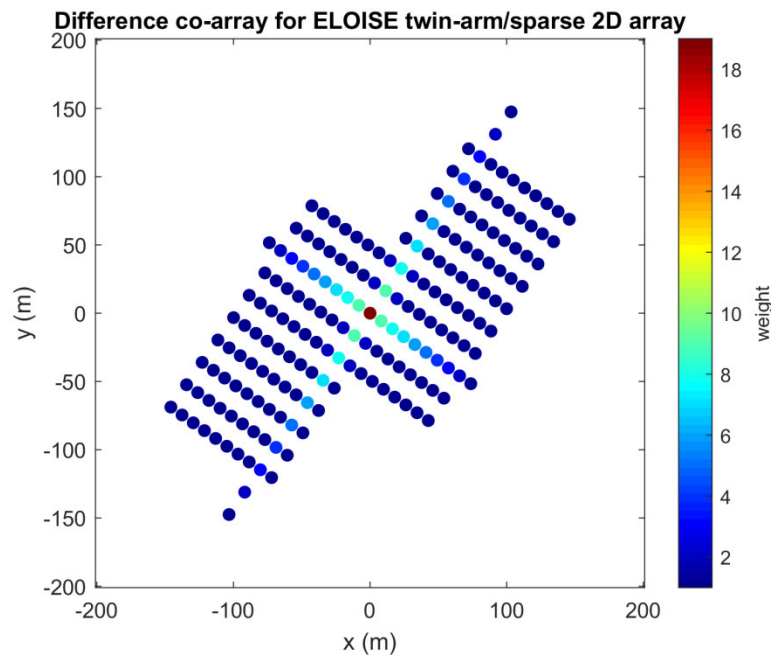
Figure B.3. Difference co-array for the ELOISE array layout in **Figure B.2**. The markers are colour-coded according to the relative occurrences of each inter-element baseline, which represent a set of co-array weights for the untapered array.
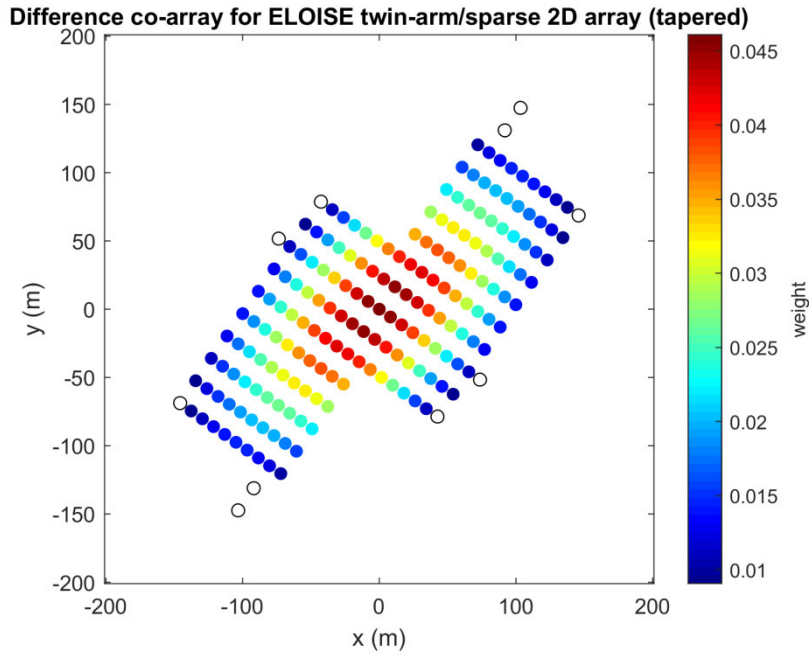
Figure B.4. Alternative co-array weighting scheme, designed to suppress side-lobes. This is based on a Dolph-Chebyshev window function with 30 dB side-lobe attenuation. Note that certain baselines are missing (i.e. effectively zeroed) due to the asymmetric shape of the array; this will degrade the ideal side-lobe performance.

There are two global maxima in the top left panel of **Figure B.5**: one peak, in the northeastern quadrant, represents the main-lobe or true signal direction (with an approximate bearing of 68 °T and elevation of 17°), while the other peak, in the south-eastern quadrant, is a grating lobe ambiguity, caused by spatial undersampling of the array. The two are distinguished using the known great circle path bearing (dashed radial line) and the predicted spread in equivalent elevation angles for the oblique ground range, as described in **Section 4.2.2.1**. An enlarged plot of the main-lobe is presented in the top right panel. Note that both the beamwidth and position of the grating lobe(s) are dependent on frequency, which in this case was 13.18 MHz.

Applying the standard ELOISE AoA estimation algorithm from **Section 4.2.2** gives the peak signal direction indicated by the black circle marker; this is computed as a power-weighted mean of AoA estimates from all the unmasked pixels in **Figure B.1**, each plotted as a black cross in the enlarged (top right) panel. Clearly the agreement between the image and AoA estimates is excellent, as is to be expected for a single propagation
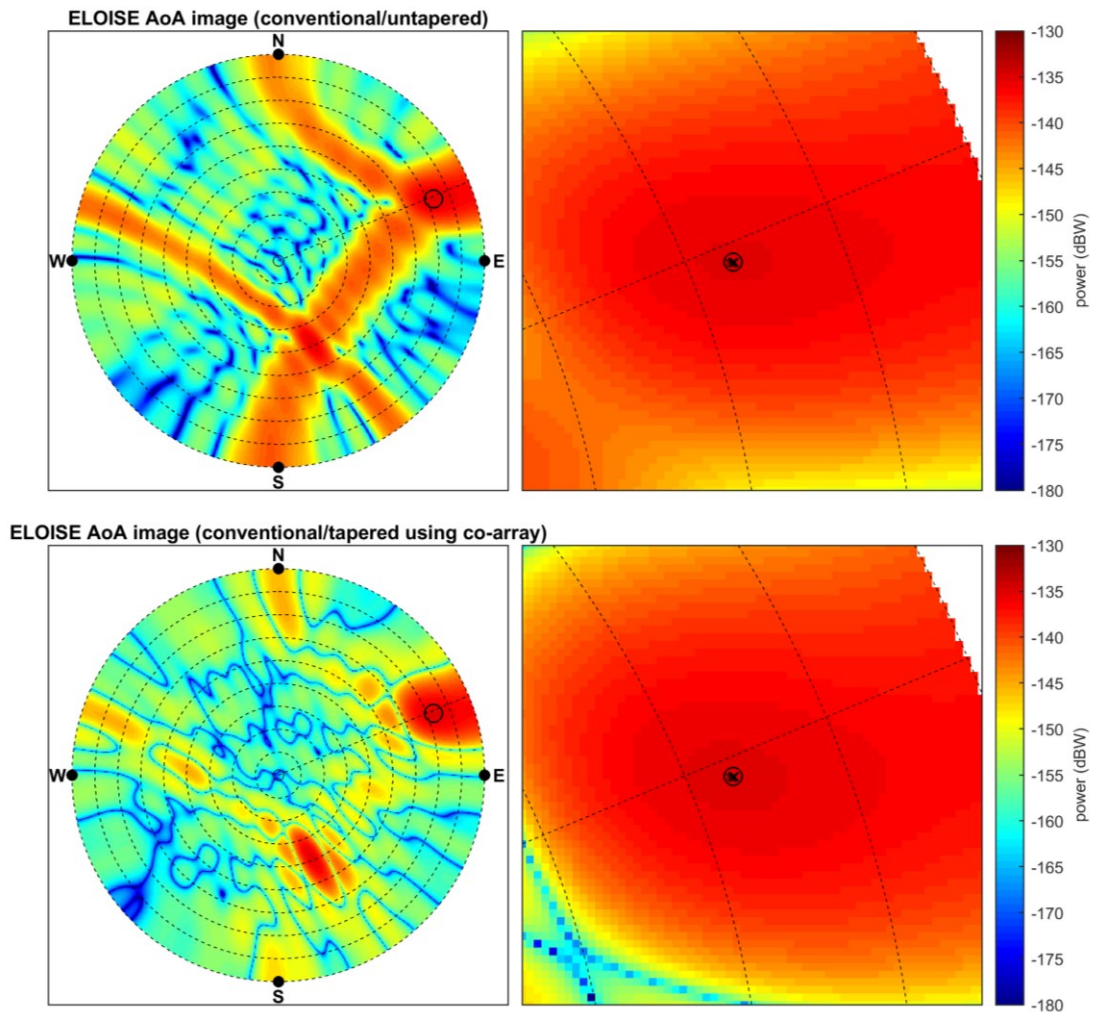
Figure B.5. Polar AoA images showing the results of conventional beamforming with two different weighting schemes. In the top row, the array elements are untapered, which corresponds to the effective co-array weights shown in **Figure B.3**. In the bottom row, an alternative set of co-array weights from **Figure B.4** were adopted to suppress side-lobes. The centre of each circular image represents the zenith, and the outer edge the horizon, with dashed lines of concentric circles marking each 10° elevation increment. The great circle path bearing for the Harts Range transmitter (67 °T) is shown by a further radial dashed line. Each image is enlarged on the right, to focus on the region surrounding the main-lobe peak, with the black cross and circle markers representing individual and averaged AoA estimates, respectively, derived as part of the standard ELOISE on-board processing.

mode. Therefore, in this case, there is no obvious benefit to be gained from the extra computational effort of beamforming, peak detection, and peak fitting, compared with just a planar fit to the phase angle measurements.

An alternative conventional beamformer, constructed in the co-array domain with the modified weighting scheme in **Figure B.4**, was implemented to produce the bottom row of polar diagrams in **Figure B.5**. For a sparse 2D array, this amounts to a form of aperture synthesis. Note that the weights themselves were not optimised in any rigorous way; they were merely selected to exemplify the benefits of using the difference co-array for beamforming on a sparse array. In this case, the highest side-lobe is now over 8 dB below the main-lobe, which by no means great, is nevertheless an improvement over the untapered array.

While there is little doubt as to the effectiveness of the ELOISE AoA estimation algorithm for a single mode resolved in delay and Doppler, imaging techniques have the potential to deliver value by providing an additional (spatial) dimension to support mode resolvability. Unfortunately, the size of the ELOISE array aperture and, thus, the widths of the conventional beams at typical oblique incidence elevation angles are poorly suited to resolving even F2 low- and high-ray modes (separated by ~14° in this example), let alone O and X geomagnetic components. It is for these kinds of problems that high-resolution adaptive processing techniques offer promise; leading candidates include the optimal Minimum Variance Distortionless Response (MVDR) beamformer [Capon, 1969], and subspace-based methods, such as the Multiple Signal Classification (MUSIC) algorithm [Schmidt, 1986]. Many studies have previously compared the relative merits of these different techniques (as summarised in Krim & Viberg [1996]) and these will not be addressed here.

The challenge with any such high-resolution technique is constructing a robust estimate of the signal-plus-noise covariance matrix from the outputs of each array element (receiver). In using the full delay and Doppler resolution of the CSF dwell, there are very few samples to work with and, as these are mostly a product of spectral leakage, they are not independent. For instance, in **Figure B.1**, there are just 16 unmasked pixels (samples) characterising the F2-low mode, so with 19 elements, an average over these samples cannot produce a full-rank covariance matrix estimate. Time-averaging, under the assumption of stationarity, can potentially be employed to overcome this, but to do

so requires the coherent integration time of the fixed-length CSF dwell to be reduced, sacrificing Doppler resolution.

Figure B.6 gives polar diagrams for the MVDR beamformer (top row) and MUSIC algorithm (bottom row), for the same data as in **Figure B.5**. Owing to the rank-deficient covariance matrix estimate, regularisation (diagonal loading) was necessary to gain meaningful results from the MVDR beamformer and, by its very nature, this raises the apparent white noise background. In the process of nulling the highly correlated side-lobe structure, it appears the peak power in the MVDR image has also been suppressed to some extent. Again, the peak position is consistent with the standard (phase-only) AoA approach. Results are similar for the MUSIC pseudo-spectrum, assuming just one signal source ($p = 1$).

Introducing now all three distinct F2 propagation modes from **Figure B.1** into the processing, the MVDR and MUSIC output is shown in **Figure B.7**. The MUSIC algorithm uses knowledge of the additional signal components ($p = 3$), which can also be deduced from the eigenspectrum of the covariance matrix estimate. Of course, this is a rather contrived example, as separation of the F2-low and F2-high modes is routinely possible based on the delay dimension alone. However, the results do demonstrate a modest ability to resolve signals in the spatial dimension, despite the fact that the high-ray O- and X- components (separated by only ~3° in elevation) remain indistinguishable in the image. It is interesting to also note a shift in the position of the F2-low peak output by MUSIC.

One avenue that has not been explored is the use of a Bayesian framework for adaptive beamforming. As described in Bell et al. [2000], this can take the form of a weighted sum of MVDR beamformers, incorporating a known *a priori* probability density function that characterises the uncertainty in the signal source direction. For the HF radar application, assembling such a function would need to take into account details of ionospheric variability, so long-term ionosonde observations would likely be well-suited to such a study. This is acknowledged as an area of future work.

In conclusion, preliminary results from imaging on the ELOISE array have revealed limitations in the array geometry that may be difficult to overcome for direction finding applications. Impressions are that the spatial domain offers very little in terms of extra resolvability of one-way oblique modes compared to the delay and Doppler dimensions

that are already used in the standard ELOISE on-board processing, and many of these imaging techniques may be less robust and more computationally expensive. Should further experimentation be possible, other sparse 2D array geometries, such as circular or elliptical arrays (e.g. Kozick & Kassam [1992]), and/or larger apertures may prove to be better candidates for imaging purposes.
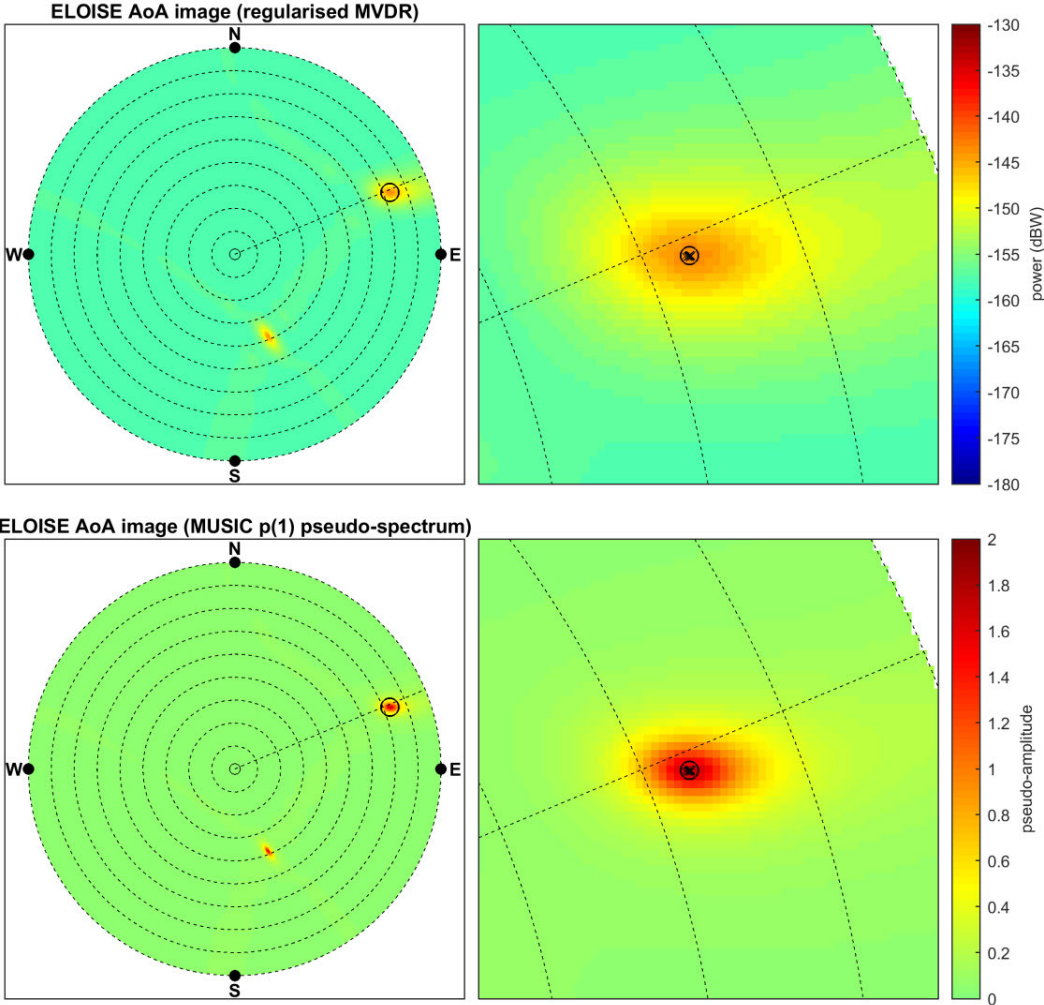
Figure B.6. Polar AoA images showing the results of the regularised MVDR beamformer (top row) and MUSIC algorithm (bottom row).
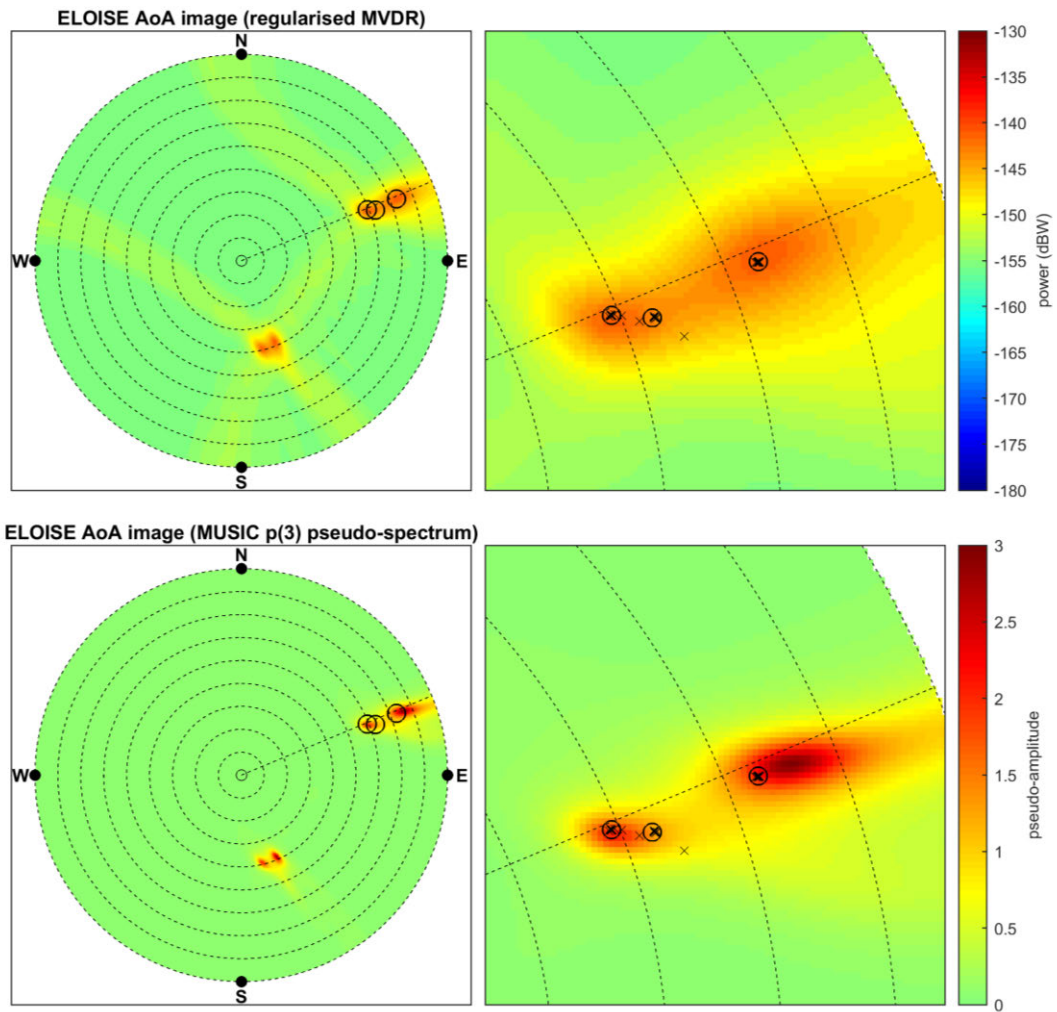
Figure B.7. As in **Figure B.6**, but with both F2 low- and high-ray modes combined. Standard AoA estimates for each of three modes, indicated by the black cross and circle markers, are resolved using the CSF delay and Doppler dimensions only.

Appendix C   Near-field extension to AoA fitting

For direct line-of-sight propagation from a near-field source, the wave vector varies over the array aperture and defines a spherical wavefront to first order. As such, the standard planar fitting approach for AoA estimation from **Section 4.2.2**, applicable to far-field (sky-wave) sources, is no longer appropriate. Assuming free space propagation from a source at $r_s = [x_s, y_s, z_s]^T$, with phase path P equal to the geometric path length, the spherical function of position vector $r = [x, y, z]^T$ can be written in the following form:

$$P^2 = \left(\frac{\lambda}{2\pi}\Phi\right)^2 = \|r - r_s\|^2 = (x - x_s)^2 + (y - y_s)^2 + (z - z_s)^2, \qquad \text{(C.1)}$$

where $\Phi = -(\varphi + \varphi_{\text{ref}})$ is the (unwrapped) total phase change along the path, for some unknown reference phase angle φ_{ref} common to all relative phase angles φ. The negative sign reflects the fact that the measured phase φ decreases (i.e. phase lags) as the path length from the source increases (cf. the negative sign in the $-i2\pi f T'_{g,n}$ exponent of equation (A.5) from **Appendix A**).

This can be rewritten in terms of vector dot products as

$$(\varphi + \varphi_{\text{ref}})^2 = \left(\frac{2\pi}{\lambda}\right)^2 (r \cdot r - 2r \cdot r_s + r_s \cdot r_s). \qquad \text{(C.2)}$$

At the origin (i.e. $r = 0$), equation (C.2) reduces to

$$\varphi_0 + \varphi_{\text{ref}} = -\frac{2\pi}{\lambda} r_s, \qquad \text{(C.3)}$$

where $r_s = \sqrt{x_s^2 + y_s^2 + z_s^2}$. If the source position $\boldsymbol{r}_s$ is known, and phase angle φ_0 estimated from a preliminary planar fit, then φ_{ref} can be determined.

Inserting equation (C.3) and the wave vector at the origin $\boldsymbol{k}_0 = \frac{2\pi}{\lambda}\hat{\boldsymbol{r}}_s$ (for unit vector $\hat{\boldsymbol{r}}_s \equiv \boldsymbol{r}_s/r_s$) back into equation (C.2) gives the following:

$$(\varphi + \varphi_{\text{ref}})^2 = \left(\frac{2\pi}{\lambda}\right)^2 r^2 + 2(\varphi_0 + \varphi_{\text{ref}})(\boldsymbol{k}_0 \cdot \boldsymbol{r}) + (\varphi_0 + \varphi_{\text{ref}})^2, \qquad \text{(C.4)}$$

or equivalently,

$$\boldsymbol{k}_0 \cdot \boldsymbol{r} = \frac{1}{2(\varphi_0 + \varphi_{\text{ref}})}\left(\varphi^2 - \varphi_0^2 + 2(\varphi - \varphi_0)\varphi_{\text{ref}} - \left(\frac{2\pi}{\lambda}\right)^2 r^2\right). \qquad \text{(C.5)}$$

Thus, for array element n,

$$\boldsymbol{k}_0 \cdot \boldsymbol{r}_n = \varphi_n^* - \varphi_0, \qquad \text{(C.6)}$$

where "flattened" phase angle $\varphi_n^* = \frac{1}{2(\varphi_0 + \varphi_{\text{ref}})}\left(\varphi_n^2 + \varphi_0^2 + 2\varphi_n\varphi_{\text{ref}} - \left(\frac{2\pi}{\lambda}\right)^2 r_n^2\right)$. Note that for far-field sources, as $r_s \to \infty$, $\varphi_{\text{ref}} \to -\infty$ and $\varphi_n^* \to \varphi_n$; i.e. equation (C.6) reverts to defining a standard planar wavefront.

The planar wavefront fit from **Section 4.2.2** can thus be generalised to a spherical wavefront fit by solving the least-squares problem in an iterative fashion with a "corrected" phase field across the array. At each step, the measured phase angles are effectively flattened into a plane, using the previous AoA estimate, to which the standard (planar) least-squares fit yields an updated AoA estimate. This process is repeated until the fitted parameters converge to within some predefined tolerance.

Appendix D JORN ionospheric profile parameterisation

The Jindalee Operational Radar Network (JORN) comprises three sky-wave over-the-horizon radars in central Australia. It relies on a Real-Time Ionospheric Model (RTIM) for waveform advice and coordinate registration of targets, which assimilates data from a dense network of vertical and oblique incidence sounders. The inputs to the RTIM are the automatically-scaled electron density profiles from the sounder network, parameterised as a sequence of multiple quasi-parabolic (QP) segments representing each 1-hop mid-point. The output is a set of spatial maps for each of the profile parameters, in a form suitable for either analytical or numerical ray tracing. Sporadic-E is handled separately based on a sector model of mirror reflectors. The D layer is omitted from the profile altogether, noting that non-deviative absorption effects can be modelled alongside the ray tracing if required (e.g. Pederick & Cervera [2014]).

The QP layer formulation and ray path integrals were originally developed by Croft & Hoogasian [1968] and later generalised to multi-segment QP (MQP) profiles by Hill [1979], Dyson & Bennett [1988], Chen et al. [1990], and Bennett et al. [1991]. The advantage of QP layers over parabolic layers is the fact that exact analytical ray equations exist for a spherical Earth geometry. The JORN parameterisation is a constrained form of the MQP profile, which includes inverted joining segments but no valleys; that is, the electron density is assumed to be monotonically increasing with height. It also imposes geophysical limits on the E, F1 and F2 layer parameters and the amount of overlap between layers to ensure a valid and realistic profile shape.

The JORN MQP profile building rules [Gardiner-Garden et al., 2018], described in further detail below, reduce the number of independent profile parameters to just ten, as listed in **Table D.1**. These are the critical frequencies, peak heights (radii), and semi-thicknesses for each of the three QP layers (E, F1 and F2), along with a single index defining the start point of the F2 join. The F1 join has its start point fixed to the peak of

the E layer. The joining segments are all calculated by imposing continuity conditions on the electron density and its height gradient, including at the base of the E layer. Under most ionospheric conditions, this approach is able to provide good fits to the key features observed in ionogram traces, with only a modest number of degrees of freedom.

The fundamental building block for the MQP profile is the QP layer [Croft & Hoogasian, 1968], for which the square of the plasma frequency f_p closely approximates a parabola at heights $h_0 \leq h \leq h_m$. It can be expressed mathematically as

$$f_p(h)^2 = f_c^2\left(1 - \zeta_{qp}^2\right), \text{ where } \zeta_{qp} \equiv \frac{r_0}{r(h)}\left(\frac{h - h_m}{y_m}\right), \ |\zeta_{qp}| \leq 1. \tag{D.1}$$

Table D.1. List of JORN quasi-parabolic (QP) profile parameters, specifying the E, F1 and F2 layers.

Parameter	Description
f_{cE}	E layer critical frequency (MHz)
h_{mE}	E layer peak height (km)
y_{mE}	E layer semi-thickness (km)
f_{cF1}	F1 layer critical frequency (MHz)
h_{mF1}	F1 layer peak height (km)
y_{mF1}	F1 layer semi-thickness (km)
f_{cF2}	F2 layer critical frequency (MHz)
h_{mF2}	F2 layer peak height (km)
y_{mF2}	F2 layer semi-thickness (km)
q	F2 join index (scaled from 0 to 1)

Here f_c is the peak (critical) frequency, h_m is the peak height, and y_m is the semi-thickness of the layer. The factor of $(r_0/r(h))^2$, which distinguishes the QP layer from the true parabolic layer, is the squared ratio of the Earth-centred radius at the bottom of the layer (i.e. $r_0 = h_0 + R_E$, where $h_0 = h_m - y_m$ and R_E is the spherical Earth radius) and the radius at the current height (i.e. $r(h) = h + R_E$).

Note that electron density N_e (in m^{-3}) is directly proportional to the square of the (electron) plasma frequency f_p (in MHz):

$$N_e(h) = f_p(h)^2 \left(\frac{4\pi^2 \varepsilon_0 m_e}{e^2} \right) = \frac{f_p(h)^2}{80.6 \text{ m}^3\text{s}^{-2}}, \tag{D.2}$$

where ε_0 is the permittivity of free space, m_e is the electron rest mass, and e is the elementary charge of an electron.

Adopting the notation in Dyson & Bennett [1988], the QP layer from equation (D.1) can also be represented as

$$f_p(h)^2 = a - b\left(1 - \frac{r_m}{r(h)}\right)^2, \tag{D.3}$$

where $a = f_c^2$ and $b = f_c^2(r_0/y_m)^2$ are coefficients describing the vertex and curvature of the layer, respectively. This is the form that will be used in the JORN profile description that follows.

The JORN MQP profile consists of three QP layer segments, representing E, F1 and F2, and three QP join segments (i.e. inverted QP layers), smoothly connected to give a continuously differentiable function from the base of the E layer to the peak of the F2 layer. The formal definition in equation (D.4) below is illustrated for a sample set of parameters in **Figure D.1**.

$$
f_p(r)^2 = \begin{cases}
a_{j0} + b_{j0}\left(1 - \dfrac{r_{\mathrm{mj0}}}{r}\right)^2, & \text{for } r_{i0} \le r < r_{j0} \text{ (E join)}, \\[2ex]
a_E - b_E\left(1 - \dfrac{r_{\mathrm{mE}}}{r}\right)^2, & \text{for } r_{j0} \le r < r_{i1} \text{ (E layer)}, \\[2ex]
a_{j1} + b_{j1}\left(1 - \dfrac{r_{\mathrm{mj1}}}{r}\right)^2, & \text{for } r_{i1} \le r < r_{j1} \text{ (F1 join)}, \\[2ex]
a_{F1} - b_{F1}\left(1 - \dfrac{r_{\mathrm{mF1}}}{r}\right)^2, & \text{for } r_{j1} \le r < r_{i2} \text{ (F1 layer)}, \\[2ex]
a_{j2} + b_{j2}\left(1 - \dfrac{r_{\mathrm{mj2}}}{r}\right)^2, & \text{for } r_{i2} \le r < r_{j2} \text{ (F2 join)}, \\[2ex]
a_{F2} - b_{F2}\left(1 - \dfrac{r_{\mathrm{mF2}}}{r}\right)^2, & \text{for } r_{j2} \le r \le r_{\mathrm{mF2}} \text{ (F2 layer)}.
\end{cases}
\tag{D.4}
$$
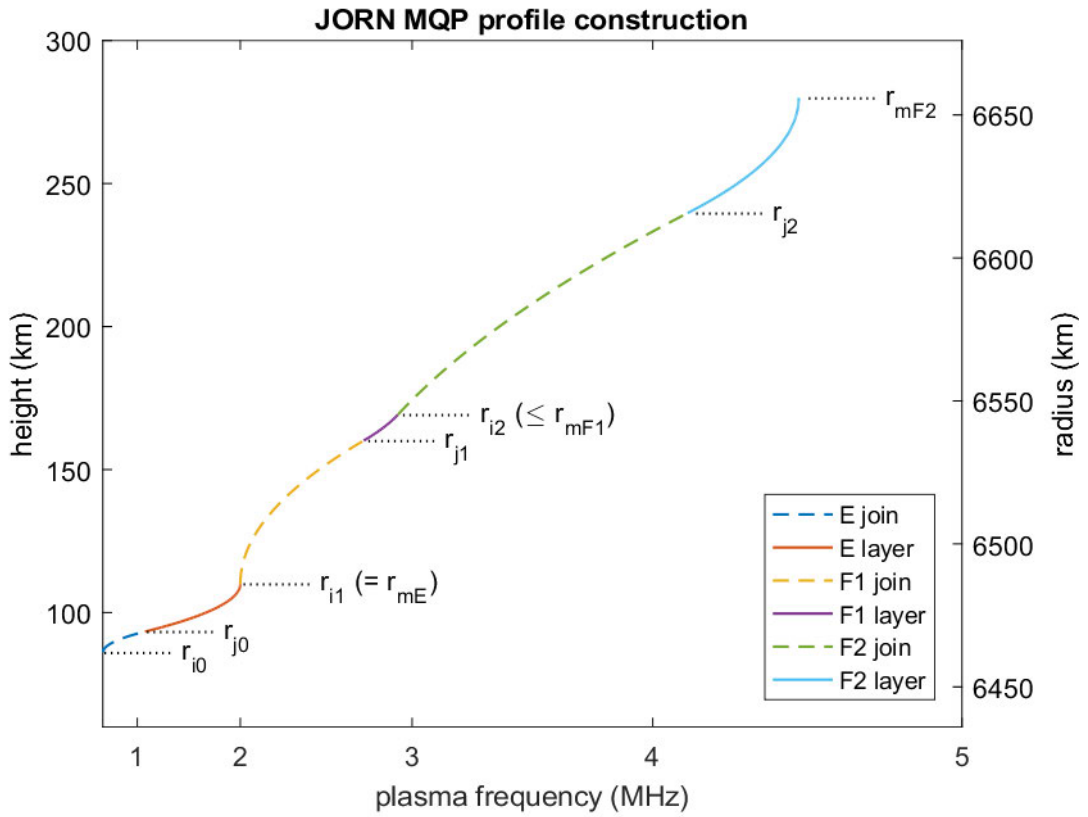
Figure D.1. A sample JORN MQP profile, with segment joining points annotated as Earth-centred radii.

The three physical layer segments are defined by their respective critical frequencies, peak heights, and semi-thicknesses, which map to the QP coefficients for E, F1 and F2 layers as follows:

$$a_E = f_{cE}^2, \ a_{F1} = f_{cF1}^2, \ a_{F2} = f_{cF2}^2$$

(vertex coefficients for E, F1 and F2 layers), (D.5)

$$b_E = f_{cE}^2 \frac{r_{0E}}{y_{mE}}, \ b_{F1} = f_{cF1}^2 \frac{r_{0F1}}{y_{mF1}}, \ b_{F2} = f_{cF2}^2 \frac{r_{0F2}}{y_{mF2}}$$

(curvature coefficients of E, F1 and F2 layers), (D.6)

$$r_{mE} = h_{mE} + R_E, \ r_{mF1} = h_{mF1} + R_E, \ r_{mF2} = h_{mF2} + R_E$$

(radii of E, F1 and F2 layers). (D.7)

By matching adjacent layer and join segments, and their gradients, at each of the join points, the equations for the QP join coefficients can easily be derived. In the majority of cases, a valid solution for each join segment exists, although in certain cases adjustment of the E, F1 or F2 layer parameters may be necessary. This can occur if the E/F1/F2 critical frequencies and heights do not form an increasing sequence, if there is too much overlap to construct a join, or if a join would otherwise be excessively large. A series of checks and corrections ensures that a valid profile will always be generated, and that smoothly varying input parameters result in smoothly varying output, to a level of precision suitable for numerical optimisation.

The E join starts at the vertex of the inverse QP segment, where $f_p = 0$ and $r = r_{mj0}$. The following QP coefficients are defined:

$$a_{j0} = 0 \ \text{(vertex coefficient of E join)}, \quad\quad \text{(D.8)}$$

$$r_{mj0} = r_{mE} - 1.2 y_{mE} \ \text{(radius of E join vertex)}, \quad\quad \text{(D.9)}$$

$$r_{i0} = r_{mj0} \ \text{(radius of lower E join point).} \quad\quad \text{(D.10)}$$

The radius of the upper E join point r_{j0} and the curvature coefficient of the E join b_{j0} can then be derived by applying the continuity conditions stated above (cf. equation (7) from Dyson & Bennett [1988]):

$$r_{j0} = \frac{r_{\mathrm{mE}} b_E \left(\frac{r_{\mathrm{mE}}}{r_{\mathrm{mj0}}} - 1 \right)}{\left(a_E - a_{j0} + b_E \left(\frac{r_{\mathrm{mE}}}{r_{\mathrm{mj0}}} - 1 \right) \right)} \qquad (D.11)$$

(radius of upper E join point),

$$b_{j0} = -\frac{r_{\mathrm{mE}} b_E \left(1 - \frac{r_{\mathrm{mE}}}{r_{j0}} \right)}{r_{\mathrm{mj0}} \left(1 - \frac{r_{\mathrm{mj0}}}{r_{j0}} \right)} \qquad (D.12)$$

(curvature coefficient of E join).

Similar equations apply to the F1 join, which starts at the E layer peak with $f_p = f_{\mathrm{cE}}$ and $r = r_{\mathrm{mj1}} = r_{\mathrm{mE}}$:

$$a_{j1} = a_E \quad \text{(vertex coefficient of F1 join),} \qquad (D.13)$$

$$r_{\mathrm{mj1}} = r_{\mathrm{mE}} \quad \text{(radius of F1 join vertex),} \qquad (D.14)$$

$$r_{i1} = r_{\mathrm{mj1}} \quad \text{(radius of lower F1 join point),} \qquad (D.15)$$

$$r_{j1} = r_{\mathrm{mF1}} b_{F1} \left(\frac{r_{\mathrm{mF1}}}{r_{\mathrm{mj1}}} - 1 \right) / \left(a_{F1} - a_{j1} + b_{F1} \left(\frac{r_{\mathrm{mF1}}}{r_{\mathrm{mj1}}} - 1 \right) \right) \qquad (D.16)$$

(radius of upper F1 join point),

$$b_{j1} = -r_{\mathrm{mF1}} b_{F1} \left(1 - \frac{r_{\mathrm{mF1}}}{r_{j1}} \right) / r_{\mathrm{mj1}} \left(1 - \frac{r_{\mathrm{mj1}}}{r_{j1}} \right) \qquad (D.17)$$

(curvature coefficient of F1 join).

The final F2 join differs in construction from the E and F1 joins, in that the segment may start below the F1 layer peak. As such, the F1 segment may only consist of an

inflection in the profile, rather than a distinct peak. The arbitrary lower join point is controlled by the q index, which is normalised to vary from 0 (corresponding to a larger join segment starting at $r = r_{i2,\text{min}}$) to 1 (corresponding to a smaller join segment starting at $r = r_{\text{mF1}}$). Lower values of q will tend to produce weaker cusps in the synthetic ionogram traces, especially when there is significant overlap between the F1 and F2 layers. The starting point of the F2 join is defined by

$$r_{i2} = r_{i2,\text{min}} + q\left(r_{\text{mF1}} - r_{i2,\text{min}}\right)$$

(radius of lower F2 join point).

(D.18)

Note that $r_{i2,\text{min}}$ is the lowest possible join point to which an inverse QP function can be fitted; this corresponds to a near-linear F2 join segment. The maximum possible join point $r_{i2,\text{max}}$ may be less than r_{mF1} if the F1 peak is engulfed by the F2 layer (e.g. during the night), in which case the maximum valid q will be less than 1.

With r_{i2} specified, the unique F2 join is then the solution of four simultaneous equations in a_{j2}, b_{j2}, r_{mj2} and r_{j2} (from continuity conditions evaluated at r_{i2}). The first step is to solve for vertex radius r_{mj2}, which is given by the following equation after variable elimination:

$$r_{\text{mj2}} = \frac{\text{B} b_{F2} r_{\text{mF2}}^2}{\text{A} b_{F1} r_{\text{mF1}}}$$

(radius of F2 join vertex),

(D.19)

where

$$A = a_{F1} - a_{F2} - b_{F1} + b_{F2} + (a_{F1} - a_{F2} - b_{F1})b_{F2}\gamma/b_{F1}\alpha$$
$$+ (a_{F2} - a_{F1} + 2b_{F1} - b_{F2})\beta - b_{F1}\beta^2 + b_{F2}\gamma^2,$$

(D.20)

$$B = a_{F1} - a_{F2} - b_{F1} + b_{F1}(2\alpha - 2\alpha\beta + \beta^2) - (b_{F1}\alpha\beta)^2/b_{F2}$$
$$+ 2(b_{F1}\alpha)^2\beta/b_{F2} - (b_{F1}\alpha)^2/b_{F2},$$

(D.21)

$$\alpha = \frac{r_{\text{mF1}}}{r_{\text{mF2}}}, \tag{D.22}$$

$$\beta = \frac{r_{\text{mF1}}}{r_{i2}}, \tag{D.23}$$

$$\gamma = \frac{r_{\text{mF2}}}{r_{i2}}. \tag{D.24}$$

The upper F2 join point r_{j2} and coefficients a_{j2} and b_{j2} can then be calculated:

$$r_{j2} = \frac{r_{\text{mF2}} b_{F2} \left(\frac{r_{\text{mF2}}}{r_{\text{mj2}}} - 1 \right)}{\left(a_{F2} - a_{j2} + b_{F2} \left(\frac{r_{\text{mF2}}}{r_{\text{mj2}}} - 1 \right) \right)} \tag{D.25}$$

(radius of upper F2 join point),

$$b_{j2} = -\frac{r_{\text{mF1}} b_{F1} \left(1 - \frac{r_{\text{mF1}}}{r_{i2}} \right)}{r_{\text{mj2}} \left(1 - \frac{r_{\text{mj2}}}{r_{i2}} \right)} \tag{D.26}$$

(curvature coefficient of F2 join),

$$a_{j2} = a_{F1} - b_{F1} \left(1 - \frac{r_{\text{mF1}}}{r_{i2}} \right)^2 - b_{j2} \left(1 - \frac{r_{\text{mj2}}}{r_{i2}} \right)^2 \tag{D.27}$$

(vertex coefficient of F2 join).

Unlike the vertex coefficients for the other layers and joins, a_{j2} may be negative.

It is important to note that, being a monotonically increasing profile with a fixed number of layers, the F1 segment persists through the night in the JORN parameterisation, even though a separate F1 layer is unlikely to exist in reality. However, the JORN profile building rules are such that the night-time (pseudo) F1 layer can be increased in height to maximise F1/F2 overlap and minimise the extent of the F2 join; thus, it ultimately has very little effect on the synthetic ionogram traces.

Appendix E Ionospheric response to atmospheric gravity waves

Travelling ionospheric disturbances (TIDs) are most commonly thought to be the ionospheric response to atmospheric gravity waves (AGWs), which frequently propagate up to the thermosphere from sources at lower altitudes. In a series of seminal papers, Hooke [1968; 1970] applied a linearised perturbation approach, with careful consideration given to each of the ion production, chemical loss and transport effects, to derive the governing physical equations for the dynamics and propagation of TIDs. In the F2 region, it is the collisional interactions of ions with neutral gas motions that dominate, leading to a simplified form of these equations which is outlined here, along with specific model implementation details relating to the ELOISE analysis. This has been adapted from earlier work by Cervera & Harris [2014], using code provided by these authors as a starting point. A full discussion of the assumptions and limitations can be found in Section 7 of Hooke [1968].

The fundamental dispersion relation for acoustic-gravity waves, relating the (intrinsic) angular wave frequency ω to the horizontal and vertical wave numbers k_h and k_z, is given by [Hines, 1960, eq. (14); Kelley, 2009, eq. (6.3)]

$$\omega^4 - \omega^2 c_a^2 (k_h^2 + k_z^2) + (\gamma - 1)g^2 k_h^2 + i\gamma g \omega^2 k_z = 0, \qquad (E.1)$$

where c_a is the speed of sound, γ is the ratio of specific heats, and g is the acceleration due to gravity. The assumptions involved are described by Kelley [2009, ch. 6.2].

In terms of the neutral atmospheric scale height $H \equiv c_a^2/\gamma g$ [Hines, 1960, eq. (12)] and Brunt-Väisälä (buoyancy) frequency $\omega_B = \sqrt{\frac{\gamma-1}{\gamma}\frac{g}{H}}$, the dispersion relation can be rewritten as

$$\frac{\omega^2}{c_a^2} - (k_h^2 + k_z^2) + \frac{\omega_B^2 k_h^2}{\omega^2} + i\frac{k_z}{H} = 0. \tag{E.2}$$

In practice, the scale height at the path midpoint is calculated based on the standard formula $H = k_B T/\overline{m}g$, in which k_B is the Boltzmann constant, using temperature T and mean molecular mass $\overline{m}$ derived from the NRLMSISE-00 reference model [Picone et al., 2002]. The ratio of specific heats $\gamma \equiv c_p/c_v$ (i.e. between the specific heat of dry air at constant pressure c_p and that at constant volume c_v) is interpolated from the look-up table in Hilsenrath et al. [1955] and is approximately equal to 1.4. The speed of sound and Brunt-Väisälä frequency follow from the definitions above.

On the assumption that k_h is real, this implies k_z is complex; i.e. $k_z = k_{zr} + ik_{zi}$. Inserting this into equation (E.2) and solving for the real and imaginary components gives [Hines, 1960, eq. (21)]

$$k_{zr} = -\sqrt{k_h^2\left(\frac{\omega_B^2}{\omega^2} - 1\right) + \frac{\omega^2 - \omega_a^2}{c_a^2}}, \tag{E.3}$$

and

$$k_{zi} = 1/2H, \tag{E.4}$$

where the acoustic cut-off frequency $\omega_a \equiv c_a/2H$. Note that the value of k_{zi} above is purely a growth factor and allows for no dissipation (e.g. due to molecular viscosity, thermal diffusion and ion drag); k_{zi} can be reduced to partially compensate for the effects of dissipation if required. The following modified expression was applied to the ELOISE analysis, with the dissipation factor $0 \leq \delta \leq 1$ typically assigned a value of 0.3 based on previous work by Dr Manuel Cervera:

$$k_{zi} = (1 - \delta)/2H. \tag{E.5}$$

For internal atmospheric gravity waves with $\omega < \omega_B < \omega_a$, equation (E.3) reduces to

$$k_{zr} \approx -k_h \sqrt{\frac{\omega_B^2}{\omega^2} - 1}. \qquad \text{(E.6)}$$

The wave solution in 3D space and time, which governs the neutral atmospheric motions (wind velocities), takes the following form [Hooke, 1970, eq. (1)]:

$$\boldsymbol{u}(\boldsymbol{r}, t) = u(\boldsymbol{r}, t)\hat{\boldsymbol{u}} = U_0 \exp\left(k_{zi}(z - z_0)\right) \exp\left(i(\omega t - \boldsymbol{k_r} \cdot \boldsymbol{r} + \varphi_0)\right)\hat{\boldsymbol{u}}, \qquad \text{(E.7)}$$

where $U_0 = U(z_0)$ is the magnitude of the wind variation at reference height z_0 (taken to be 250 km in the ELOISE analysis), and $\hat{\boldsymbol{u}}$ is a constant unit vector along the axis of motion. In this local East-North-Up (ENU) Cartesian coordinate system, $\boldsymbol{r} = [x, y, z]^T$ is the position vector, t is the time, and $\boldsymbol{k_r} = [k_x, k_y, k_{zr}]^T$ is the real 3D wave vector (separating out k_{zi}). For a given wave azimuth θ, measured clockwise from True North, and forward tilt angle $\psi = \text{atan}(-k_{zr}/k_h) \approx \text{atan}\left(\sqrt{\omega_B^2/\omega^2 - 1}\right)$, measured between the direction of phase propagation and the local horizontal, $k_x = k_h \sin\theta$, $k_y = k_h \cos\theta$ and $\hat{\boldsymbol{u}} = [\sin\theta\chi\sin\psi, \cos\theta\chi\sin\psi, \cos\psi]^T$. By definition, $\hat{\boldsymbol{u}} \perp \boldsymbol{k_r}$ for such a transverse wave.

At this point, it is worth noting that the curvature of the Earth and ionosphere can be included to first order by letting x and y represent great circle distances to the East and North, respectively, at reference height z_0 above a spherical Earth.

Travelling ionospheric disturbances are created primarily by ion motion under the influence of ion-neutral collisions and the geomagnetic field (Lorentz forcing). On this basis, the ion drift velocities due to the neutral winds are given by [MacLeod, 1966, eq. (5); Hooke, 1970, eq. (2)]

$$\boldsymbol{u_i}(\boldsymbol{r}, t) = \frac{u(\boldsymbol{r}, t)}{1 + \rho_{in}^2}\left(\rho_{in}^2\hat{\boldsymbol{u}} + \rho_{in}\hat{\boldsymbol{u}} \times \hat{\boldsymbol{b}} + (\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}})\hat{\boldsymbol{b}}\right), \qquad \text{(E.8)}$$

where $\rho_{in}(z) \equiv \nu_{in}/\omega_i$ is the normalised ion-neutral collision frequency (i.e. the ratio of the collision frequency ν_{in} to the ion gyrofrequency ω_i), and $\hat{\boldsymbol{b}}$ is a unit vector in the direction of the geomagnetic field. Note that $u(\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}}) \neq \boldsymbol{u} \cdot \hat{\boldsymbol{b}}$ in general, as $\boldsymbol{u}$ includes a

complex phase component that should not be conjugated as part of the inner product; to make this explicit, the notation above deviates slightly from Hooke [1970] and only includes the real $\hat{\boldsymbol{u}}$ in the inner product.

The ion-neutral collision frequency ν_{in} can be calculated using the formula from Davies et al. [1997, eq. (9)], modelling the collisions between an ionosphere comprising NO^+, and a neutral atmosphere comprising N_2, O_2 and O, with densities given by NRLMSISE-00. The NO^+ gyrofrequency $\omega_i = q_i B_0 / m_i$ is calculated using ion charge $q_i = 1.602 \times 10^{-19}$ C, and ion mass $m_i = 30$ amu $= 4.98 \times 10^{-26}$ kg. The magnitude of the geomagnetic field B_0 is derived from the International Geomagnetic Reference Field (IGRF) with 2015 coefficients [Thébault et al., 2015].

The value of ρ_{in} decays exponentially with height, being of the order of 1–10 in the E region [MacLeod, 1966], but close to zero in the F region. At lower altitudes, the first term on the right-hand side of equation (E.8) dominates, resulting in $\boldsymbol{u}_i \approx \boldsymbol{u}$; i.e. the ion and neutral wind velocities are approximately the same. At higher altitudes, the third term dominates, and ion motion is restricted to the direction of the geomagnetic field.

Divergence of the ion motion leads to electron density perturbations, on the assumption that the electrons move in tandem with the positive ions due to ambipolar diffusion [Francey, 1963]. Given a background (unperturbed) electron density distribution $N_{e0}(\boldsymbol{r})$, which may be as simple as a spherically-symmetric profile, the perturbations can be represented to first order by [Hooke, 1970, eq. (4)]

$$N_e'(\boldsymbol{r}, t) = i\omega^{-1}\boldsymbol{\nabla} \cdot (N_{e0}\boldsymbol{u}_i) = i\omega^{-1}(\boldsymbol{u}_i \cdot \boldsymbol{\nabla}N_{e0} + N_{e0}\boldsymbol{\nabla} \cdot \boldsymbol{u}_i). \qquad \text{(E.9)}$$

This is just a statement of the continuity equation, assuming production and loss terms are in equilibrium. The gradient of N_{e0}, which appears in the first term on the right-hand side of equation (E.9), is dominated by its vertical component. In the absence of strong horizontal gradients (e.g. around the dawn/dusk terminators), it is reasonable to make the following approximation:

$$\boldsymbol{u}_i \cdot \boldsymbol{\nabla}N_{e0} \approx u_{iz}\frac{\partial N_{e0}}{\partial z}. \qquad \text{(E.10)}$$

The divergence of $\boldsymbol{u}_i$, which appears in the second term on the right-hand side of equation (E.9), can be expanded out using the product rule:

$$\boldsymbol{\nabla} \cdot \boldsymbol{u}_i = \boldsymbol{\nabla} u \cdot \left(\frac{\boldsymbol{u}_i}{u}\right) + u\boldsymbol{\nabla} \cdot \left(\frac{\boldsymbol{u}_i}{u}\right), \tag{E.11}$$

where

$$\boldsymbol{\nabla} u = \left[-ik_x u, -ik_y u, (k_{zi} - ik_{zr})u\right] = [0,0, k_{zi}u] - i\boldsymbol{k}_r u, \tag{E.12}$$

and

$$
\begin{aligned}
\boldsymbol{\nabla} \cdot \left(\frac{\boldsymbol{u}_i}{u}\right) &= \frac{\partial}{\partial z}\left(\frac{u_{iz}}{u}\right) \\
&= \frac{\partial}{\partial z}\left(\frac{1}{1+\rho_{in}^2}\left(\rho_{in}^2 \hat{u}_z + \rho_{in}(\hat{\boldsymbol{u}} \times \hat{\boldsymbol{b}})_z + (\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}})\hat{b}_z\right)\right) \\
&= \frac{1}{1+\rho_{in}^2}\frac{\partial \rho_{in}}{\partial z}\left(-\frac{2\rho_{in}u_{iz}}{u} + \left(2\rho_{in}\hat{u}_z + \hat{u}_x\hat{b}_y - \hat{u}_y\hat{b}_x\right)\right).
\end{aligned}
\tag{E.13}
$$

Inserting equations (E.10)–(E.13) back into equation (E.9) gives

$$N_e'(\boldsymbol{r}, t) \approx i\omega^{-1}\left(u_{iz}\frac{\partial N_{e0}}{\partial z} + N_{e0}\boldsymbol{\nabla} \cdot \boldsymbol{u}_i\right), \tag{E.14}$$

where

$$
\begin{aligned}
\boldsymbol{\nabla} \cdot \boldsymbol{u}_i &= k_{zi}u_{iz} - i\boldsymbol{k}_r \cdot \boldsymbol{u}_i \\
&+ \frac{1}{1+\rho_{in}^2}\frac{\partial \rho_{in}}{\partial z}\left(-2\rho_{in}u_{iz} + u\left(2\rho_{in}\hat{u}_z + \hat{u}_x\hat{b}_y - \hat{u}_y\hat{b}_x\right)\right).
\end{aligned}
\tag{E.15}
$$

In short, equation (E.14) expresses the ionospheric electron density perturbation in terms of the AGW frequency and wave vector, along with the background geomagnetic field, ion-neutral collision frequency, and electron density distribution. In the F region,

the normalised collision frequency $\rho_{in} \ll 1$. Setting both $\rho_{in} \approx 0$ and $\frac{\partial \rho_{in}}{\partial z} \approx 0$ reduces equation (E.8) to

$$\boldsymbol{u}_i(\boldsymbol{r}, t) \approx u(\boldsymbol{r}, t)(\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}})\hat{\boldsymbol{b}}, \tag{E.16}$$

and equation (E.14) to the standard approximation presented by Hooke [1968, eq. (50)]:

$$
\begin{aligned}
N'_e(\boldsymbol{r}, t) &\approx i\omega^{-1}\left(u_{iz}\frac{\partial N_{e0}}{\partial z} + N_{e0}\left(k_{zi}u_{iz} - iu(\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}})(\boldsymbol{k}_r \cdot \hat{\boldsymbol{b}})\right)\right) \\
&= i\omega^{-1}u_b\left(\sin I\left(\frac{\partial N_{e0}}{\partial z} + N_{e0}k_{zi}\right) - iN_{e0}k_{br}\right),
\end{aligned}
\tag{E.17}
$$

where $u_b = u(\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}})$ and $k_{br} = \boldsymbol{k}_r \cdot \hat{\boldsymbol{b}}$ are the components of $\boldsymbol{u}$ and $\boldsymbol{k}_r$, respectively, in the direction of the geomagnetic field, and $u_{iz} \approx u(\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}})\hat{b}_z = u_b\sin I$ from equation (E.16). Such a form is well suited to TID model simulations. Here I is the geomagnetic inclination (dip angle), defined in the same sense as Hooke [1968]; that is, positive in the Southern Hemisphere (contrary to the usual convention). Note that equation (E.16) implies all ion motion is constrained to the geomagnetic field lines under this approximation, unlike the more general equation (E.8). The factor of ω^{-1} means that the ionospheric response to larger-scale AGWs is greater than smaller-scale AGWs.

Naturally it is the real part of $N'_e(\boldsymbol{r}, t)$ in equations (E.14) and (E.17) that is of practical use and which has been used in this thesis to produce synthetic "observables" via ray tracing through realistic ionospheric perturbations. Full 3D magneto-ionic ray tracing is carried out in WGS-84 coordinates, using the PHaRLAP toolbox developed by Dr Manuel Cervera [Cervera & Harris, 2014; Cervera, 2019]. Rays are fired off from the transmit site through a grid of triangulated vertices on a unit sphere, to give a uniform distribution in solid angle that spans all feasible path directions. Total absorption loss through the D and E regions is estimated using the method of George & Bradley [1974], while deviative absorption above the E layer is calculated by integrating the imaginary component of the complex wave number along the ray path [Davies, 1990, eq. (7.14); Pederick & Cervera, 2014; Zawdie et al., 2017]. The latter makes use of an empirical expression for the electron-neutral collision frequency as a function of height [Coleman, 1997].

Those ray triplets found to straddle the receive site are linearly interpolated from the Delaunay triangulation [Press et al., 2007] to yield the synthetic observables. Focusing/defocusing losses were calculated from the PHaRLAP outputs using the effective distance method outlined in Davies [1990, ch. 7.3], but with ray flux tubes of triangular cross-section; see Nickisch [1988] and Västberg & Lundborg [1996] for more details. Direct comparisons of path losses estimated by such methods against oblique HF propagation measurements have been carried out by Netherway & Gardiner-Garden [2015].

Note that in the computational form given by Hooke [1968, eq. (51)], the $\sin I$ term is pulled out of the amplitude factor, which in effect adds an extra 180° phase shift in the Northern Hemisphere. Hooke appears to cancel this out by adopting a strictly positive value for $U_{b0} = U_b(z_0) = U_0 \left(\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{b}}(z_0) \right)$, which would otherwise be negative for most latitudes in the Northern Hemisphere. However, this approach becomes problematic more generally, including in the Southern Hemisphere, so to avoid confusion and preserve the phase relationship between the AGW and TID, equations (E.14) and (E.17) are chosen as the preferred form for calculating the electron density perturbation.

Under the approximation in equation (E.17), the AGW/TID phase offset becomes

$$\varphi = \frac{\pi}{2} - \operatorname{atan2}\left(u_b N_{e0} k_{br}, \; u_b \sin I \left(\frac{\partial N_{e0}}{\partial z} + N_{e0} k_{zi} \right) \right), \tag{E.18}$$

where $\operatorname{atan2}(y, x)$ represents the four-quadrant arctangent. This varies by up to ±90° in the vertical dimension [Hooke, 1970].

The forward tilt angle of the TID phase-front, which is not the same as the corresponding AGW tilt angle $\psi = \operatorname{atan}(-k_{zr}/k_h)$ due to the second height-varying term in equation (E.18), is then given by [Morgan & Calderón, 1978; Ballard, 1981]

$$\psi_{TID} = \operatorname{atan}\left(\frac{-k_{zr,TID}}{k_h} \right), \tag{E.19}$$

where

$$k_{zr,TID} = k_{zr} - \frac{d\varphi}{dz}. \tag{E.20}$$

Given a discretely sampled electron density field, the derivative of the phase offset in equation (E.20) is estimated as a centred finite difference at each height.

For a typical 20 min MSTID period at mid-latitudes, both AGW and TID tilt angles are of the order of 45° at F2 heights [Hines, 1960; Davies & Jones, 1971; Clarke, 1972, p. 267], while for longer period LSTIDs this can increase up to around 80°; that is, such that the phase propagation direction is more vertical and phase-fronts more horizontal. The tilt angle tends to decrease with height through the F2 region, and for magnetic dip angles over Australia, the TID tilt angle will usually be a little larger than that of the forcing AGW (by up to ~10°).

Note that in discussing TIDs, some authors (e.g. Bowman [1992]) refer to tilts in the iso-ionic contours, between the phase-fronts and the local horizontal plane, which is essentially the complement of the angle defined above. This is consistent with the usage in **Chapter 7**, in relation to large-scale gradients, but not **Chapter 8**, in the context of TID characterisation, where the more common definition above is adopted.

Appendix F   Publications

The following list of publications relate to or include work undertaken as part of this PhD project and have been reproduced in the subsequent sections of this appendix as noted below. Certain parts of these papers have been directly incorporated into the body of the thesis and are cited as such in the relevant chapters.

- Heitmann, A. J., Cervera, M. A., Gardiner-Garden, R. S., Holdsworth, D. A., MacKinnon, A. D., Reid, I. M., & Ward, B. D. (2017), 'Observations and modeling of traveling ionospheric disturbance signatures from an Australian network of oblique angle-of-arrival sounders' (summary paper), *Proceedings of the 32nd URSI General Assembly and Scientific Symposium*, Montreal, Canada. doi:10.23919/URSIGASS.2017.8105329

- Heitmann, A. J., Cervera, M. A., Gardiner-Garden, R. S., Holdsworth, D. A., MacKinnon, A. D., Reid, I. M., & Ward, B. D. (2018), 'Observations and modeling of traveling ionospheric disturbance signatures from an Australian network of oblique angle-of-arrival sounders', *Radio Science*, 53(9), pp. 1,089–1,107. doi:10.1029/2018RS006613 [**Appendix F.1**]

- Ayliffe, J. K., Durbridge, L. J., Frazer, G. J., Gardiner-Garden, R. S., Heitmann, A. J., Praschifka, J., et al. (2019), 'The DST Group high-fidelity, multi-channel oblique incidence ionosonde', *Radio Science*, 54(1), pp. 104–114. doi:10.1029/ 2018RS006681 [**Appendix F.2**]

- Heitmann, A. J., & Gardiner-Garden, R. S. (2019), 'A robust feature extraction and parameterized fitting algorithm for bottom-side oblique and vertical incidence ionograms', *Radio Science*, 54(1), pp. 115–134. doi:10.1029/2018RS006682 [**Appendix F.3**]

- Turley, M. D. E., Heitmann, A. J., & Gardiner-Garden, R. S. (2019), 'Ionogram RFI rejection using an auto-regressive interpolation process', *Radio Science*, 54(1), pp. 135–150. doi:10.1029/2018RS006683 [**Appendix F.4**]

- Gardiner-Garden, R., Cervera, M., Debnam, R., Harris, T., Heitmann, A., Holdsworth, D., et al. (2019), 'A description of the Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE)', *Advances in Space Research*, 64(10), pp. 1,887–1,914. doi:10.1016/j.asr.2019.03.036 [**Appendix F.5**]

F.1   Heitmann et al., *Radio Science*, 2018

The following is a reprint of Heitmann, A. J., Cervera, M. A., Gardiner-Garden, R. S., Holdsworth, D. A., MacKinnon, A. D., Reid, I. M., & Ward, B. D. (2018), 'Observations and modeling of traveling ionospheric disturbance signatures from an Australian network of oblique angle-of-arrival sounders', *Radio Science*, 53(9), pp. 1,089–1,107.

This paper was presented at the 32nd URSI General Assembly and Scientific Symposium, Montreal, Canada, initially appearing in summary form as Heitmann et al. [2017]. Its contents, which include an overview of the ELOISE AoA ionosonde, preliminary observations, and synthetic TID comparisons, have largely been distributed among **Chapters 3, 4, 5** and **8** of this thesis.

The manuscript was written by myself, as the lead and corresponding author, and covers research directly conducted as part of my PhD candidature.

Observations and Modeling of Traveling Ionospheric Disturbance Signatures From an Australian Network of Oblique Angle-of-Arrival Sounders

Andrew J. Heitmann[1,2] (ID), Manuel A. Cervera[1,2] (ID), Robert S. Gardiner-Garden[3] (ID), David A. Holdsworth[1,2] (ID), Andrew D. MacKinnon[2] (ID), Iain M. Reid[2] (ID), and Bruce D. Ward[1,2]

[1]Defence Science and Technology (DST) Group, Edinburgh, South Australia, Australia, [2]School of Physical Sciences, The University of Adelaide, Adelaide, South Australia, Australia, [3]Defence Science and Technology (DST) Group, Eveleigh, New South Wales, Australia

Abstract A network of oblique angle-of-arrival (AoA) ionosondes was installed as part of the Elevation-scanned Oblique Incidence Sounder Experiment (ELOISE) in September 2015. The ELOISE experimental campaign was designed to study the spatial and temporal structure of ionospheric variability at midlatitudes, of which traveling ionospheric disturbances are a key component. The new AoA sounder makes use of Defence Science and Technology Group's direct-digital high-frequency transmitter and receiver technology, to enable multichannel collection of both ionograms and channel scattering functions (Doppler spectra) on a common 2-D array. In this paper, the array design and onboard signal processing for the AoA sounder is described, along with a sample of results showing typical disturbance signatures across the delay, Doppler frequency, bearing, and elevation measurements. Realistic parameterized models of electron density perturbations, along with geometric ray tracing, were used to synthesize the effects of medium- to large-scale traveling ionospheric disturbances on the sounder observables and aid in interpreting the measured signatures.

Plain Language Summary The ionosphere is an outer charged layer of the Earth's atmosphere that is used by technologies such as over-the-horizon radar to "bounce" (refract) high-frequency radio waves out to very long distances beyond the horizon. However, as a propagation medium, the ionosphere is far from uniform, and to improve the performance of such radars, it is necessary to better understand the nature and sources of variability that affect high-frequency propagation. One of the aims of the Elevation-scanned Oblique Incidence Sounder Experiment (ELOISE), conducted by Defence Science and Technology Group in September 2015, was to study the wave-like perturbations known as traveling ionospheric disturbances by receiving a network of low-power ionospheric sounder (or ionosonde) transmissions on two new antenna arrays. These were designed to measure path delay, angle-of-arrival and Doppler shifts simultaneously on multiple paths. This paper describes the ELOISE array and its signal processing and provides examples of typical observations. It is shown that the characteristic signatures of traveling ionospheric disturbances in the measurements, combined with realistic synthetic models, can provide information about the structure and behavior of such ionospheric disturbances in space and time.

1. Introduction

Traveling ionospheric disturbances (TIDs) are one of the key contributors to spatial and temporal variability of high-frequency (HF) radio propagation through the ionosphere. First observed and explained by Munro in the 1950s (e.g., Munro, 1950; Munro & Heisler, 1956), these wave-like perturbations in the background electron density are typically interpreted as the ionospheric manifestation of atmospheric gravity waves (AGWs) in the thermosphere (Hines, 1960; Hocke & Schlegel, 1996; Hunsucker, 1982). Although many candidate source mechanisms have been proposed, within the F2 layer, medium-scale TIDs (MSTIDs; with periods of 15–60 min) are broadly thought to originate from regional tropospheric sources (Gossard, 1962), while large-scale TIDs (LSTIDs; with periods of 30 min to 3 hr) are often associated with geomagnetic disturbances in distant auroral regions (Richmond, 1978). TIDs are frequently responsible for large range and angle-of-arrival perturbations in HF propagation paths, along with additional discrete modes, and their occurrence is also closely linked to nighttime spread-F in ionograms and scintillation effects caused by small-scale plasma bubbles (e.g., Abdu et al., 2009; Booker, 1979; Bowman & Munro, 1988; Kelley et al., 1981).

In recent years there has been increased interest in monitoring and forecasting these disturbances to support a diverse range of technologies sensitive to space weather, including Global Navigation Satellite Systems, low-frequency radio astronomy, and HF sky wave radar systems. Projects such as Net-TIDE (Reinisch et al., 2018) have been established to develop regional TID diagnostic and warning capabilities to such users. Over-the-horizon radars (OTHRs) designed for long-range aircraft and ship detection, such as those comprising the Jindalee Operational Radar Network (JORN) in Australia (Colegrove, 2000; Johnson et al., 2008; Sinnott, 2015), may have their ability to detect and geolocate targets degraded in the presence of TIDs. This is usually due to ionospherically induced perturbations in the backscattered returns from targets and inadequate representation of the disturbance fields in propagation models constructed from sparsely sampled ionospheric measurements. A more detailed discussion of the challenges posed by TIDs for OTHR and potential mitigation strategies can be found in Fridman and Nickisch (2001) and Nickisch et al. (2006, 2007).

Defence Science and Technology (DST) Group ran an experimental campaign (ELOISE, the Elevation-scanned Oblique Incidence Sounder Experiment) in September 2015 to observe spatial and temporal ionospheric variability across a number of different radar and optical systems in Australia and assess the performance of data-assimilative ionospheric models (such as the JORN Real-Time Ionospheric Model) in capturing the effects of this variability on HF propagation. The ELOISE experiment included two 19-element oblique incidence sounder (OIS) receiving arrays in Laverton, Western Australia and Coondambo, South Australia, observing nine midlatitude paths across the HF band. These systems were capable of measuring the received power (path loss), group delay, angle-of-arrival (AoA), and Doppler frequency for each ionospheric mode and formed part of a broader sounder network comprising 17 vertical (and quasi-vertical) incidence paths and 43 oblique incidence paths across central and northern Australia. As a whole, ELOISE thus provided a rich data set for analyzing the effect of disturbances such as TIDs on HF propagation.

This paper presents an outline of the ELOISE AoA array design and capabilities, along with a sample of the results collected during the 2015 campaign, showing the off-angle returns and perturbed Doppler spectra characteristic of disturbances and dynamic behavior in the ionosphere. Examples of quasiperiodic and spatially moving disturbances from across the sounder network are interpreted as TIDs, and one of the aims of the subsequent analysis was to evaluate how well these TID signatures can be synthesized by ray tracing through realistic parameterized models of electron density perturbations. Alternative TID models, from the simple corrugated mirror reflector to the AGW-seeded formulation of Hooke (1968), are evaluated in terms of their ability to reproduce the AoA observations.

2. Array Design and Calibration

While there exist a number of commercially available vertical incidence sounders with an AoA and/or Doppler capability, for example, the Lowell Digisonde (Reinisch et al., 2009) and Dynasonde (Rietveld et al., 2008), such systems are not as common for one-way oblique incidence as they require comparatively large arrays. Many earlier oblique experiments (e.g., Rice, 1973; Sherrill & Smith, 1977; Sweeney, 1970) focused on AoA observations over narrow frequency channels only, although several more modern studies have collected oblique AoA soundings over the full HF band (e.g., Black et al., 1993; Rogers et al., 2003; Sherrill & Brown, 1997; Vertogradov et al., 2013; Wright & Kressman, 1983). For the ELOISE campaign, a new AoA sounding system was built using DST Group's direct-digital HF transmitter (DOTS) and receiver (DORS) designs, which employ a low-power (20 W) wideband chirp waveform to function as a high-fidelity OIS (Gardiner-Garden et al., 2008, 2011).

The two ELOISE receiving arrays at Laverton and Coondambo were designed to simultaneously measure a number of intersecting oblique paths with midpoints over central Australia (geomagnetic latitudes from 30–40°S). The use of multiple aspect angles to observe propagation through the same region of the ionosphere was a core component of the ELOISE experiment, on the basis that this would offer improved characterization of the HF response to medium- and large-scale ionospheric disturbances; in particular, any anisotropic properties. A map of the two array locations and nine key AoA paths are shown in Figure 1. The great circle path lengths ranged from 900 to 2,700 km, with 1-hop elevations spanning 0 to 30°. The schedule timeline (described in more detail below) included both a 2- to 36-MHz wideband sounding and three channel scattering function (CSF) measurements of the Doppler spectra every 3.75 min on each of these nine
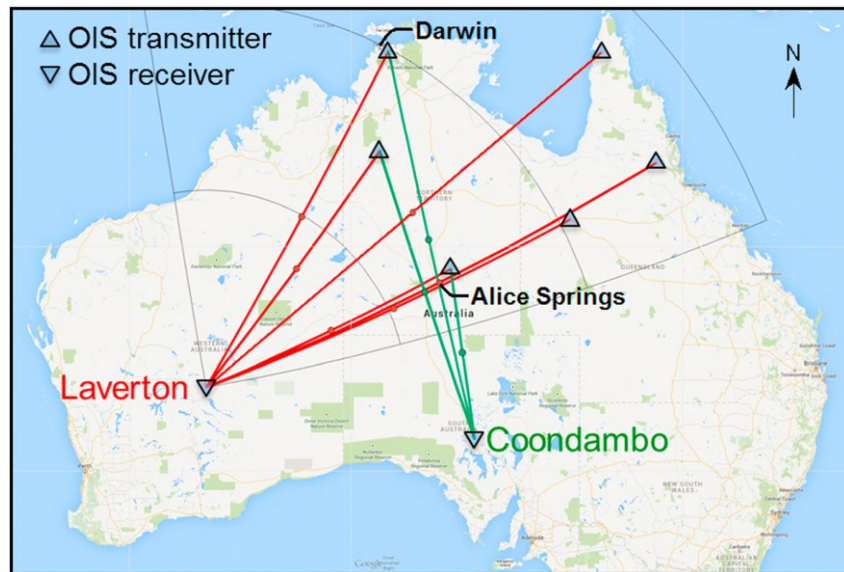
Figure 1. Locations of the two receiving arrays (Laverton and Coondambo) and nine oblique paths observing angle-of-arrival and Doppler during the ELOISE campaign. The black lines define the eastern arc of coverage for the co-located JORN OTHR at Laverton.

paths. Thus, up to 144 soundings and 432 CSF measurements were made per hour during the main collection period from 27 August to 21 September 2015.

Each ELOISE array consisted of 19 antenna elements on a horizontal plane, arranged as an orthogonal pair of uniform linear arms with a multichannel direct-digital HF receiver per element. The twin-arm geometry of the Laverton array is shown in Figure 2; all path bearings are with respect to True North (°T), a convention that is used throughout this paper. The asterisks represent each of the 6.5 m vertically polarized aluminum monopoles, pictured in Figure 3, which were embedded on a wire-mesh ground screen and connected to a central hut with phase-matched feed cables. One 10-element arm (of length 180 m) was aligned close to end-fire for the oblique paths of interest, while the other 10-element arm (of length 90 m and sharing a common element) was aligned close to broadside, providing the ability to estimate both elevation and bearing (azimuth) from both arms combined. Polarization discrimination of ordinary (O-mode) and extraordinary (X-mode) components was not possible with this array. Apart from a simple rotation and mirroring, the Coondambo array layout was otherwise identical to the Laverton array.

Figure 2. Overhead layout of the Laverton ELOISE array. The asterisks indicate the locations of the 19 monopole antenna elements, arranged in two uniform linear arms. Great circle bearings for the six key paths (red lines in Figure 1) are shown by the inset compass arrows.

The positioning of the elements was chosen to make best use of the available (cleared) land area and digital receiver units, while maximizing the effective aperture. For the selected interelement spacings, grating lobe ambiguities were able to be suitably managed using the 1-hop great circle geometry as a reference, with at least a ±20° unambiguous window in both bearing and elevation (with respect to the reference predictions) over the entire HF band. In practice, the window is actually much wider than this at typical propagation frequencies (e.g., at least ±80° at 10 MHz). Based on DST Group's prior experience in HF radar, and subsequent results from the ELOISE experiment itself, only the most extreme disturbances could lead to AoA deviations which fall outside this

Figure 3. Photographs of the Laverton ELOISE array, from the end of the short arm, and (inset) the digital receiver racks.

unambiguous window for the oblique paths in Figure 1. Although such a sparse 2-D array is not well suited to imaging applications (with conventional phased array beamforming offering poor sidelobe performance in general), its purpose in this case was simply to operate as an interferometer, that is, adopting a robust phase-fitting technique to estimate (but not fully resolve) AoA for the discrete one-way propagation at each frequency, delay, and Doppler value. Similar techniques have had a long history of use in radio direction-finding and imaging problems (e.g., Adams et al., 1985; Burtnyk & Wolfe, 1967).

Figure 3 shows photographs of the Laverton array and the digital receiver racks, which also contain the common GPS-disciplined time and frequency reference unit used by the DORS receivers. In this case, each receiver was capable of supporting eight independent down-converter channels (six being assigned to the red paths in Figure 1 and the remaining two assigned to other paths without a CSF capability). The Coondambo array used an older digital receiver model, supporting only three channels (being assigned to the green paths in Figure 1).

Array calibration was carried out using a combination of known line-of-sight, surface wave, and aggregated sky wave signals, covering different angles of arrival. Both the overall array alignment and interelement amplitude and phase delay offsets were corrected on scales of up to ±1°, ±1 dB, and ±2 ns, respectively, to compensate for small positioning, feed delay, and receiver mismatch errors. With this approach, bearing and elevation uncertainties of a few tenths of a degree were typical, although uncertainties increased for elevations below 10° and in the presence of unseparated multimode returns (e.g., the complex mode structure associated with spread-F and, to a lesser extent, unresolvable O/X components). A directly injected reference waveform was also used to calibrate individual receivers, and a side-by-side comparison was made with the co-located JORN OTHR, to give further confidence in the quality of results.

3. Onboard Processing and Communications

A suite of onboard signal processing was deployed on the main server of each ELOISE array, designed to keep up with the ~4,000 AoA ionograms and ~10,000 CSF dwells per day at the peak of the experiment. The onboard processing included the following components, some of which were adapted from DST Group's existing DORS oblique incidence ionosonde receiver (Gardiner-Garden et al., 2008, 2011):

1. a robust radio frequency interference rejection scheme, to improve the clarity of ionogram trace features;
2. an ionogram feature extraction and fitting algorithm, to extract key image features such as the maximum observed frequency (MOF) of F2 propagation and parameterize the midpoint electron density profile;
3. a clear channel evaluation (CCE) algorithm and adaptive scheduler for making CSF observations, to characterize the ionospheric Doppler spectra at key operating frequencies; and

4. a 2-D AoA technique to extract bearing and elevation estimates for each range and frequency cell combination in the ionogram/CSF data, based on a least-squares planar wavefront fit to the elemental phases.

All transmitter and receiver sites communicated via a virtual private network, which enabled schedules to be distributed and processed results to be retrieved in near real-time throughout the campaign. The standard files that were transferred to DST Group's central data server in Edinburgh, South Australia, contained both thresholded and byte-scaled ionogram and CSF images, as well as the ionogram features and fitted parameters. Each AoA record was represented by a 3-D data cube consisting of up to five pages of frequency versus group delay images; these contained amplitude, bearing, and elevation estimates, along with bearing and elevation uncertainties. Raw I/Q data from each element were also recorded locally to support postprocessing and further algorithm development.

The following subsections outline the specific properties, including scheduling and processing steps, which relate to the ELOISE CSF collection and AoA estimation.

3.1. Channel Scattering Function Scheduling

Direct measurements of ionospherically induced Doppler shifts on OIS paths provide information about the dynamics of the ionosphere, including TIDs. The ability to measure the full Doppler spectra, rather than just a single (dominant) Doppler component at each ionogram frequency and delay, was favored for the ELOISE experiment, as it provided an extra dimension for separating propagation modes and characterizing spread-Doppler conditions. The CSF mode, using a repetitive linear frequency modulated continuous wave waveform centered on a fixed frequency, enabled range-Doppler images to be collected alongside each OIS ionogram, without the need for any additional hardware. A similar technique has been demonstrated in previous experiments by DST Group (Harris & Frazer, 2005), although not with the same level of automation developed for ELOISE.

Six of the existing DOTS OIS transmitter sites had their software upgraded ahead of the experiment to support an interleaved CSF operating mode, whereby each wideband (ionogram) sweep was preceded by three short (6.6 s) CSF dwells. The same output power (nominally 20 W) applied to both wideband and CSF waveforms. To accommodate the additional processing and CSF dwell time, the frequency limits of the wideband sweep were reduced from 2–45 MHz to 2–36 MHz, without any change to the 3.75-min revisit rate or 250-kHz/s sweep rate.

Adaptive scheduling of the CSF frequencies (within the ionogram sweep limits) and bandwidths (from 7.5 to 30 kHz) was controlled by two CCE processes running onboard the Laverton and Coondambo receiving systems. The purpose of having real-time CCE was to ensure that, as a secondary user, ELOISE CSF transmissions were made on a noninterference basis with other active HF users and avoided known forbidden frequencies. At any given time, each DOTS transmitter had its CSF scheduling assigned to either one of these two sites, while the other site remained slaved.

The CCE algorithm is based on the approach used for frequency advice on the JORN radars (Earl & Ward, 1986, 1987). Each CCE system was fed by background narrowband noise data, extracted from the "guard band" of DORS ionograms (i.e., at group delays less than the great circle ground range), for the single omnidirectional antenna at the center of the array. All 8-kHz channels in which the background power spectral density is less than 7 dB above the lower decile over the surrounding 1-MHz interval are deemed to be *unoccupied*. Furthermore, those 8-kHz channels that remain unoccupied for at least three of the last four observations (within 15 min) are deemed to be *clear*. Only contiguous blocks of clear and not forbidden channels spanning up to 30 kHz in bandwidth are candidates for CSF transmission.

Figures 4 and 5 show examples of the background narrowband noise data and CCE output, respectively, for 24 hr of Laverton observations. The F2 MOF for the Cloncurry to Laverton path is overlaid, exhibiting the typical pattern of diurnal variation embedded with smaller-period ionospheric disturbances. At a remote site such as this, most of the strong interferers (horizontal striations) and broadband background noise in Figure 4 are ionospherically propagated sources at very long distances. Radio sources above the ionosphere also contribute to the background noise at higher frequencies. CCE identifies the green (clear/not forbidden) bands in Figure 5 in which CSF dwells may be scheduled, and only those below the F2 MOF offer propagation support. Note that frequencies that are found to be internally noise limited on the DORS receiver are rejected
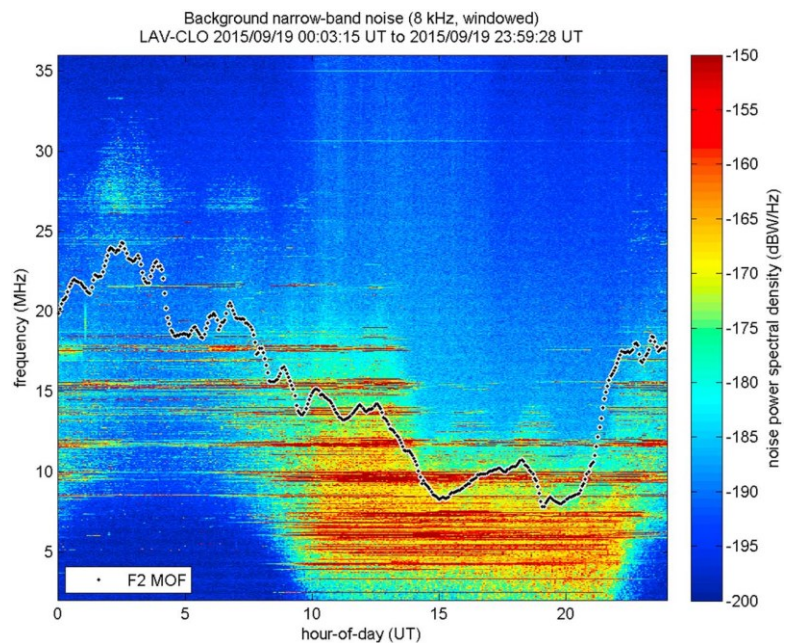
Figure 4. Background narrowband noise data from Laverton on 19 September 2015, with the F2 maximum observed frequency for the Cloncurry to Laverton (2,056 km) path overlaid. The local time at Laverton is universal time + 8 hr.

and shown in black; however, these are not operationally useful frequencies, being contained to the lower HF band during the day and upper HF band during the night.

Target CSF frequencies are based on fixed logarithmic ratios of the F2 MOF, extracted automatically from the previous ionogram, and the frequency selection maintains temporal persistence whenever permitted by
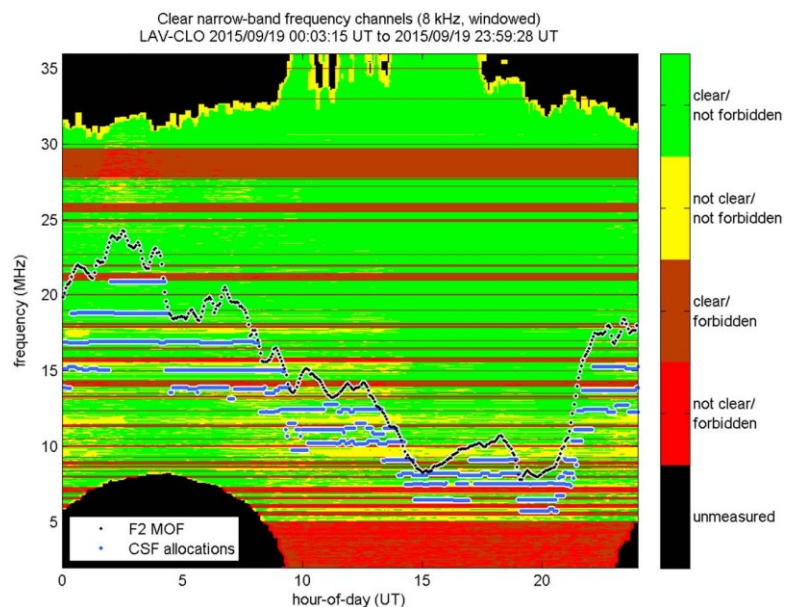
Figure 5. Clear channel evaluation output from Laverton on 19 September 2015, with the F2 maximum observed frequency and three concurrent channel scattering function frequency allocations for the Cloncurry to Laverton path overlaid.

channel availability rules. Except at night, when the lower HF band becomes significantly congested, the preferred CSF bandwidth of 30 kHz was usually available in the CCE output. A time-staggered lane scheme, with a fixed 10-Hz waveform repetition frequency, was implemented to allow the six CSF transmitters to potentially share some of the same frequency bands and further minimize the likelihood of interference.

Figure 5 shows the agile frequency allocations for each of the three CSF dwells, which are updated every 3.75 min. At each time, the frequencies represent 72.9%, 81.0%, and 90.0% of the median MOF over the last 15 min (subject to some tolerance for channel availability); this is thought to cover the band of most operational relevance to a JORN-like radar.

Although the CSF frequencies were selected with a specific oblique path in mind, the ELOISE virtual private network facilitated schedule distribution to multiple OIS transmitter and receiver sites, enabling the collection of CSF dwells wherever ionospheric propagation existed. As such, the six transmitters were typically received at both Laverton and Coondambo simultaneously, giving up to 12 CSF paths. However, only nine of these paths (six into Laverton and three into Coondambo) were measured on the AoA arrays. Onboard range-Doppler processing used a conventional fast Fourier transform-based approach (Fabrizio, 2013), although high-resolution spectral analysis techniques have been applied in offline processing.

A sample CSF range-Doppler image is shown in Figure 6, alongside an ionogram that probed the same frequency 69 s later. The kink-like distortions on the 1-hop F2 trace (indicated by the red arrows) are the characteristic signatures of MSTIDs (e.g., Harris et al., 2012). The same group delay structure is clearly seen in both images, albeit at lower resolution in the CSF owing to its lower processing bandwidth (i.e., 30 kHz versus 200 kHz). This particular case is representative of midlatitude ionospheric conditions in the early evening, when the F2 layer height is increasing, as evident from the slightly negative Doppler shifts in the CSF image.

3.2. Angle-of-Arrival Estimation

Both ionogram and CSF observations received on the AoA arrays were subject to onboard AoA estimation, which applied an interferometric technique to fit (in a least squares sense) a planar wavefront jointly across the phase angles on all 19 elements of both array arms. This was carried out independently for each range and frequency cell combination in the ionogram/CSF data, using the predicted great circle bearing and 1-hop equivalent mirror elevation (based on group delay only) to resolve grating lobe ambiguities. The fundamental assumption is that each range/frequency cell contains a single dominant mode; separation of modes must therefore make use of information from the other data dimensions (i.e., frequency, delay, and Doppler). Given point-to-point propagation via discrete modes, this is believed to be a reasonable assumption for a significant proportion of the ELOISE data set. Where there are unresolved modes (e.g., O/X components), either one mode will tend to dominate or else the resultant nonplanarity will be flagged by a corresponding increase in the fit uncertainty.

In short, the AoA estimation algorithm consists of the following steps:

1. Resolve 2π phase ambiguities on each arm of the array using predicted geometry.
2. Fit a planar wavefront jointly to the unwrapped phase angles across the two arms.
3. Transform the least-squares fit in wave vector space, including uncertainties, to bearing and elevation.

With the 2-D array being undersampled at frequencies above 7.5 MHz (i.e., antenna spacing $d > \lambda/2$ on the long arm for wavelengths $\lambda < 40$ m), it is first necessary to resolve 2π ambiguities in the measured phase angles (with calibration corrections applied). Using the known great circle bearing and 1-hop equivalent mirror elevation at the measured delay, the closest valid solution in wave vector space is chosen to be the correct one. This enables the phase angles to be effectively unwrapped along each array arm, as illustrated in Figure 7. While there remain some small variations in the interelement phase differences, due to both ionospheric distortions and uncorrected calibration errors, the phase angles defining the wavefront are very close to linear on each arm. Noting that this is a fairly typical example, the assumption of a planar wavefront is felt to be justified. In this case, the slopes are related to the steer angles on each arm; the frequency dependence may be removed by converting phase angle to phase delay.

The linear equations that relate the phase angles on each of the 19 elements can be written in matrix form as follows:
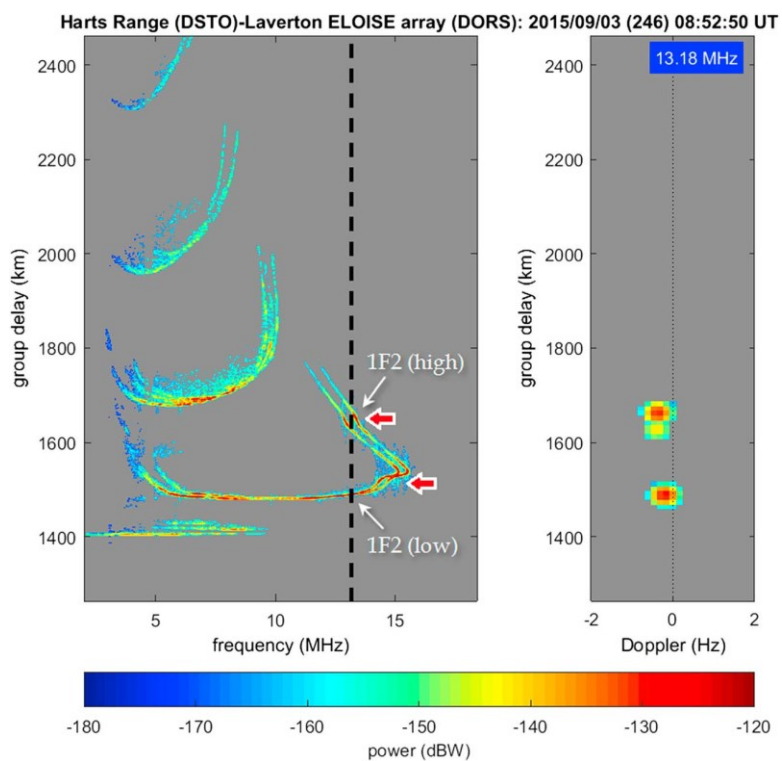
Figure 6. (left) A sample ionogram and (right) corresponding channel scattering function (CSF) range-Doppler image at 13.18 MHz, for a 1,383-km oblique incidence sounder path from Harts Range (near Alice Springs, Northern Territory) to Laverton. The key 1-hop F2 low- and high-angle propagation modes at the CSF frequency (dashed black line) are annotated on the ionogram; see, for example, Davies (1990) or McNamara (1991) for a more detailed description of ionograms and their interpretation. The red arrows highlight the trace distortions ("kinks") attributed to medium-scale traveling iono-spheric disturbances. Note that each pixel combines all elements of the array using a conventional phased array beam in the direction of the estimated angle-of-arrival.
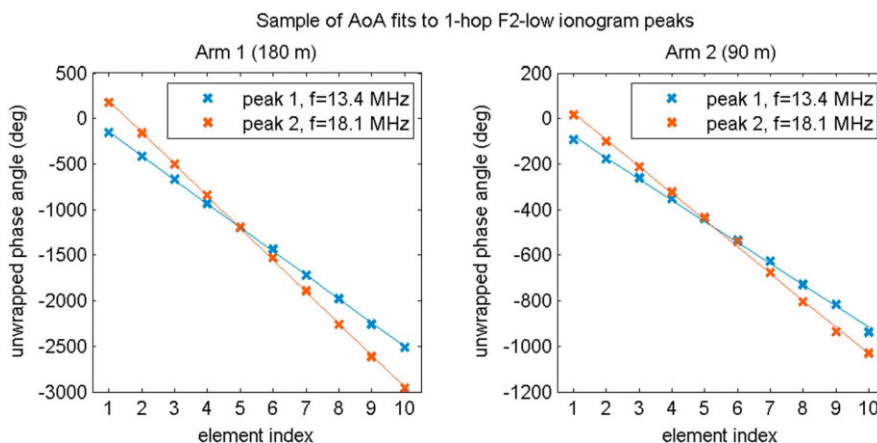
Figure 7. Measurements of unwrapped phase angles for two ionospheric modes at different frequencies on the 1-hop F2 (low) mode from Cloncurry to Laverton. Observations are shown for both the (left) long (180 m) and (right) short (90 m) array arms. The lines of best fit are overlaid to provide a sense of linearity, although it is noted that the actual angle-of-arrival estimation algorithm fits both arms jointly.
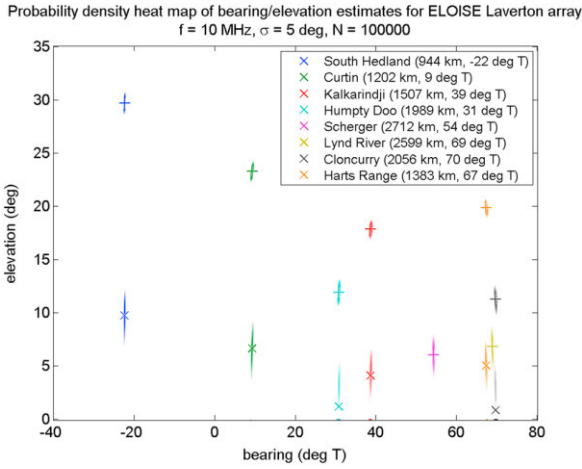
Figure 8. Simulated distributions of angle-of-arrival (AoA) estimates in bearing/elevation space at a frequency of 10 MHz, assuming white Gaussian phase noise with a standard deviation of 5°. For each oblique path, returns from both 1-hop E-mode (100-km virtual height) and 1-hop F-mode (300-km virtual height) are evaluated, using 100,000 randomized samples (without regard for the antenna gain patterns). The true AoA positions are indicated by the "x" (E-mode) and "+" (F-mode) markers.

$$X\boldsymbol{p} = \boldsymbol{b},$$

where design matrix $X = \begin{bmatrix} x_1 & y_1 & 1 \\ \vdots & \vdots & \vdots \\ x_{19} & y_{19} & 1 \end{bmatrix}$ contains the element positions in the local horizontal coordinate system (x_n, y_n for $n = 1, \ldots, 19$), parameter vector $\boldsymbol{p} = \begin{bmatrix} k_x \\ k_y \\ \varphi_0 \end{bmatrix}$ contains the corresponding components of the incident wave vector (k_x, k_y) and the phase angle at the origin (φ_0), and measurement vector $\boldsymbol{b} = \begin{bmatrix} \varphi_1 \\ \vdots \\ \varphi_{19} \end{bmatrix}$ contains the unwrapped phase angles at each of the elements (φ_n). The axes of the (x, y) coordinate system are aligned with the semimajor and semiminor axes of the least-squares uncertainty ellipse for the ELOISE array geometry.

The least-squares solution to the normal equations above can then be expressed by the standard formula

$$\widehat{\boldsymbol{p}} = \left(X^T X\right)^{-1} X^T \boldsymbol{b},$$

where the hat notation is used to denote an estimated variable. In practice, the Cholesky decomposition of $X^T X$ is used for numerical stability, rather than actually computing the pseudo inverse as written above.

Estimates of bearing (azimuth) $\widehat{\theta}$, measured clockwise from True North, and elevation $\widehat{\beta}$ are related to the wave vector components by the simple trigonometrical equations

$$\widehat{\theta} = \text{atan2}\left(\widehat{k}_x, \widehat{k}_y\right) + \theta_y,$$

and

$$\widehat{\beta} = \text{acos}\left(\frac{\lambda}{2\pi} \sqrt{\widehat{k}_x^2 + \widehat{k}_y^2}\right),$$

where θ_y defines the bearing of the y axis.

Note that for a physical solution,

$$\sqrt{\widehat{k}_x^2 + \widehat{k}_y^2} \leq \frac{2\pi}{\lambda}.$$

Assuming no uncertainties in the element positions, the standard deviations for each component of the parameter vector estimate $\widehat{\boldsymbol{p}}$ are directly related to the sum of squared residuals in the fit. This allows uncertainties in the corresponding bearing and elevation estimates to be characterized. Higher uncertainties may be indicative of multimode structure unresolved in frequency, delay, and Doppler, for example. Besides unresolved multimode, calibration errors, receiver noise, residual radio frequency interference, and spatial incoherency across the array aperture can also contribute to AoA uncertainties. Array misalignment and antenna gain patterns are further sources of potential error but manifest as fixed biases rather than contributing to the variance.

Simulated distributions of AoAs estimated using the above technique are shown in Figure 8, based on spherical mirror model propagation across eight oblique paths into Laverton, and for two representative virtual heights: 100 km (E-mode) and 300 km (F-mode). All but South Hedland and Curtin were among the key ELOISE paths with CSF functionality. The effects of receiver noise are modeled by introducing a random
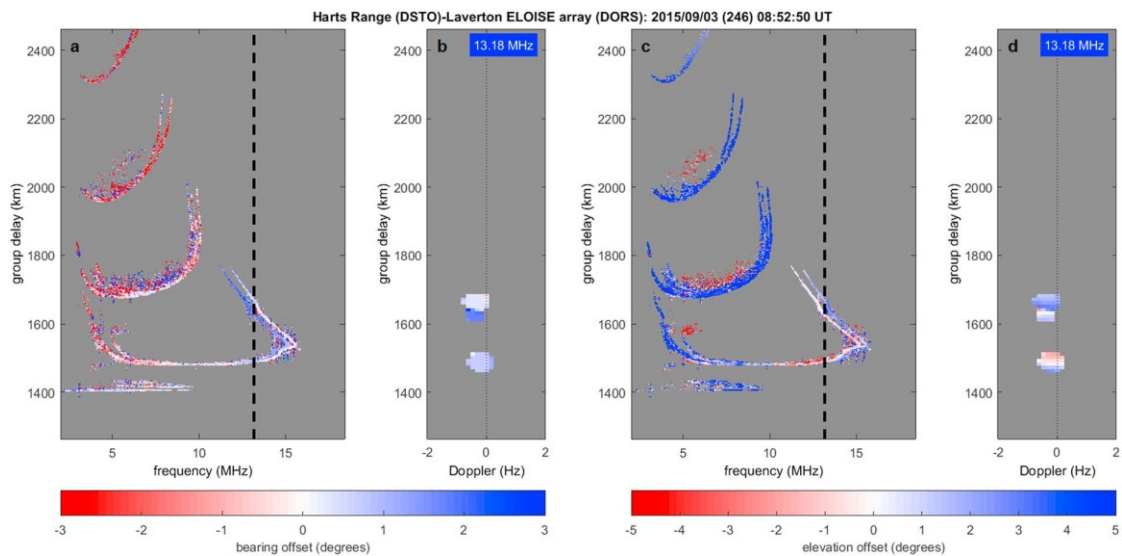
Figure 9. (a and b) Estimates of bearing offset (measured clockwise with respect to the great circle bearing) and (c and d) elevation offset (measured with respect to the 1-hop equivalent mirror elevation) for the same ionogram and channel scattering function images as in Figure 6.

(Gaussian) phase error to each element. A standard deviation of 5° is adopted for this phase error, based on closed-loop receiver tests at around 20 dB SNR, and results in the bearing and elevation spread pictured. This plot highlights both the overall distribution of oblique paths in AoA-space, as well as some of the challenges faced by an array of this size operating at low elevations, with representative E-mode returns of 0–10° and F-mode returns of 5–30° in elevation. Below 10° in particular, the distributions are very broad in elevation and potentially become biased (overestimated) as symmetrically distributed errors in wave vector space map to a skewed distribution in bearing/elevation space, and unphysical solutions are clipped at an elevation of zero. The null in the antenna gain pattern at 0° elevation is believed to further exacerbate this problem in real observations. Whereas the standard deviation in bearing is typically around 0.2°, for elevation, this increases to 0.4–0.7° (above 10°) and 1.0–1.5° (below 10°). Overall performance will generally be better at higher SNR, although the geometrical limitations remain.

A sample of the onboard AoA output is shown in Figure 9, for the same sounding as in Figure 6. Bearing and elevation offsets are defined as the differences between the measured angles and the predictions, based on the great circle and 1-hop mirror geometry. For the 1-hop F2 low and high ray modes observed in the CSF image (as annotated in Figure 6), the magnitudes of the bearing offsets are mostly less than 1° in Figure 9; that is, they are received in line with the great circle bearing of 67°T (measured clockwise from True North), as is typical under fairly benign ionospheric conditions. However, note that the ordinary (O-mode) and extraordinary (X-mode) polarization components of the 1F2(high) mode are subtly distinct, with the O-mode being more perturbed by the passage of the MSTID at this particular time (although the reverse is true at other times). While geomagnetic effects alone can influence the AoA, simulations by Dao et al. (2016) have indicated that these O/X deviations are insignificant compared to the potential effects of horizontal gradients and disturbances on longer oblique paths such as this.

The corresponding elevation offsets in Figure 9 are slightly larger, and again, there is an O/X-mode difference, but the perturbations are still relatively small on the whole (less than 3°). As the predicted geometry varies with delay, an artificial delay-dependent gradient is found in the elevation offsets of the modes in the lower-resolution CSF image (panel d); this is purely an artifact introduced by subtracting the equivalent elevation and does not relate to unresolved mode components. The multihop returns in the ionogram (panel c) also tend to have positive offsets, as the 1-hop predicted geometry underestimates elevation for these multiple hops.
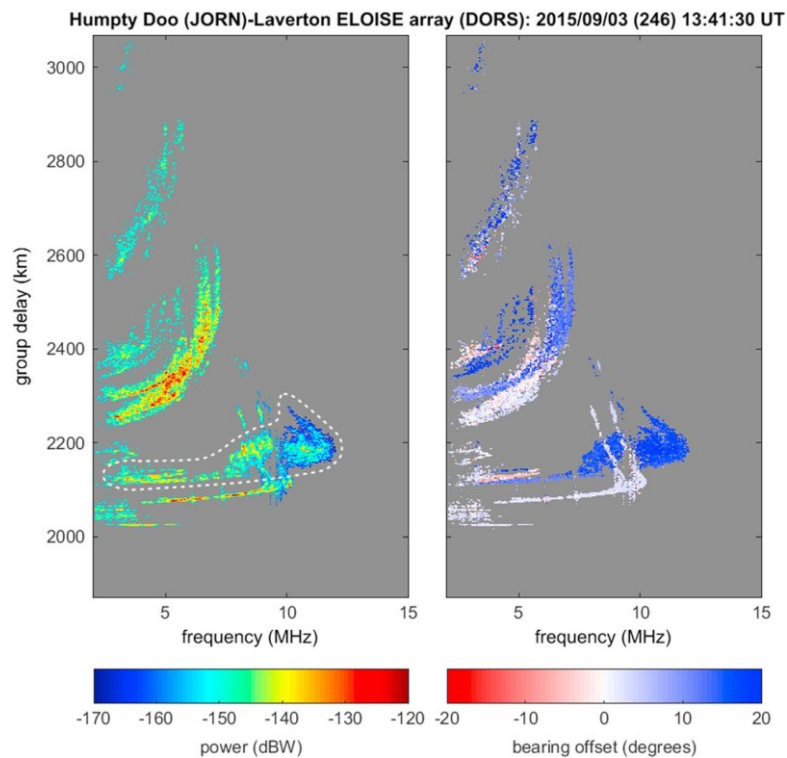
Figure 10. Sample of an ELOISE ionogram scaled in terms of both (left) power and (right) bearing offset. The local time at Laverton is universal time + 8 hr (i.e., 9:41 p.m.). The dashed white line in the left panel identifies an additional ("satellite") 1F2 trace, characterized by large deviations in bearing.

4. Sample Disturbance Signatures

ELOISE AoA ionogram data were collected almost continuously on the Laverton array from late July 2015 to mid-January 2016, while the Coondambo array was operated on a shift basis (8–16 hr per day) across the first 3 weeks of September 2015, owing to its dependence on a local diesel generator. Interleaved CSF dwells were produced between 27 August and 21 September 2015, with automatic scheduling of the six CSF transmitters shared between the Laverton and Coondambo sites. Throughout this period, there were a moderate number of geomagnetically disturbed days, with Kp indices frequently in the 4–7 (active to storm) range on a scale of 0–9.

A sample ionogram from the Laverton array is shown in Figure 10, with colors scaled according to both received power (left) and bearing offset (right). This example represents relatively disturbed nighttime conditions (approximately 3 hr after sunset at F2 heights) on the oblique path from Humpty Doo (near Darwin) to Laverton, with ground range 1,989 km and great circle bearing 31°T. Elevation measurements are also available, although not shown here for sake of brevity. The off-angle ("satellite") trace that appears across most frequencies at group delays between 2,100 and 2,200 km, with a bearing offset of almost 20° in parts, is a signature of LSTIDs frequently seen in the postsunset ionosphere at midlatitudes. It is often found to be a precursor to spread-F irregularities (Bowman, 1981; Bowman & Monro, 1988; Lynn et al., 2013), which pervade the postmidnight ionosphere, particularly around the equinoxes during low sunspot conditions. Over the surrounding 45-min period, shown in Figure 11, the additional off-angle traces move outward and inward in group delay as their bearing offsets shift from negative (red) to positive (blue) values. At times, these traces partially overlap in the ionograms (e.g., 1337 UT), and the bearing offset represents only the stronger of the two returns.

An enlarged portion of the same ionogram from Figure 10 is shown in Figure 12 alongside a neighboring CSF dwell at 5.7 MHz. This frequency corresponds to the dashed vertical line on the ionogram, which captures the
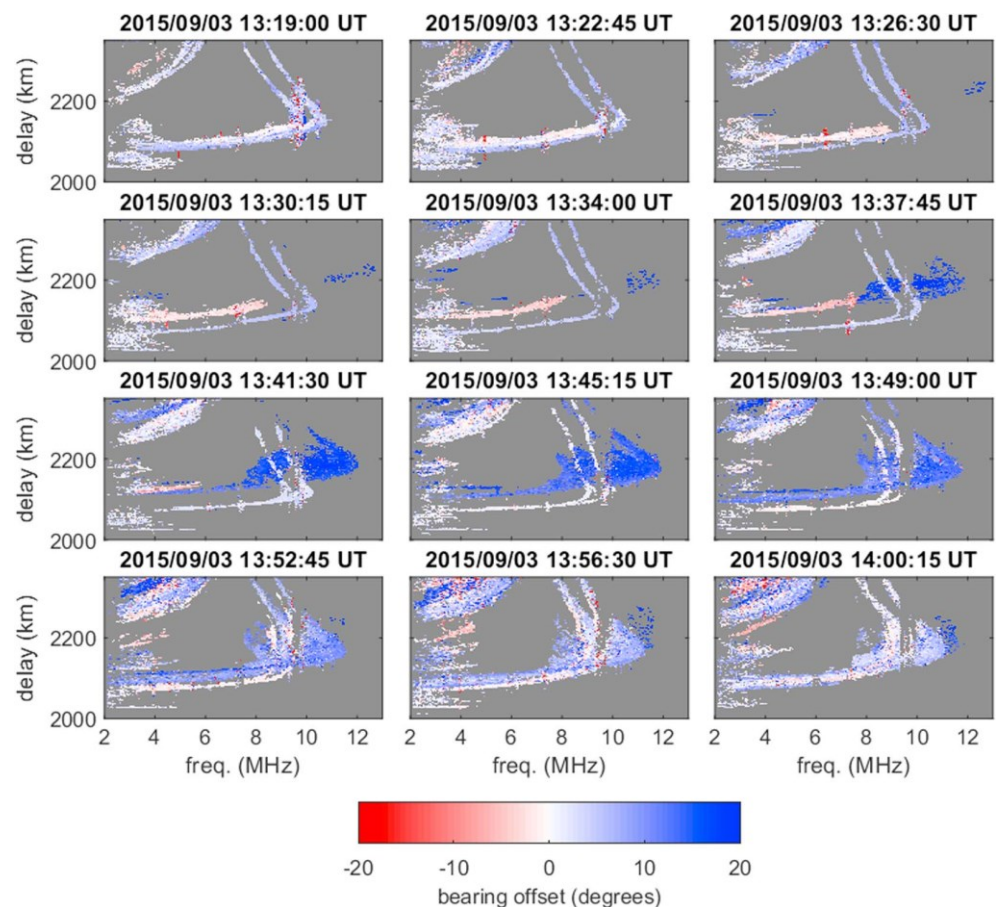
Figure 11. A time sequence of bearing offset images, measured before and after the ionogram in Figure 10. From left to right and top to bottom, the panels show one "satellite" trace breaking away from the unperturbed (direct) path and moving outward in group delay (1319–1341 UT), followed by another trace, from the opposite direction, moving inward in group delay and merging with the unperturbed path (1337–1400 UT).

lower frequency tail of the off-angle structure identified in Figure 10. An automatic peak detection and fitting algorithm, operating across each page of the AoA data cube, has identified three 1-hop F2 returns in the CSF image, including the unperturbed shorter path (at 0.0-Hz Doppler and 2,076-km group delay), and two off-angle longer paths: one with its reflection point advancing (at +0.8 Hz and 2,123 km) and one receding (at −0.7 Hz and 2,134 km). Their great circle bearing offsets were +10° and −7°, respectively, indicative of a TID structure propagating transversely (roughly northwestward) to the direction of the HF reception (northeastward). The time evolution of ionograms in Figure 11 supports this interpretation. Furthermore, it is consistent with the systematic westward and equatorward propagation directions reported elsewhere for nighttime TIDs at midlatitudes (e.g., Shiokawa et al., 2009).

The 2-hop F2 mode structure (above 2,300-km group delay) is considerably more complicated, although also contains both advancing and receding components in the CSF dwell.

The quasiperiodic nature of TID signatures becomes most evident in time sequences of the multidimensional peak data. Mode classification of peaks is carried out using the parameterized midpoint profiles derived from each ionogram (Gardiner-Garden et al., 2011), leaving just the 1F2 (low) returns in Figure 13. The five rows represent each of the observables, for both ionogram peaks (left column) and CSF peaks (right column) at a frequency of around 8 MHz. For the most part, the ionogram and CSF data sets capture the same ionospheric variability, although the higher delay resolution of the ionograms means that O- and X-mode
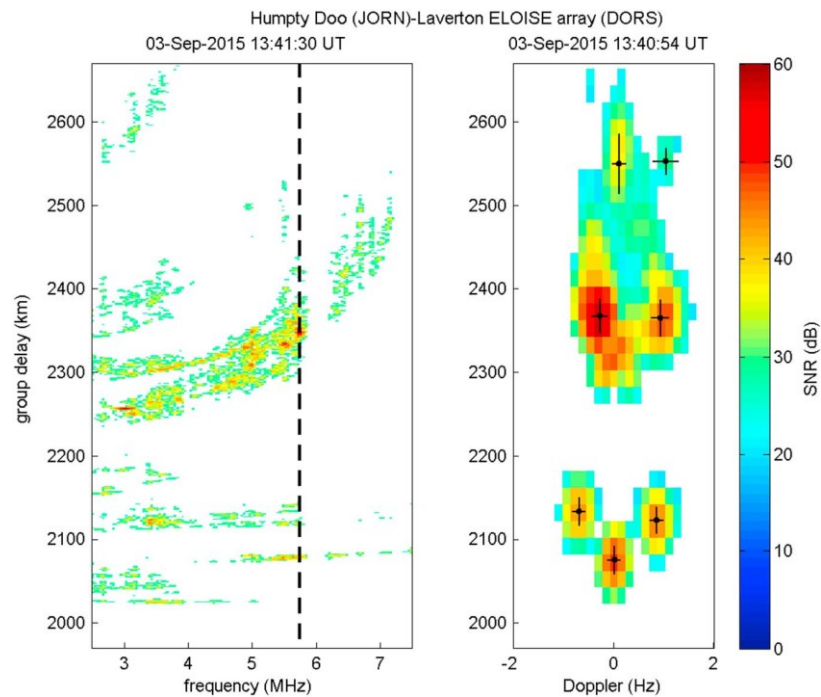
Figure 12. (left) Enlarged portion of an ELOISE ionogram (rescaled from Figure 10) alongside the (right) matching channel scattering function dwell at 5.7 MHz (with fitted peaks).

components can more often be resolved (e.g., between 1200 and 1330 UT). The vertical error bars, representing the fitted width of the peaks, capture both this instrumental resolution and ionospheric spread.

Common periodicities across group delay, Doppler, and AoA are a good indicator of TID activity, although it is important to note that not all quasiperiodic signatures are TIDs; for example, mesoscale plasma structure deposited by the equatorial fountain or front-like structures oscillating over the sounder midpoint may produce similar effects (e.g., Lynn et al., 2013; Lynn et al., 2016). Spatial lags between multiple sounder paths and a phase pattern descending as a function of height (and plasma frequency) are additional characteristics that may be used to identify AGWs manifested as TIDs (Hines, 1960). In the case of Figure 13, a long sequence of MSTIDs was clearly observed as "kinks" in the traces of the individual ionograms, although many appeared to be solitary and lacked a distinct period. This was typical of the majority of nights during the ELOISE campaign. While these TIDs did not manifest as large perturbations in the group delay of the fitted peaks, the effects are more pronounced in Doppler and elevation (and to a lesser extent, bearing). During this period, elevation predicted from equivalent geometry (group delay only) was found to differ by up to ±5°. The bearing also appeared to be shifted slightly away from the equator throughout the night, which may be attributable to persistent large-scale horizontal gradients in the background ionosphere.

5. Synthetic TID Modeling

To aid the interpretation of results, realistic TID models and associated ray tracing/homing algorithms were developed to synthesize the observables from the AoA array (i.e., group delay, Doppler, bearing, and elevation). The effects of TIDs from far-field sources on HF propagation have been widely synthesized in the past using such models, which seek to parameterize the TID (or its underlying AGW) in terms of amplitude, period, wave vector, and phase. Commonly adopted TID model variants include the following:

1. a corrugated (sinusoidal) mirror reflector (e.g., Huang et al., 2016; Lyon, 1979);
2. a sequence of tilted Gaussian peak or trough perturbations applied to a background ionosphere (e.g., Lobb & Titheridge, 1977);
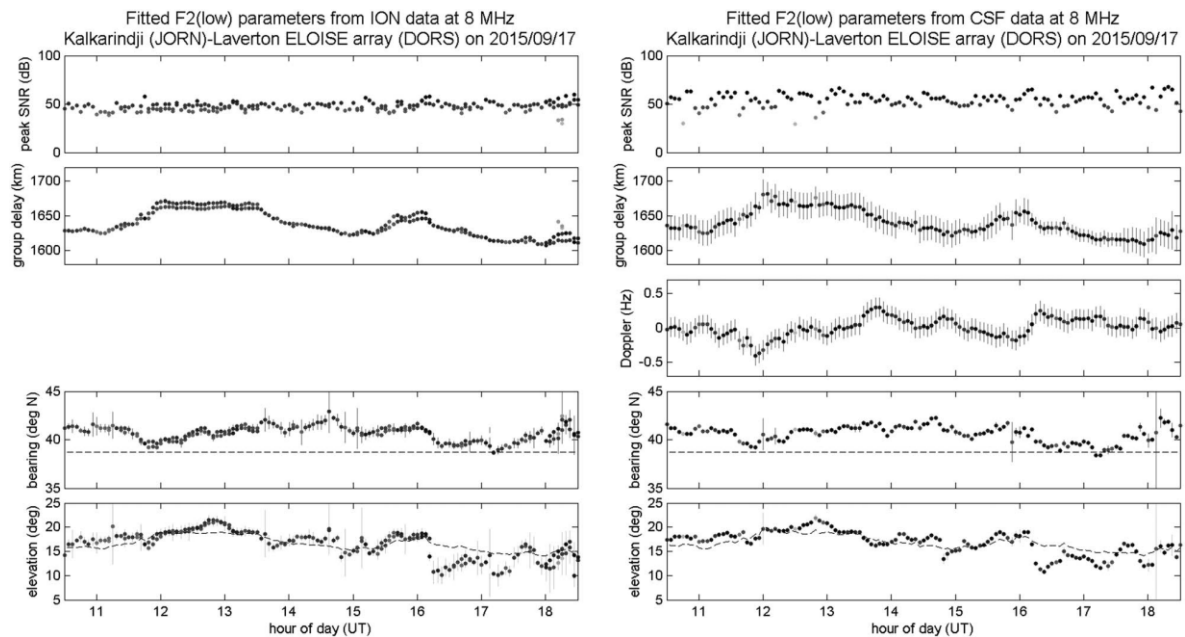
Figure 13. A nighttime sequence of multidimensional angle-of-arrival peak data across each of the (left) ionogram and (right) channel scattering function (CSF) observables, for 1-hop F2 (low) propagation at around 8 MHz on the Kalkarindji to Laverton path (1,507 km, 39°T). The dashed lines in the fourth and fifth rows of plots represent the great circle bearing and equivalent mirror elevation (from group delay), respectively. Note that Doppler measurements are only available from the CSF dwells.

3. an AGW-seeded physics-based ("Hooke") model, driven by collisional interactions with the neutral atmospheric constituents (e.g., Cervera & Harris, 2014; Hooke, 1968). Similar physics-based perturbation model treatments have also been described by Clark (1970), Clark et al. (1971), Davis (1973), and Kirchengast (1996).

In this paper, only models 1 and 3 are considered, with monochromatic TID perturbation fields adopted as a starting point. Oblique propagation synthesis is carried out in 3-D using basic specular ray homing for model 1 and full numerical ray tracing for model 3.

The corrugated spherical mirror reflector (model 1) is the simplest of the candidate TID models, consisting of a spherical shell ionosphere with a sinusoidal perturbation superposed. It has the advantage of being quick to compute (being specular reflection only) and directly invertible from observations under certain conditions, for example, using the generalized frequency and angular sounding technique (Beley et al., 1995; Galushko et al., 2003; Paznukhov et al., 2012). It can also be relatively easily extended to support multiple harmonic components and large-scale tilts. Such a model is therefore useful for classifying disturbances in the observed data and providing first-order estimates of the TID parameters.

As an example of the utility of model 1, Figure 14 shows a time series of CSF observables (fitted F2 peak data from Kalkarindji to Laverton, 1,507 km, 39°T) with synthesized values overlaid for a corrugated mirror moving transversely to the HF propagation plane at a mean height of 280 km (untilted). The modeled TID, constructed by manually adjusting the parameters to match (by eye) the observables on this path, had a 15-km amplitude, 70-min period, 400-km horizontal wavelength, and midpoint bearing of −53°T (northwestward). This period is characteristic of a LSTID, potentially of southern auroral origin, and similar to the disturbance signature shown previously in Figure 12. While such disturbances were less common than the medium-scale perturbations shown in Figure 13, similar "satellite" trace characteristics were seen on approximately a third of the nights during the monthlong campaign. This example has been chosen for its single dominant harmonic component, which makes it a less complicated modeling problem than Figure 13 (although multiple components are supported in the TID model implementation).

313
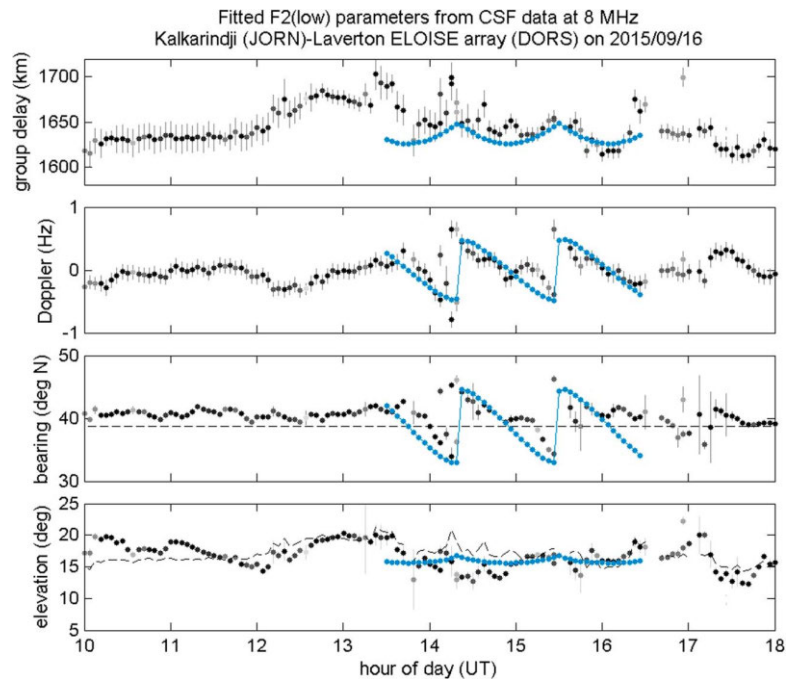
Figure 14. A nighttime sequence of channel scattering function peaks (black dots) for 1-hop F2 (low) propagation at an operating frequency of around 8 MHz. Synthesized values for the corrugated mirror traveling ionospheric disturbance model are overlaid on the most disturbed interval from 1330–1630 UT (blue dots, showing shortest path only). From top to bottom, the panels are group delay, Doppler, bearing, and elevation. Note that moderate range spread was observed for much of this period, and at times, the frequency was very close to the F2 leading edge, causing the group delay estimates to be rather noisy.

The apparent horizontal wavelength includes the effects of background neutral winds, and as such, it is likely that the intrinsic AGW wavelength is in fact far greater than 400 km (Chum et al., 2012). Indeed, the empirical Horizontal Wind Model HWM14 (Drob et al., 2008; Drob et al., 2015) predicts a strong prevailing east-north-eastward wind of over 100 m/s at this height, which largely opposes the apparent north-westward velocity of the TID model. Furthermore, if the modeled corrugation were to be extended to a "tilted" sinusoid, this could allow for a longer horizontal wavelength, more in line with the ≥1,000-km LSTID wavelengths quoted in Hunsucker (1982); as it stands, the pure sinusoid must be kept reasonably short in length to reproduce the sawtooth-like properties of the observables. Thus, it is fair to say that the precise TID model parameters should be interpreted with a degree of caution, as they are highly coupled to the model construction and an element of personal judgment. While not a perfect match to the observations in Figure 14, the model is successful at reproducing the dominant features, including the magnitude of the variations in Doppler and bearing, and the fact that there is very little of the same quasiperiodic signature in elevation due to the transverse geometry (as the reflection point is constrained to the transverse plane through the midpoint; Bramley, 1953).

The AGW-seeded approach (model 3), originally developed by Hooke (1968), applies perturbations to the continuity equations governing ion production, loss, and transport and is combined with 3-D magneto-ionic ray tracing (Cervera & Harris, 2014) to produce a TID representation that is more physically realistic. Although more computationally demanding, it supports multiple ionospheric layers/modes (e.g., using the International Reference Ionosphere as the background electron density field) and the iso-ionic contours correctly capture the characteristic forward-tilting wavefront of the AGW/TID. Over time, this appears as a height-descending phase progression, illustrated in Figure 15. Both O- and X-mode components can be synthesized using this approach.

The synthetic observables from ray tracing at 8 MHz (corresponding to the CSF observations) are shown in Figure 16, for the perturbation field in Figure 15. The modeled LSTID is again moving transversely to

314
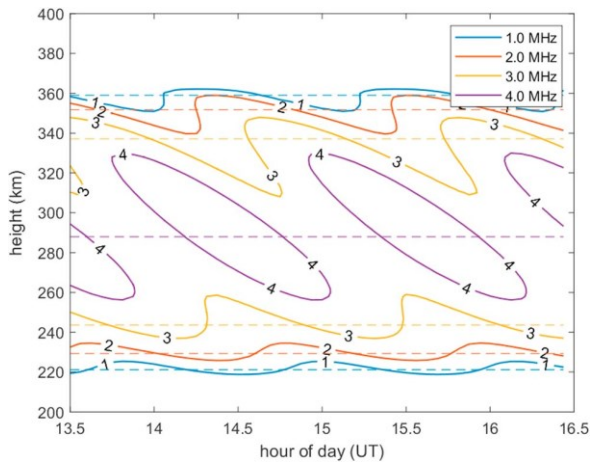
Figure 15. Illustration of iso-ionic contours (i.e., lines of constant plasma frequency in MHz), overhead a fixed point for the atmospheric gravity wave-seeded model.

the HF propagation plane with a period of 70 min and horizontal wavelength of 400 km, and with amplitude governed by the neutral wind perturbation and dissipative terms. The background ionosphere is a fixed spherically symmetric profile based on the mean of parameters fitted to the surrounding 3-hr ionogram sequence (the same observations as in Figure 14), and the perturbation parameters were chosen to synthesize a broadly similar class of disturbance as the one in Figure 14. Not all scales of AGWs are able to penetrate to F2 heights (Hines, 1960), and in this nighttime example, it appears to be difficult to realistically reproduce the full amplitude of the observed perturbation with the AGW-seeded model, without resorting to physically unrealistic perturbation parameters. Many of the higher elevation synthetic features in Figure 16 (for rays approaching the F2 peak height and undergoing stronger refraction) are also not readily apparent in the observations, despite the additional estimated path loss being less than 15 dB (within the SNR margin of the instrument). More work is required to better understand the source of these discrepancies and the limits of the Hooke model over a broader range of TIDs, some of which may not in fact be AGW-generated.
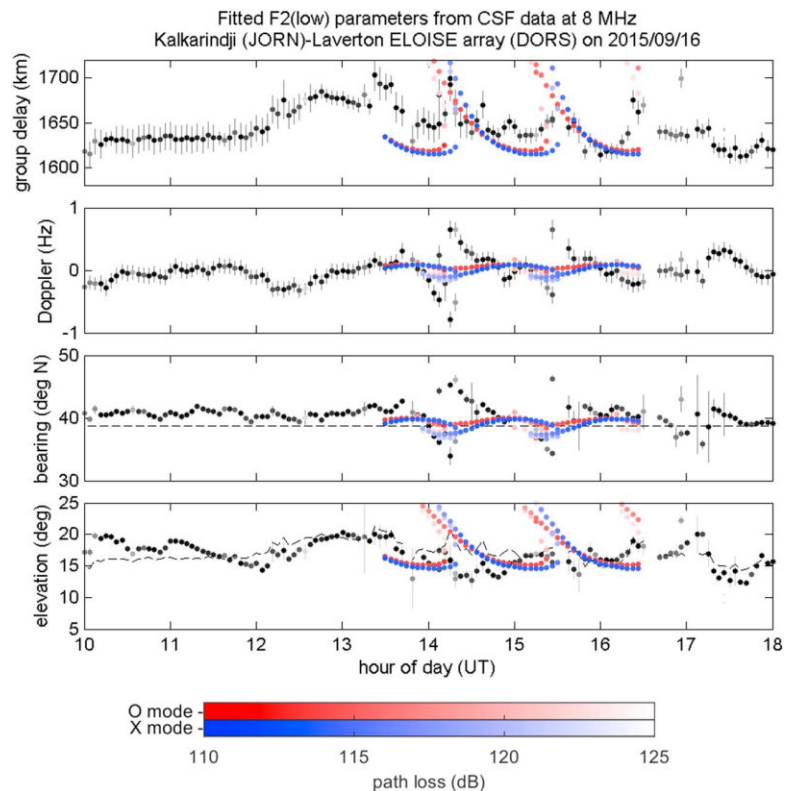
Figure 16. Observables synthesized from the atmospheric gravity wave-seeded traveling ionospheric disturbance model in Figure 15 (red/blue dots), overlaid on the channel scattering function observations from Figure 14 (black dots). Again, from top to bottom, the panels depict group delay, Doppler, bearing offset, and elevation, for a fixed operating frequency of 8 MHz. Both O-mode (red) and X-mode (blue) components are calculated by 3-D numerical ray tracing (Cervera & Harris, 2014), with color saturation indicating the estimated path loss.

Note that, for this 1,507-km path, the ray apogee heights for an 8-MHz operating frequency correspond roughly to the 3-MHz plasma frequency contour in Figure 15 (yellow line centered about 240-km height). This puts it slightly beyond the sinusoidal contour regime, where the corrugated mirror reflector might be expected to have more validity. Conversely, it means that at lower frequencies/heights, or for lower amplitude perturbations, the mirror reflector may remain a good approximation.

6. Conclusions

An overview of the ELOISE oblique AoA sounder system, developed by DST Group, has been provided, along with a sample of results from a recent campaign in September 2015. By measuring both spatial and temporal characteristics, in terms of bearing, elevation, and Doppler, the effects of ionospheric variability on different scales can be studied. A description of the array design and onboard signal processing gives an insight into the quality of results, as well as some of the challenges and accuracy limitations of operating on long-range oblique paths.

Signatures of TIDs have been identified in the ELOISE AoA observations, and comparisons with synthetic results from two candidate TID models show encouraging results, in terms of their ability to classify and parameterize the disturbance field. Despite its limitations, the simple corrugated mirror can still be effectively used to identify general classes of TIDs under certain conditions, although it is noted that some of the parameters matched by eye may be relatively imprecise. On the other hand, the AGW-seeded approach can potentially offer a more realistic TID representation at higher altitudes (i.e., for which refractive effects are more significant) but may be limited in its ability to capture the properties of some larger-scale disturbances, as in the examples provided.

While broad features can be reproduced by the models, at least qualitatively and for a single plasma frequency (iso-ionic) surface, more work needs to be carried out to interpret the many examples of AoA perturbations that do not have a single dominant (and time invariant) periodicity. With TIDs being so ubiquitous, the majority of the ELOISE observations in fact represent the superposition of many components, from different sources, that may evolve over time or exist as solitary waves. Efforts to model each spectral component in a spatially consistent way across the network of ELOISE AoA paths are ongoing. The key enablers for permitting cross-spectral analysis between oblique paths are an ability to correct for the low elevation AoA bias and a method of transforming the observables to a common domain that is agnostic to the widely different path lengths.

References

Abdu, M. A., Alam Kherani, E., Batista, I. S., de Paula, E. R., Fritts, D. C., & Sobral, J. H. A. (2009). Gravity wave initiation of equatorial spread F/plasma bubble irregularities based on observational data from the SpreadFEx campaign. *Annales Geophysicae, 27*(7), 2607–2622. https://doi.org/10.5194/angeo-27-2607-2009

Adams, G. W., Edwards, D. P., & Brosnahan, J. W. (1985). The imaging Doppler interferometer: Data analysis. *Radio Science, 20*(6), 1481–1492. https://doi.org/10.1029/RS020i006p01481

Beley, V. S., Galushko, V. G., & Yampolski, Y. M. (1995). Traveling ionospheric disturbance diagnostics using HF signal trajectory parameter variations. *Radio Science, 30*(6), 1739–1752. https://doi.org/10.1029/95RS01992

Black, Q. R., Wood, J. F. Jr., Sonsteby, A. G., & Sherrill, W. M. (1993). A direction finding ionosonde for ionospheric propagation research. *Radio Science, 28*(5), 795–809. https://doi.org/10.1029/93RS00690

Booker, H. G. (1979). The role of acoustic gravity waves in the generation of spread-F and ionospheric scintillation. *Journal of Atmospheric and Terrestrial Physics, 41*(5), 501–515. https://doi.org/10.1016/0021-9169(79)90074-6

Bowman, G. G. (1981). The nature of ionospheric spread-F irregularities in mid-latitude regions. *Journal of Atmospheric and Terrestrial Physics, 43*(1), 65–79. https://doi.org/10.1016/0021-9169(81)90010-6

Bowman, G. G., & Monro, P. E. (1988). Mid-latitude range spread and travelling ionospheric disturbances. *Journal of Atmospheric and Terrestrial Physics, 50*(3), 215–223. https://doi.org/10.1016/0021-9169(88)90070-0

Bramley, E. N. (1953). Direction-finding studies of large-scale ionospheric irregularities. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences, 220*(1140), 39–61. https://doi.org/10.1098/rspa.1953.0170

Burtnyk, N., & Wolfe, J. L. (1967). Direction-finding performance of a wide-aperture sampling linear array. *Proceedings of the Institution of Electrical Engineers, 114*(5), 581–588. https://doi.org/10.1049/piee.1967.0119

Cervera, M. A., & Harris, T. J. (2014). Modelling ionospheric disturbance features in quasi-vertically incident ionograms using 3D magnetoionic ray tracing and atmospheric gravity waves. *Journal of Geophysical Research: Space Physics, 119*, 431–440. https://doi.org/10.1002/2013JA019247

Chum, J., Athieno, R., Baše, J., Burešová, D., Hruška, F., Laštovička, J., et al. (2012). Statistical investigation of horizontal propagation of gravity waves in the ionosphere over Europe and South Africa. *Journal of Geophysical Research, 117*, A03312. https://doi.org/10.1029/2011JA017161

Clark, R. M. (1970). Gravity waves in the ionosphere (Doctoral dissertation). Retrieved from ProQuest Dissertations & Theses Global. (302519386). Urbana, IL: University of Illinois.

Clark, R. M., Yeh, K. C., & Liu, C. H. (1971). Interaction of internal gravity waves with the ionospheric F2-layer. *Journal of Atmospheric and Terrestrial Physics*, *33*(10), 1567–1576. https://doi.org/10.1016/0021-9169(71)90074-2

Colegrove, S. B. (2000). Project Jindalee: From bare bones to operational OTHR. Proceedings of the IEEE International Radar Conference 2000 (pp. 825–830). doi:https://doi.org/10.1109/RADAR.2000.851942

Dao, E. V., McNamara, L. F., & Colman, J. J. (2016). Magnetic field effects on the accuracy of ionospheric mirror models for geolocation. *Radio Science*, *51*, 284–300. https://doi.org/10.1002/2015RS005884

Davies, K. (1990). *Ionospheric radio, IEEE Electromagnetic Waves Series* (Vol. 31). London, UK: Peter Peregrinus Ltd.

Davis, M. J. (1973). The integrated ionospheric response to internal atmospheric gravity waves. *Journal of Atmospheric and Terrestrial Physics*, *35*(5), 929–959. https://doi.org/10.1016/0021-9169(73)90074-3

Drob, D. P., Emmert, J. T., Crowley, G., Picone, J. M., Shepherd, G. G., Skinner, W., et al. (2008). An empirical model of the Earth's horizontal wind fields: HWM07. *Journal of Geophysical Research*, *113*, A12304. https://doi.org/10.1029/2008JA013668

Drob, D. P., Emmert, J. T., Meriwether, J. W., Makela, J. J., Doornbos, E., Conde, M., et al. (2015). An update to the Horizontal Wind Model (HWM): The quiet time thermosphere. *Earth and Space Science*, *2*, 301–319. https://doi.org/10.1002/2014EA000089

Earl, G. F., & Ward, B. D. (1986). Frequency management support for remote sea-state sensing using the JINDALEE skywave radar. *IEEE Journal of Oceanic Engineering*, *OE-11*(2), 164–173.

Earl, G. F., & Ward, B. D. (1987). The frequency management system of the Jindalee over-the-horizon backscatter HF radar. *Radio Science*, *22*(2), 275–291. https://doi.org/10.1029/RS022i002p00275

Fabrizio, G. A. (2013). *High frequency over-the-horizon radar*. New York, NY: McGraw-Hill.

Fridman, S. V., & Nickisch, L. J. (2001). Generalization of ionospheric tomography on diverse data sources: Reconstruction of the three-dimensional ionosphere from simultaneous vertical ionograms, backscatter ionograms, and total electron content data. *Radio Science*, *36*(5), 1129–1139. https://doi.org/10.1029/1999RS002405

Galushko, V. G., Beley, V. S., Koloskov, A. V., Yampolski, Y. M., Paznukhov, V. V., Reinisch, B. W., et al. (2003). Frequency-and-angular HF sounding and ISR diagnostics of TIDs. *Radio Science*, *38*(6), 1102. https://doi.org/10.1029/2002RS002861

Gardiner-Garden, R., Ayliffe, J., Durbridge, L., Frazer, G., Heitmann, A., Praschifka, J., et al. (2008). A new high-fidelity oblique HF receiving system. Paper presented at the Ionospheric Effects Symposium (IES2008), Alexandria, VA.

Gardiner-Garden, R., Heitmann, A., & Frazer, G. (2011). Variability observed in a high-fidelity model of the ionosphere based on a dense network or vertical and oblique incidence sounders. Paper presented at the Ionospheric Effects Symposium (IES2011), Alexandria, VA.

Gossard, E. E. (1962). Vertical flux of energy into the lower ionosphere from internal gravity waves generated in the troposphere. *Journal of Geophysical Research*, *67*(2), 745–757. https://doi.org/10.1029/JZ067i002p00745

Harris, T. J., Cervera, M. A., & Meehan, D. H. (2012). SpICE: A program to study small-scale disturbances in the ionosphere. *Journal of Geophysical Research*, *117*, A06321. https://doi.org/10.1029/2011JA017438

Harris, T. J., & Frazer, G. J. (2005). High resolution radar measurements of ionospheric variations on an oblique path. Paper presented at the Ionospheric Effects Symposium (IES2005), Alexandria, VA.

Hines, C. O. (1960). Internal atmospheric gravity waves at ionospheric heights. *Canadian Journal of Physics*, *38*(11), 1441–1481. https://doi.org/10.1139/p60-150

Hocke, K., & Schlegel, K. (1996). A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995. *Annales Geophysicae*, *14*(9), 917–940. https://doi.org/10.1007/s00585-996-0917-6

Hooke, W. H. (1968). Ionospheric irregularities produced by internal atmospheric gravity waves. *Journal of Atmospheric and Terrestrial Physics*, *30*(5), 795–823. https://doi.org/10.1016/S0021-9169(68)80033-9

Huang, X., Reinisch, B. W., Sales, G. S., Paznukhov, V. V., & Galkin, I. A. (2016). Comparing TID simulations using 3-D ray tracing and mirror reflection. *Radio Science*, *51*, 337–343. https://doi.org/10.1002/2015RS005872

Hunsucker, R. D. (1982). Atmospheric gravity waves generated in the high-latitude ionosphere: A review. *Reviews of Geophysics and Space Physics*, *20*(2), 293–315. https://doi.org/10.1029/RG020i002p00293

Johnson, B. A., Amey, P. F., Braendler, S., & Turley, M. D. E. (2008). In-service enhancements of the Jindalee over-the-horizon radars. Proceedings of the IEEE International Radar Conference 2008 (pp. 218–223). doi:https://doi.org/10.1109/RADAR.2008.4653921

Kelley, M. C., Larsen, M. F., LaHoz, C., & McClure, J. P. (1981). Gravity wave initiation of equatorial spread F: A case study. *Journal of Geophysical Research*, *86*(A11), 9087–9100. https://doi.org/10.1029/JA086iA11p09087

Kirchengast, G. (1996). Elucidation of the physics of the gravity wave-TID relationship with the aid of theoretical simulations. *Journal of Geophysical Research*, *101*(A6), 13353–13368. https://doi.org/10.1029/96JA00750

Lobb, R. J., & Titheridge, J. E. (1977). The effects of travelling ionospheric disturbances on ionograms. *Journal of Atmospheric and Terrestrial Physics*, *39*(2), 129–138. https://doi.org/10.1016/0021-9169(77)90106-4

Lynn, K. J. W., Gardiner-Garden, R. S., & Heitmann, A. (2016). The observed compression and expansion of the F_2 ionosphere as a major component of ionospheric variability. *Radio Science*, *51*, 538–552. https://doi.org/10.1002/2016RS006036

Lynn, K. J. W., Otsuka, Y., & Shiokawa, K. (2013). Ionogram-based range-time displays for observing relationships between ionosonde satellite traces, spread F and drifting optical plasma depletions. *Journal of Atmospheric and Solar-Terrestrial Physics*, *98*(6), 105–112. https://doi.org/10.1016/j.jastp.2013.03.020

Lyon, G. F. (1979). The corrugated reflector model for one-hop oblique propagation. *Journal of Atmospheric and Terrestrial Physics*, *41*(1), 5–9. https://doi.org/10.1016/0021-9169(79)90041-2

McNamara, L. F. (1991). *The ionosphere: Communications, surveillance, and direction finding*. Malabar, FL: Krieger Publishing Company.

Munro, G. H. (1950). Travelling disturbances in the ionosphere. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, *202*(1069), 208–223. https://doi.org/10.1098/rspa.1950.0095

Munro, G. H., & Heisler, L. H. (1956). Cusp type anomalies in variable frequency ionospheric records. *Australian Journal of Physics*, *9*(3), 343–358. https://doi.org/10.1071/PH560343

Nickisch, L. J., Hausman, M. A., & Fridman, S. (2007). Traveling ionospheric disturbance mitigation for OTH radar. Proceedings of the IEEE International Radar Conference 2007 (pp. 362–366). doi:https://doi.org/10.1109/RADAR.2007.374243

Nickisch, L. J., Hausman, M. A., & Fridman, S. V. (2006). Range rate-Doppler correlation for HF propagation in traveling ionospheric disturbance environments. *Radio Science*, *41*, RS5S39. https://doi.org/10.1029/2005RS003358

Paznukhov, V. V., Galushko, V. G., & Reinisch, B. W. (2012). Digisonde observations of TIDs with frequency and angular sounding technique. *Advances in Space Research*, *49*(4), 700–710. https://doi.org/10.1016/j.asr.2011.11.012

Reinisch, B., Galkin, I., Belehaki, A., Paznukhov, V., Huang, X., Altadill, D., et al. (2018). Pilot ionosonde network for identification of traveling ionospheric disturbances. *Radio Science*, *53*, 365–378. https://doi.org/10.1002/2017RS006263

Reinisch, B. W., Galkin, I. A., Khmyrov, G. M., Kozlov, A. V., Bibl, K., Lisysyan, I. A., et al. (2009). New Digisonde for research and monitoring applications. *Radio Science*, *44*, RS0A24. https://doi.org/10.1029/2008RS004115

Rice, D. W. (1973). Phase characteristics of ionospherically propagated radio waves. *Nature Physical Science*, *244*(136), 86–88.

Richmond, A. D. (1978). Gravity wave generation, propagation, and dissipation in the thermosphere. *Journal of Geophysical Research*, *83*(A9), 4131–4145. https://doi.org/10.1029/JA083iA09p04131

Rietveld, M. T., Wright, J. W., Zabotin, N., & Pitteway, M. L. V. (2008). The Tromsø dynasonde. *Polar Science*, *2*(1), 55–71. https://doi.org/10.1016/j.polar.2008.02.001

Rogers, N. C., Warrington, E. M., & Jones, T. B. (2003). Oblique ionogram features associated with off-great circle HF propagation at high and sub-auroral latitudes. *IEE Proceedings – Microwaves, Antennas and Propagation*, *150*(4), 295–300. https://doi.org/10.1049/ip-map:20030552

Sherrill, W. M., & Brown, B. C. (1997). Azimuth spread of two-hop ionospheric modes. *Radio Science*, *32*(5), 1875–1880. https://doi.org/10.1029/97RS01267

Sherrill, W. M., & Smith, G. A. (1977). Directional dispersion of sporadic-E observations between 9 and 14 MHz. *Radio Science*, *12*(5), 773–778. https://doi.org/10.1029/RS012i005p00773

Shiokawa, K., Otsuka, Y., & Ogawa, T. (2009). Propagation characteristics of nighttime mesospheric and thermospheric waves observed by optical mesosphere thermosphere imagers at middle and low latitudes. *Earth, Planets and Space*, *61*(4), 479–491. https://doi.org/10.1186/BF03353165

Sinnott, D. H. (2015). Over the horizon radar down-under. Proceedings of the IEEE International Radar Conference 2015 (pp. 1761–1764). doi: https://doi.org/10.1109/RADAR.2015.7131284

Sweeney, L. E., Jr. (1970). Spatial properties of ionospheric radio propagation as determined with half-degree azimuthal resolution (Technical Report No. 155, SU-SEL-70-034). Stanford, CA: Stanford Electronics Laboratories.

Vertogradov, G. G., Uryadov, V. P., Vertogradova, E. G., Vertogradov, V. G., & Kubatko, S. V. (2013). Chirp ionosonde-radiodirection finder as a new tool for studying the ionosphere and radio-wave propagation. *Radiophysics and Quantum Electronics*, *56*(5), 259–275. https://doi.org/10.1007/s11141-013-9431-3

Wright, J. W., & Kressman, R. I. (1983). First bistatic oblique-incidence ionograms between digital ionosondes. *Radio Science*, *18*(4), 608–614. https://doi.org/10.1029/RS018i004p00608

F.2   Ayliffe et al., *Radio Science*, 2019

The following is a reprint of Ayliffe, J. K., Durbridge, L. J., Frazer, G. J., Gardiner-Garden, R. S., Heitmann, A. J., Praschifka, J., et al. (2019), 'The DST Group high-fidelity, multi-channel oblique incidence ionosonde', *Radio Science*, 54(1), pp. 104–114.

This paper summarises the hardware and software components of DST Group's Digital Oblique Receiving System (DORS) OIS, which were later adapted for the ELOISE AoA ionosonde. It thus provides relevant background material to supplement the descriptions in **Chapters 3** and **4** of this thesis.

The manuscript was written by myself, as the corresponding author, and covers a program of work carried out by both myself and others at DST Group. My own contributions to the work are mainly in the area of the on-board signal and image processing. The paper was published during my PhD candidature.

RESEARCH ARTICLE

Key Points:
- Defence Science and Technology Group has developed a direct-digital ionosonde for oblique incidence observations of the ionosphere
- A suite of onboard software supports flexible command and control, sophisticated signal and image processing, and data management tools
- The output ionograms and automatically scaled parameters are of high quality, and are now used routinely in regional ionospheric models

The DST Group High-Fidelity, Multichannel Oblique Incidence Ionosonde

Jim K. Ayliffe[1], Lyndon J. Durbridge[1], Gordon J. Frazer[1], Robert S. Gardiner-Garden[2] [ID], Andrew J. Heitmann[1,3] [ID], Justin Praschifka[1], Adrian D. Quinn[1], Gavin C. Scarman[1], and Mike D. E. Turley[1]

[1]Defence Science and Technology Group, Edinburgh, South Australia, Australia, [2]Defence Science and Technology Group, Eveleigh, New South Wales, Australia, [3]School of Physical Sciences, University of Adelaide, Adelaide, South Australia, Australia

Abstract Defence Science and Technology Group has developed a complete oblique incidence ionosonde system to enhance ionospheric monitoring capabilities in support of the Jindalee Operational Radar Network. The Digital Oblique Receiving System, and its counterpart, the Digital Oblique Transmitting System, feature a direct-digital, multichannel chirp-sounder design that allows flexible collection of high-resolution ionograms across many simultaneous paths, without compromising dynamic range. This paper outlines the hardware and software components of the Digital Oblique Receiving System/Digital Oblique Transmitting System, including the onboard signal processing that removes radio frequency interference and fits a parameterized electron density profile to the ionogram. Additional control and data management tools that have enabled the system to be effectively transitioned into routine operations are also described. The typical performance is illustrated through a selection of sample ionograms, highlighting the applications of Digital Oblique Receiving System/Digital Oblique Transmitting System to a wide variety of ionospheric studies, including those requiring Doppler and angle-of-arrival observations.

Plain Language Summary An oblique incidence ionospheric sounder (or ionosonde) records the ionospheric propagation characteristics for high-frequency radio waves on one-way over-the-horizon paths of up to 3,000 km or more. Observations are typically presented in the form of an ionogram image, a record of the received signal power (or path loss) as a function of frequency and delay. Defence Science and Technology Group has recently developed both transmitting and receiving subsystems that enable high-quality ionograms to be produced simultaneously on multiple paths. The sophisticated hardware and software components that underpin this system are described in this paper, and a selection of sample ionograms illustrates the typical performance and applications to ionospheric studies.

1. Introduction

Ground-based high-frequency (HF) ionospheric sounders have long been used to observe the vertical structure of the bottomside ionosphere, dating back to the pioneering experiments of Breit and Tuve (1925, 1926). Recent advances in technology have enabled a new generation of digital ionosondes to be developed, at an unprecedented level of sensitivity, and with an increased capacity for sophisticated signal and image processing. Notable examples from the last two decades include the Digisonde DPS-4D (Reinisch et al., 2009), VIPIR/Dynasonde (Grubb et al., 2008; Rietveld et al., 2008), AIS-INGV (Zuccheretti et al., 2003), IRIS (Arthur et al., 1997), Cyclone (Akchurin et al., 2009), and WMISS (Gong et al., 2016).

This paper describes a high-fidelity, multichannel system that has now gained extensive use around Australia as an oblique incidence sounder (OIS) within the Jindalee Operational Radar Network (JORN). JORN is composed of three over-the-horizon radars that provide wide area surveillance of Australia's northern approaches (Colegrove, 2000; Fabrizio, 2013). The new Digital Oblique Receiving System (DORS)/Digital Oblique Transmitting System (DOTS) ionosonde from Defence Science and Technology (DST) Group represents an evolution of the chirp-sounder technique originally pioneered by Barry and Fenwick at Stanford University in the 1960s (Fenwick & Barry, 1966), and more recently used by IRIS, WMISS, and the

Russian chirp-sounder network (Ivanov et al., 2003). With its 100% duty cycle, and linear frequency-modulated continuous-wave transmissions, only about 20 W of transmitted power is required; this makes the DST Group system (and chirp sounders in general) more spectrum-friendly than pulsed alternatives. The DST Group transmitter and receiver subsystems are also compatible with other chirp sounders. For interested readers, a comprehensive historical overview of the chirp sounder is described in Ivanov et al. (2003).

Following some background information on the development of the ionosonde and where it fits into Australia's defense program (section 2), separate overviews will be provided for the hardware (section 3) and software (section 4) components of DORS/DOTS. A small selection of sample ionograms and fitted traces from various DST Group experiments over the last decade is included to illustrate the sounder's capabilities and quality of results. Specific applications and the physical interpretation of sample results will be discussed further in section 5. The paper finishes by discussing future development (section 6) and conclusions (section 7).

2. Background

While the majority of the world's ionosondes have historically been single-site installations, configured to run as vertical incidence sounders (VIS), the use of OIS networks has a number of advantages. First, for large-scale ionospheric mapping purposes, their many combinations of point-to-point propagation paths permit a greater spatial sampling density than can be offered by a VIS network alone, given the same number of sites. Second, they may be used to observe the ionosphere at locations where a VIS is impractical (e.g., over oceans). Finally, where transmitters and receivers are not co-located, their design and installation is simplified considerably; that is, system-generated noise and emissions remain minimal at the receiver site and 100% duty cycle waveforms can be used to reduce peak power demands. Modern GNSS networks make OIS network synchronization considerably easier and more accurate than previous methods such as internal clocks or HF timing signals.

For these reasons, the Australian Department of Defence is currently in the process of expanding its OIS network in northern Australia to include up to 40 oblique propagation paths, in addition to its existing VIS network of 13 sites. This so-called Dense Integrated Network of Ionospheric Sounders (DINIS; Gardiner-Garden et al., 2011) will enhance the future frequency advice and coordinate registration capabilities of the JORN HF sky-wave radars. DINIS has also addressed obsolescence issues with the legacy JORN VIS and OIS systems, many of which had reached their end of life after continuous operation for around two decades.

Developed with the goals of DINIS in mind, the Digital Oblique Receiving System (DORS) and its counterpart, the Digital Oblique Transmitting System (DOTS), are a complete OIS solution, designed and built by Defence Science and Technology (DST) Group in Edinburgh, South Australia. The systems use state-of-the-art direct-digital receivers and waveform generators, and are accompanied by an extensive suite of onboard software and support tools, which are each summarized in the following sections. DORS and DOTS are closely related to DST Group's Portable Remote Ionospheric Monitoring Equipment (PRIME; Harris et al., 2016), a customized form of the same underlying receiver and waveform generator technology, adapted for VIS reception and analysis.

Key attributes of the DORS/DOTS ionosonde are as follows:

1. High sensitivity and fidelity, with large spurious free dynamic ranges and relatively flat gain responses.
2. Multiple simultaneous channels on a single receiver.
3. Flexible control and data acquisition, including compatibility with legacy chirp sounders.
4. Reliable and cost-effective as an OIS network element.
5. Robust onboard software for interference rejection, ionogram processing, and automatic scaling, producing ionograms and ionospheric parameters of very high quality.

Development of the sounder has now progressed from an early prototype used for short-term deployments (Gardiner-Garden et al., 2008) to a fully operational system, running as part of the JORN OIS network since 2011. Since 2017, the expanded DINIS network, comprising four DORS and seven DOTS units for OIS ionogram collection, along with a combination of Lowell Digisonde DPS-1 and DST Group PRIME units for
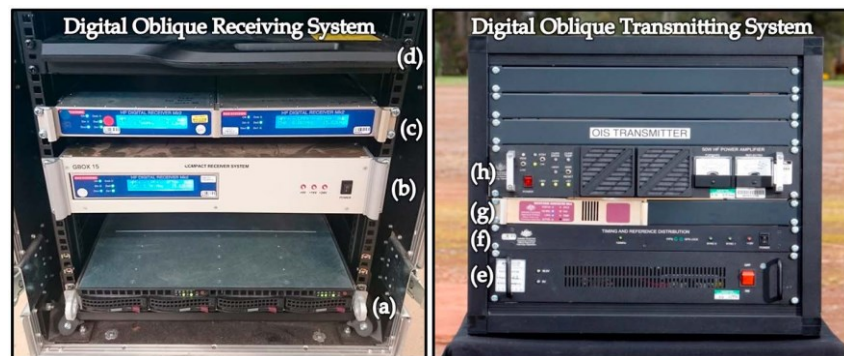
Figure 1. (left) Digital Oblique Receiving System (DORS). (right) Digital Oblique Transmitting System (DOTS). Each consists of several rack-mounted components, including (a) server, (b) GBOX, (c) additional DRx, (d) KVM drawer, (e) DRx/DWG power supply, (f) TRDU, (g) DWG, and (h) 50-W HF power amplifier. Note that the DOTS server/laptop is not pictured here.

VIS reception, has been routinely feeding parameterized data into the JORN Real-Time Ionospheric Model. This is a regional representation of the ionosphere over northern Australia, constructed by spatially filtering and mapping the ionogram profile parameters (Barnes et al., 2000). Partner organizations in Australia are now also making use of the DORS/DOTS ionosonde and its data products.

3. Hardware Overview

The DORS and DOTS ionosonde systems are modular in design, and therefore very flexible and maintainable. Each consists of a subset of the following components, most of which are rack-mountable. Typical assemblies are pictured in Figure 1.

1. A general purpose Linux (RedHat/CentOS) server or laptop, for sounder control, data acquisition, and processing.
2. Time and frequency reference distribution (TRDU) unit, with built-in GPS receiver (Jackson Labs Fury OEM) and GPS-disciplined 100-MHz reference (multiplier circuitry developed by DST Group).
3. Three-channel HF digital receiver (DRx) and power supply (DORS only).
4. Three-channel HF digital waveform generator (DWG) and power supply (DOTS only).
5. A 50-W HF power amplifier (DOTS only).
6. Network router, data modem, and keyboard-video-mouse (KVM) drawer.
7. GPS and HF antennas (with RF splitter for multiple DRx).

A compact stand-alone DRx package, with its own onboard power supply and time and frequency reference distribution unit, has also been developed; this is known colloquially as a "GBOX." It requires only a server or laptop for sounder observations on up to three simultaneous channels, and can power an additional two DRx (six channels) if required. This same stand-alone design has been adopted for the more modern NexGenR receiver, described further in section 6.

While in many cases existing antenna infrastructure can be used, DST Group has developed its own portable, broadband HF antenna suitable for short-term sounder deployment on experimental trials. The so-called "J-dart" is pictured in the left panel of Figure 2, alongside an RFS WM-245 biconical antenna typical of the OIS transmit antennas used within JORN. Both are vertically polarized, omnidirectional monopole antennas capable of operating over a broad band of frequencies, with good low to medium angle elevation performance. These are well suited to both transmission and reception on OIS paths between about 500 and 3,000 km, with typical elevation angles spanning 1 to 60° for one-hop F_2 layer propagation. However, as the gain patterns of such antennas do roll-off rapidly for elevations below 10°, reduced F_2 mode sensitivity is to be expected on those longer paths of more than 2,000 km.

The HF digital receiver is the DST Group/BAE Systems DRx MkIID, a high-fidelity wideband receiver developed specifically for HF radar applications (Frazer, 2007). The DRx MkIID features a relatively low-cost, direct-digital
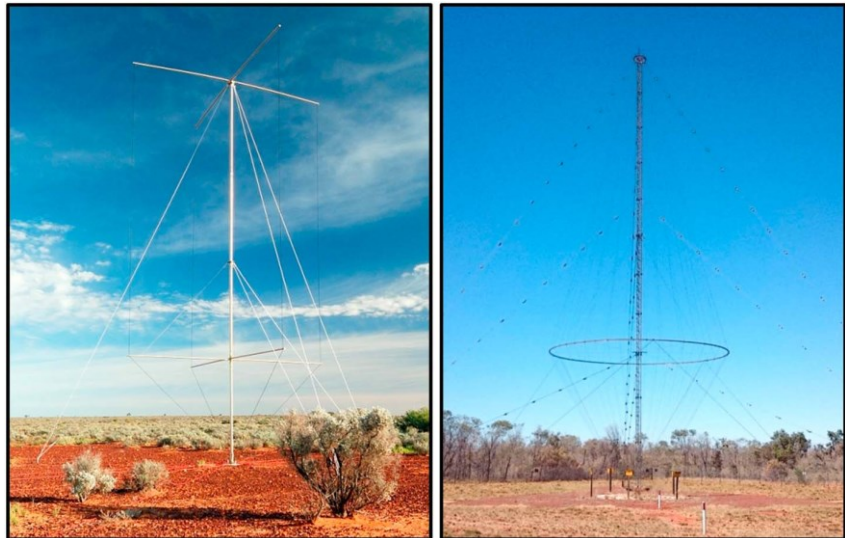
Figure 2. The typical HF antennas used for DORS/DOTS oblique soundings. (left) DST Group J-dart portable antenna. (right) RFS WM-245 biconical antenna.

design, with 14-bit analog-to-digital (A/D) converter running at a clock speed of 100 MHz. The default high-pass (broadcast band) filter and low-pass (antialiasing) filter result in an operating passband of 5–35 MHz, although modified front-end filters do exist for lower frequency reception on shorter sounder paths. The internal 32-bit data processing, implemented on a field-programmable gate array, includes three independent digital down-converter channels with sweeping numerically controlled oscillators. A general description of the down-conversion (or stretch processing) technique can be found in Fabrizio (2013).

At the output of the DRx, in-phase and quadrature (I/Q) samples are decimated to a sample rate of 15.625 kHz, giving a usable bandwidth of 12.5 kHz. Across the 5–35-MHz passband, the gain profile is flat to within ±1 dB. The resultant spurious free dynamic range has been measured to be in excess of 150 dBc across the passband under typical operating conditions, with a receiver noise figure of 7–8 dB for up to 10-dB attenuation (after which it increases at a rate approximately equal to the increase in attenuation). Timing accuracy is better than 1 μs, using the GPS-locked 1PPS signal.

The DST Group HF digital waveform generator is based on a similar design and shares common firmware with the receiver. It is also a direct-digital multichannel system, with a 16-bit digital-to-analog converter running at 100 MHz. The adjustable gain allows for RF output up to +13 dBm, with a spurious free dynamic range of 78 dBc across the 5–35-MHz passband. The DWG is used to drive a 50-W HF power amplifier, which is routinely configured to provide a nominal 20 W (or 13 ± 2 dBW, subject to antenna mismatch) for OIS soundings.

All hardware is equipped with temperature sensors and appropriate lightning protection, to cope with the hostile operating conditions of many remote sounder locations. A transmitter/receiver test harness at DST Group enables hardware calibration and diagnosis of common faults, such as failures in the analog front-end circuitry.

4. Software Overview

4.1. Command and Control

The DRx or DWG is controlled by a Linux server or laptop, which initiates each sounding, selects the gain settings, and logs the device state and temperature. A system monitor display provides a high-level interface for managing command and control, including restarting receivers or control processes, and an onboard iono-gram viewer enables real-time data inspection and replay. For DORS, the output I/Q samples are also dumped

Table 1
Standard Sounder Schedule Parameters Used by the JORN OIS Network

Sweep limits	2 to 45 MHz
Sweep rate	250 kHz/s
Total duration	172 s
Revisit interval	225 s

Note. Although the sounder routinely sweeps from 2 to 45 MHz, the usable dynamic range is reduced outside the 5–35-MHz standard DRx/DWG passband.

to a raw file, ready for subsequent onboard processing (described in the later sections of this paper). Connectivity is provided via Gigabit Ethernet.

A predefined schedule file controls the fixed start times, sweep rates, and transmitter-receiver delay offsets of the soundings, as well as specifying the path coordinates. A schedule editor ("schedit") provides a simple user interface to help populate this file. The linear frequency–modulated continuous-wave sweep rate is fully adjustable, although rates commonly chosen are 125, 156.25, 250, and 312.5 kHz/s, being integer multiples of the 15.625-kHz output sample rate for the DRx. Faster sweep rates reduce the revisit interval and delay-Doppler coupling effects (i.e., the ambiguity caused by ionospheric Doppler shifts on chirp-sounder delay measurements), but also reduce the signal-to-noise ratio (SNR) for the same transmitted power. Standard JORN OIS schedule parameters are listed in Table 1. It should be noted that none of these parameters are fixed and can be customized as required.

Except on very short paths, DORS OIS ionograms typically only retain a 4-ms group delay window, so additional decimation is usually applied to reduce raw data volumes before writing to disk. For example, at a sweep rate of 250 kHz/s, 10X decimation reduces the sample rate from 15.625 to 1.5625 kHz, leaving 1.25-kHz usable bandwidth or 5-ms usable group delay. The transmitter-receiver delay offset is chosen to preserve a small guard band of negative delays (less than the great circle ground range) for background noise estimation.

The automatic gain control function of the control software manages the dynamic range of the receiver against diurnal changes in the background RF environment. It adjusts the DRx attenuation setting whenever the system is inactive (between soundings), to preserve a fixed headroom and protect against overloads. As the A/D converter is open to the whole HF band, there is some compromise between maintaining linearity and operating above the internal noise limit. With this in mind, a target A/D headroom of 6 dB is the standard DORS configuration, resulting in a very small number of DRx overloads, but remaining externally noise limited at most operational frequencies within the passband (5–35 MHz).

To quantify the above statement, median noise levels at an exceptionally quiet Australian receiving site near Alice Springs were analyzed for the month of September 2015. Across six medium-length OIS paths (700–1,600 km), less than 0.1% of ionogram frequency bins within the passband and below the maximum observed frequency (MOF) of F_2 propagation contained noise within 3 dB of the internal noise limit. In other words, the vast majority of these observations (>99.9%) were deemed to be externally noise limited. Above the MOF, the time-averaged percentage of internally noise limited frequencies rose to 10–20%; however, this is of lesser consequence for ionogram mode detection and scaling. For other sites with higher external HF noise environments, the percentages of internally noise limited observations are reduced, so Alice Springs is believed to represent a worst-case scenario.

4.2. Ionogram Signal Processing

The fundamental data product from each ionospheric sounding is the ionogram: an image of the received power (in dBW), across all propagated modes, as a function of frequency and group delay. Given the downconverted I/Q samples from a single DORS receiver channel, the ionogram is formed by simply performing spectral analysis on overlapped blocks of the coherent data sequence. A given frequency component in the data is equal to the product of the ionospheric group delay and the sweep rate, plus any ionospheric Doppler shift (usually <1 Hz or, equivalently, <1 km in delay at 250 kHz/s). In practice, the FFT algorithm is used, with a Hann window applied for modest sidelobe reduction (Harris, 1978). An 80% overlap is chosen to produce smoothly connected ionogram traces in the presence of windowing; this greatly assists subsequent scaling and interpretation. Table 2 lists the standard JORN OIS ionogram characteristics.

Prior to the ionogram calculation, an autoregressive (AR) interpolation technique is used to mitigate against narrowband radio frequency interference (RFI) in the I/Q samples (Turley et al., 2019). It first detects

Table 2
Standard Ionogram Characteristics Used by the JORN OIS Network, for the Schedule Parameters in Table 1

Coherent processing time/bandwidth	0.8 s/200 kHz
Ionogram dimensions	1073 pixels (frequency) × 801 pixels (group delay)
Frequency limits	2.00 to 44.88 MHz
Group delay limits (with respect to great circle ground range)	−0.4 to 3.6 ms (−120 to 1080 km)
Frequency step size	40 kHz (80% overlap)
Group delay step size	5 μs (1.5 km)
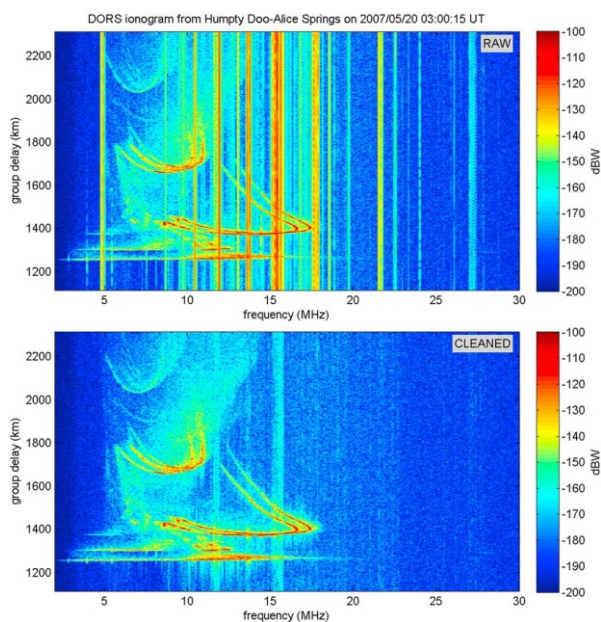
Figure 3. Comparison of DORS ionograms (top) before and (bottom) after RFI rejection (cf. Turley et al., 2019, Figure 1). Impulsive RFI in the I/Q samples manifests as the vertical striations in the top image.

additive impulsive noise in the data, using a low-order AR-model-based finite impulse response filter, then replaces the bad samples with a higher-order AR model of the desired signal (i.e., representing the ionospherically propagated modes from the sounder transmitter). In typical cases, up to 40–60 dB of SNR can be recovered in the vicinity of RFI, and the trace continuity that would otherwise be broken into discontinuous segments is maintained. Figure 3 shows corresponding ionograms before and after RFI rejection; the clarity of ionospheric trace features is improved markedly by this preliminary processing step.

Additional ionogram preprocessing has also been developed to improve compatibility with other chirp sounders. This includes compensation for known timing biases (e.g., a 33-μs timing offset between the TCI/BR Communications TCS-7 transmitter and DORS receiver) and AR interpolation across known transmitter dropouts (often on the scale of tens of milliseconds, associated with mechanical filter switching). Longer dropouts, on the same scale as the coherent processing time/bandwidth (e.g., due to barred frequencies), cannot be effectively remedied.

To highlight the combined improvements of state-of-art hardware and signal processing on the final ionogram, Figure 4 compares a DORS ionogram (after RFI rejection) with an older-generation TCI/BR Communications RCS-7 ionogram, received simultaneously on the same path and antenna. As well as the increased image resolution in both frequency and group delay, which reveals finer and higher-order ionospheric mode structure, the DORS ionogram also contains much smoother traces (due to overlapped processing), with better SNR (due to RFI rejection), and far fewer spectral artifacts (due to the direct-digital receiver design having minimal analog circuitry). One such artifact is the spurious "ghost" trace appearing in the left panel, but absent in the right panel, which replicates the true one-hop F_2 trace (at around 10–15 MHz and 900–1,000 km) with an extra delay of 120 km; it is noted that this offset corresponds to the 50-Hz grid power frequency in Australia.

4.3. Feature Extraction and Profile Fitting

While the high-quality DORS ionograms provide a rich characterization of the ionospheric propagation environment, in many modeling applications it is desirable to reduce these observations down to a more simple parameterized description of the electron density profile. DORS has its own robust, heuristic technique for automatic scaling that runs onboard the sounder, known as the DST-IIP (Ionogram Image Processing) algorithm. It is based on extracting a limited set of key image features relating to the E, sporadic E (E_s), and F_2 layers, then fitting a multi-segment quasi-parabolic profile (Dyson & Bennett, 1988), with embedded mirror reflectors for up to two E_s layers, that reproduces the features through analytic trace synthesis. A spherically symmetric ionosphere at the path midpoint is assumed, and climatological models for E and F_1 are used to complete the profile where critical information is unable to be extracted from the image itself. A full description of the algorithm can be found in Heitmann and Gardiner-Garden (2019).

Figure 4. A daytime sounding from Kalkarindji to Alice Springs (736 km) in northern Australia. (left panel) The ionogram produced by an older-generation chirp-sounder receiver (TCI/BR communications RCS-7). (right panel) The corresponding DORS ionogram, received simultaneously on the same antenna.

The primary outputs of the DST-IIP algorithm are the midpoint electron density profile (fitted up to the F_2 peak) and the associated synthetic traces for each of the E, F_1, F_2, and E_s layers. These synthetic traces are shown overplotted on a sample DORS ionogram in Figure 5. Only the one-hop returns are strictly fitted, but under relatively benign ionospheric conditions, the two-hop synthetic traces can also match the image remarkably well. Note that a first-order model for geomagnetic effects (Bennett

Figure 5. A day-time DORS ionogram from Humpty Doo to Alice Springs (1,234 km) in northern Australia. Overplotted are the synthetic traces, based on a parameterized profile automatically fitted to the one-hop image features (cf. Heitmann & Gardiner-Garden, 2019, Figure 4).

et al., 1991; Chen et al., 1992) is sufficient to represent the bulk of the ordinary (o) and extraordinary (x) mode splitting, without requiring explicit separation of polarization components in the ionogram.

The analysis of tens of millions of DORS ionograms since 2007 has revealed that, under good propagation conditions (>90% of medium-length observations), the automatic results are robust and reliable. The estimation of feature and fit uncertainty is an important part of the algorithm design, and critical to the subsequent use of parameters in data-assimilative ionospheric models. The F_2 critical frequency, for example, is typically accurate to within ±0.1 MHz. The challenge of producing satisfactory output in the presence of strong midlatitude spread F (e.g., Bowman, 1990), intense multihop/blanketing E_s (e.g., Whitehead, 1989), depressed F_2 ionization during storm conditions (e.g., Mendillo, 2006), and large trace perturbations caused by traveling ionospheric disturbances (e.g., Hocke & Schlegel, 1996) is the subject of ongoing algorithm development.

4.4. Data Management and Network Displays

For each sounding, both a raw data file and a processed data file are generated onboard the DORS server. The former contains the I/Q samples on the down-converted and decimated channel, while the latter contains the ionogram (typically byte-scaled and thresholded) along with its automatically scaled features and fitted parameters. Average file sizes tend to be around 4 MB and 70 kB, respectively. In routine JORN operations, only the processed data files are transferred to a central data server, where they undergo further analysis in constructing the Real-Time Ionospheric Model. Allowing for the typical 20–30 s of onboard processing time, ionograms can be transferred to the central server well within a minute of the end of the sounding observation.

The Sounder Health Display, shown in Figure 6, is a tool created for monitoring the incoming DORS OIS ionograms (and indeed any other OIS/VIS ionograms) at a network level. It offers a number of different time
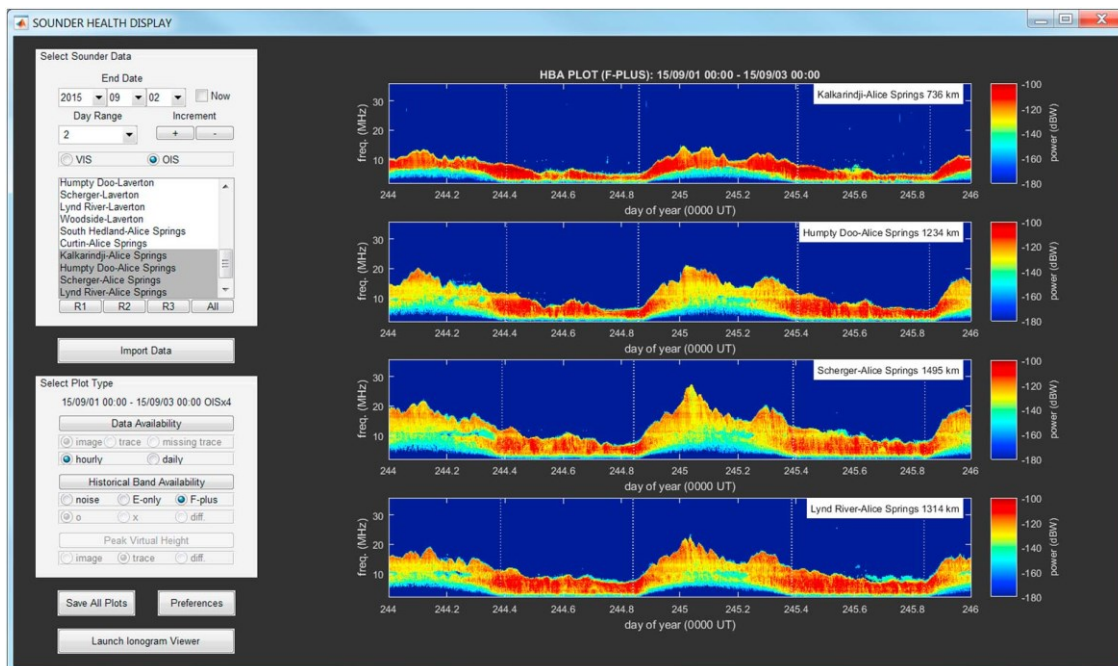
Figure 6. A screenshot of the Sounder Health Display tool, which is used to inspect the quality of the incoming DORS ionograms at a network level.
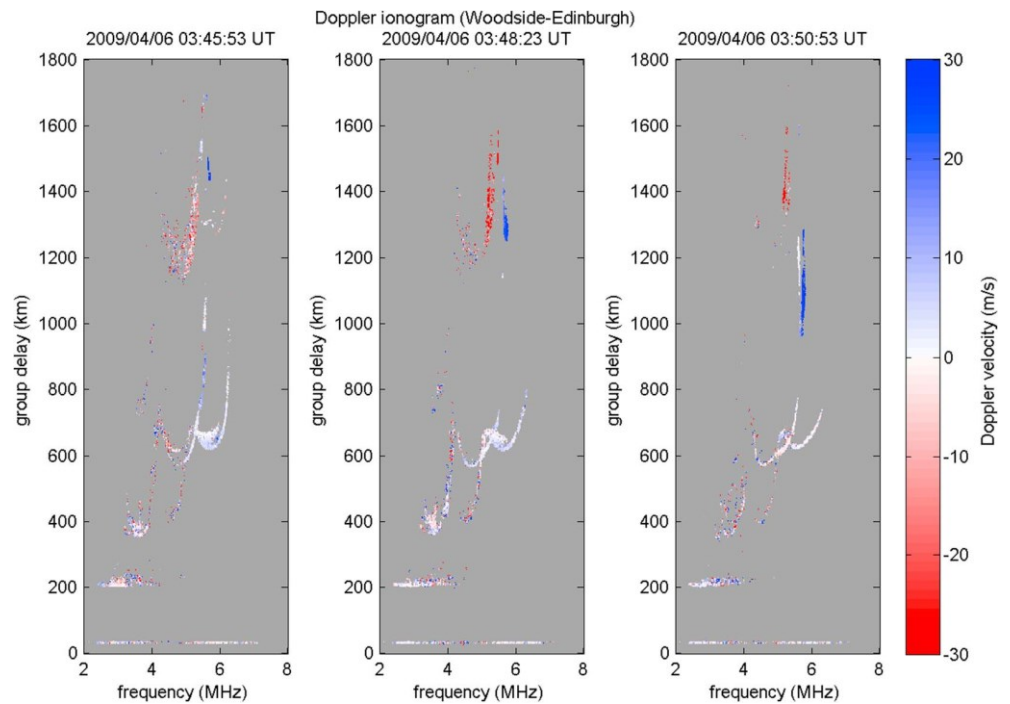
Figure 7. A 5-min sequence of DORS ionograms color coded by Doppler velocity, for a near-vertical incidence path from Woodside to Edinburgh (34 km), near Adelaide in South Australia. The descending F_2 trace feature, appearing at group delays between 1,000 and 1,500 km and with a positive (blue) Doppler velocity, is characteristic of a medium-scale traveling ionospheric disturbance.
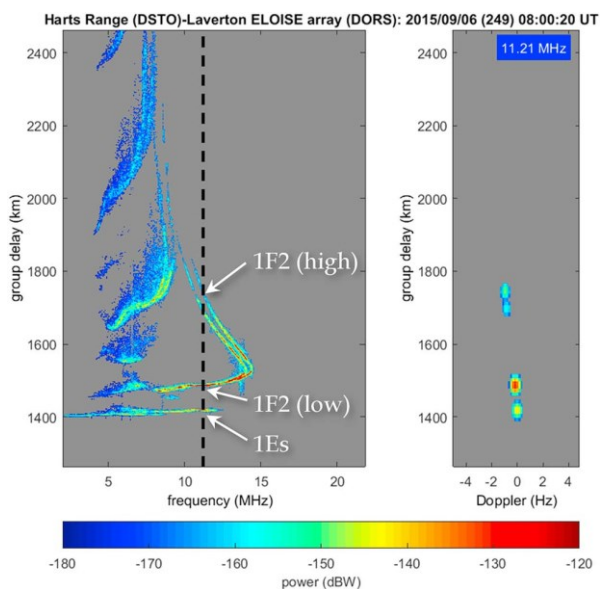
Figure 8. (left) A sample DORS ionogram and (right) corresponding CSF range-Doppler map at 11.21 MHz, for a 1,383-km OIS path from Harts Range to Laverton, in central Australia. The key one-hop propagation modes at the CSF frequency (dashed black line) are annotated on the ionogram.

series representations of the data, including background noise and peak power at each frequency, that highlight any abnormal behavior or temporal performance variation in the sounder that may not be obvious from individual ionograms. Besides its maintenance role, the tool is also indispensable for reviewing propagation conditions, particularly the network-wide response to ionospheric storms and disturbances. A coupled ionogram viewer also permits manual inspection of the synthetic traces fitted automatically by the DORS onboard processing.

4.5. Auxiliary Functions

In addition to the core sounder functionality (i.e., the generation of iono-grams and fitted parameters), a number of auxiliary functions have also been developed by DST Group for special experimental applications. The precise measurement of phase data is critical to many of these functions. They include the following:

1. The collection of ionospheric Doppler information; either Doppler ionograms, using a replicated wideband sounder waveform offset in time, or full Doppler spectra at a discrete set of frequencies, using a repetitive narrowband waveform (so-called channel scattering functions). A sequence of DORS Doppler ionograms is shown in Figure 7, while an example of a DORS channel scattering function (CSF) dwell is shown alongside the ionogram in Figure 8. The interpretation of these Doppler observations will be discussed further in section 5.
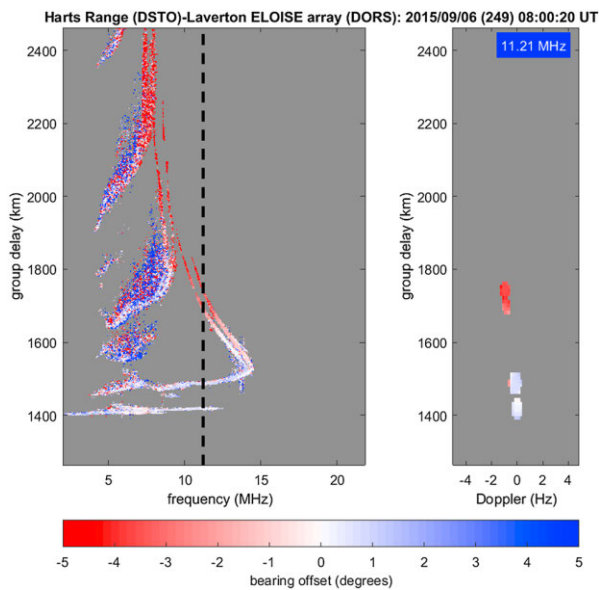
Figure 9. Estimates of bearing offset, with respect to the great circle path geometry, for the same ionogram and CSF images as in Figure 8.

2. The measurement of angle-of-arrival on reception, using 2-D HF arrays to estimate both bearing and elevation. An example of a DORS ionogram, color-scaled according to the bearing offset from great circle geometry, is shown in Figure 9. A 19-element twin-arm array, with a DORS receiver per element, was used in this case (Heitmann et al., 2018).

3. A passive spectrum monitoring tool, which can run alongside sounder operations on one of the independent down-converter channels. It features a waterfall display, showing the time history of spectra, to better identify other HF users and assess channel occupancy.

5. Applications and Further Examples

The DORS/DOTS ionosonde has been used for many years in HF propagation studies by DST Group and is now a key operational component of the JORN sounder network. Its excellent dynamic range and high-resolution ionograms enable very fine mode structure to be observed, while the flexible scheduling and control allows for the study of ionospheric dynamics on shorter time scales than traditionally recorded (including subminute revisit rates) and with optional Doppler information. The hardware and antennas are also relatively portable, making DORS/DOTS well suited to short-term experimental campaigns. Recent DST Group projects to use DORS/DOTS and its onboard processing capabilities have included the deployment of ionosonde networks to study ionospheric spatial variability (Gardiner-Garden et al., 2011), oblique angle-of-arrival reception (Heitmann et al., 2018), mode-selective radar trials (Abramovich et al., 2013), and HF geolocation experiments (Fabrizio & Heitmann, 2013).

As well as the climatological patterns of ionospheric layer variability that are captured in long-term collections of fitted profile parameters, the DORS ionograms themselves reveal many interesting ionospheric phenomena and disturbance signatures that are deserving of further study. In the predominantly midlatitude paths observed to date, this includes the manifestation of traveling ionospheric disturbances, spread F irregularities, and tidal ion layer (sporadic E) morphology.

For example, the high-resolution ionograms in Figure 4 (right panel) and Figure 10 are able to separate fine mode structure in the two-hop E_s/F_2 combination modes (at around 1,000–1,200 km) and one-hop E_s modes (at around 1,250–1,350 km), respectively. The multitude of apparent E_s reflectors in Figure 10 is thought to be due to the presence of patchy ionization at two or more physical E_s layer heights, forming multiple intraionospheric reflections between the different heights. Furthermore, the large dynamic range of the receiver permits weaker modes, associated with micromultipath effects or off-angle paths, to be frequently detected. An extreme example, during postmidnight spread F conditions, is shown in Figure 11. The weaker (off-angle) one-hop F_2 returns extend 8 MHz beyond what appears to be the maximum frequency cutoff for the great circle path at 10 MHz.

Although not currently in routine use, the Doppler capability of DORS/DOTS has been an essential part of recent DST Group experiments looking at ionospheric dynamics. Figure 7 shows the characteristic descending "kink" signature of a traveling ionospheric disturbance in a sequence of DORS Doppler ionograms. Such disturbances are typically the manifestation of atmospheric gravity waves in the thermosphere, causing wave-like perturbations in the electron density that propagate horizontally across the field of view with a forward-tilted phase-front (Hocke & Schlegel, 1996; Hunsucker, 1982). Just as ubiquitous is the evening terminator behavior in Figures 8 and 9, whereby the ionospheric Doppler shifts of the $1F_2$ (low) and $1F_2$ (high) modes in the channel
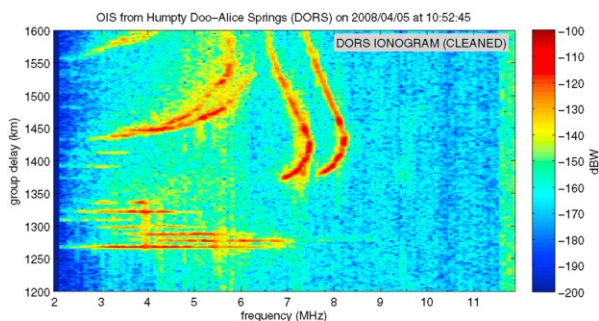
Figure 10. An enlarged DORS ionogram for a 1,234-km OIS path from Humpty Doo to Alice Springs, showing a multitude of apparent sporadic E (E_s) reflectors. Many of these are presumed to be combination, multihop, and/or off-angle modes, derived from at least two distinct E_s layer heights.
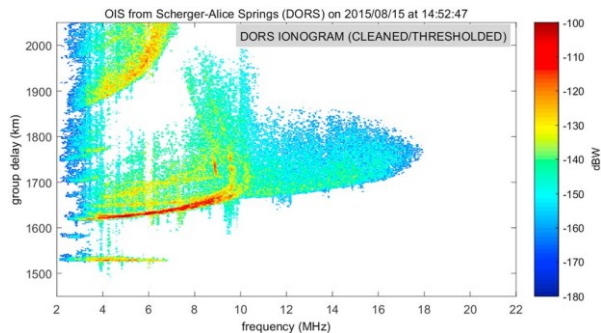
Figure 11. An enlarged DORS ionogram for a 1,495-km OIS path from Scherger to Alice Springs, showing severe off-angle propagation and post-midnight spread *F* conditions. These propagation features remain well within the DORS sensitivity limits, even after thresholding is applied.

scattering function range-Doppler map are distinctly negative, corresponding to the upward motion of the F_2 layer during the day-to-night transition.

6. Future Development

A future upgrade to DORS will see the DRx MkIID replaced with DST Group's NexGenR receiver. This represents the evolution of the direct-digital HF receiver design, with a 160-MHz A/D converter, optimized RF front-end, and a total of eight down-converter channels. Although not yet operational, recent experimental testing has demonstrated that the NexGenR will offer improved dynamic range over a wider part of the HF band (2–45 MHz). The DOTS DWG will be upgraded in tandem to run at the same 160-MHz clock speed.

Commercialization discussions for DORS and DOTS are also underway, to ensure the viability of maintenance and support into the future.

7. Conclusions

A high-fidelity, multichannel oblique incidence ionosonde has been developed by DST Group for routine observations of the ionosphere. The DORS receiving subsystem supports the simultaneous collection of multiple high-resolution ionograms, and offers excellent SNR even with a relatively fast sounder revisit rate. It has many potential applications to problems requiring accurate knowledge of the bottomside electron density profile, such as HF radar and communications, space weather monitoring, and low-frequency radio astronomy. As a sounder network component, it represents a reduced cost per spatial sampling point compared to traditional networks of vertical incidence sounders, and its low-power, direct-digital design offers an exceptional level of quality and robustness.

Key aspects of the DORS/DOTS ionosonde have been described, including its modular hardware design, flexible command and control, sophisticated onboard processing, data management, and auxiliary functions. Example ionograms and fitted traces speak to the overall quality of the data products. Interested readers are referred to companion papers by Heitmann and Gardiner-Garden (2019) and Turley et al. (2019) for more detailed descriptions of the feature extraction and fitting (DST-IIP) algorithm and RFI rejection algorithm, respectively.

References

Abramovich, Y. I., Frazer, G. J., & Johnson, B. A. (2013). Principles of mode-selective MIMO OTHR. *IEEE Transactions on Aerospace and Electronic Systems*, 49(3), 1839–1868. https://doi.org/10.1109/TAES.2013.6558024

Akchurin, A. D., Bochkarev, V. V., Ryabchenko, E. Y., & Sherstyukov, O. N. (2009). Improved precision of virtual height measurements with coherent radio pulse sounding based on the maximum likelihood method. *Advances in Space Research*, 43(11), 1595–1602. https://doi.org/10.1016/j.asr.2008.03.033

Arthur, P. C., Lissimore, M., Cannon, P. S., & Davies, N. C. (1997). Application of a high quality ionosonde to ionospheric research. Paper presented at the 7th International Conference on HF Radio Systems and Techniques, IEE Conference Publication No. 441, 135–139. https://doi.org/10.1049/cp:19970776

Barnes, R. I., Gardiner-Garden, R. S., & Harris, T. J. (2000). Real time ionospheric models for the Australian Defence Force. Paper presented at the Workshop on the Applications of Radio Science WARS-2000, 122–135. Retrieved from http://www.sws.bom.gov.au/IPSHosted/NCRS/wars/wars2000/

Bennett, J. A., Chen, J., & Dyson, P. L. (1991). Analytic ray tracing for the study of HF magneto-ionic radio propagation in the ionosphere. *Applied Computational Electromagnetics Society Journal*, 6(1), 192–210.

Bowman, G. G. (1990). A review of some recent work on mid-latitude spread-*F* occurrence as detected by ionosondes. *Journal of Geomagnetism and Geoelectricity*, 42(2), 109–138. https://doi.org/10.5636/jgg.42.109

Breit, G., & Tuve, M. A. (1925). A radio method of estimating the height of the conducting layer. *Nature*, 116(2914), 357. https://doi.org/10.1038/116357a0

Breit, G., & Tuve, M. A. (1926). A test of the existence of the conducting layer. *Physical Review*, 28(3), 554–575. https://doi.org/10.1103/PhysRev.28.554

Chen, J., Bennett, J. A., & Dyson, P. L. (1992). Synthesis of oblique ionograms from vertical ionograms using quasi-parabolic segment models of the ionosphere. *Journal of Atmospheric and Terrestrial Physics*, 54(3–4), 323–331. https://doi.org/10.1016/0021-9169(92)90011-9

Colegrove, S. B. (2000). Project Jindalee: From bare bones to operational OTHR. *Proceedings of the IEEE International Radar Conference 2000*, (pp. 825–830). https://doi.org/10.1109/RADAR.2000.851942

Dyson, P. L., & Bennett, J. A. (1988). A model of the vertical distribution of the electron concentration in the ionosphere and its application to oblique propagation studies. *Journal of Atmospheric and Terrestrial Physics*, 50(3), 251–262. https://doi.org/10.1016/0021-9169(88)90074-8

Fabrizio, G., & Heitmann, A. (2013). A multipath-driven approach to HF geolocation. *Signal Processing*, *93*(12), 3487–3503. https://doi.org/10.1016/j.sigpro.2013.01.026

Fabrizio, G. A. (2013). *High Frequency Over-The-Horizon Radar*. New York, NY: McGraw-Hill.

Fenwick, R. B., & Barry, G. H. (1966). Sweep-frequency oblique ionospheric sounding at medium frequencies. *IEEE Transactions on Broadcasting*, *BC-12*(1), 25–27. https://doi.org/10.1109/TBC.1966.265861

Frazer, G. J. (2007). Forward-based receiver augmentation for OTHR. *Proceedings of the IEEE International Radar Conference 2007* (pp. 373–378). https://doi.org/10.1109/RADAR.2007.374245

Gardiner-Garden, R., Ayliffe, J., Durbridge, L., Frazer, G., Heitmann, A., Praschifka, J., et al. (2008). A new high-fidelity oblique HF receiving system. Paper presented at the Ionospheric Effects Symposium (IES2008), Alexandria, VA.

Gardiner-Garden, R., Heitmann, A., & Frazer, G. (2011). Variability observed in a high-fidelity model of the ionosphere based on a dense network of vertical and oblique incidence sounders. Paper presented at the Ionospheric Effects Symposium (IES2011), Alexandria, VA.

Gong, W., Cui, X., & Pan, L. (2016). Design and application of the digital multifunctional ionosonde. *IET Radar, Sonar and Navigation*, *10*(7), 1303–1309. https://doi.org/10.1049/iet-rsn.2015.0561

Grubb, R. N., Livingston, R., & Bullett, T. W. (2008). A new general purpose high performance HF radar. Paper presented at the 29th URSI General Assembly, Chicago, IL.

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. https://doi.org/10.1109/PROC.1978.10837

Harris, T. J., Quinn, A. D., & Pederick, L. H. (2016). The DST group ionospheric sounder replacement for JORN. *Radio Science*, *51*, 563–572. https://doi.org/10.1002/2015RS005881

Heitmann, A. J., Cervera, M. A., Gardiner-Garden, R. S., Holdsworth, D. A., MacKinnon, A. D., Reid, I. M., & Ward, B. D. (2018). Observations and modeling of traveling ionospheric disturbance signatures from an Australian network of oblique angle-of-arrival sounders. *Radio Science*, *53*, 1089–1107. https://doi.org/10.1029/2018RS006613

Heitmann, A. J., & Gardiner-Garden, R. S. (2019). A robust feature extraction and parameterized fitting algorithm for bottom-side oblique and vertical incidence ionograms. *Radio Science*, *54*. https://doi.org/10.1029/2018RS006682

Hocke, K., & Schlegel, K. (1996). A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995. *Annales Geophysicae*, *14*(9), 917–940. https://doi.org/10.1007/s00585-996-0917-6

Hunsucker, R. D. (1982). Atmospheric gravity waves generated in the high-latitude ionosphere: A review. *Reviews of Geophysics and Space Physics*, *20*(2), 293–315. https://doi.org/10.1029/RG020i002p00293

Ivanov, V. A., Kurkin, V. I., Nosov, V. E., Uryadov, V. P., & Shumaev, V. V. (2003). Chirp ionosonde and its application in the ionospheric research. *Radiophysics and Quantum Electronics*, *46*(11), 821–851. https://doi.org/10.1023/B:RAQE.0000028576.51983.9c

Mendillo, M. (2006). Storms in the ionosphere: Patterns and processes for total electron content. *Reviews of Geophysics*, *44*, RG4001. https://doi.org/10.1029/2005RG000193

Reinisch, B. W., Galkin, I. A., Khmyrov, G. M., Kozlov, A. V., Bibl, K., Lisysyan, I. A., et al. (2009). New Digisonde for research and monitoring applications. *Radio Science*, *44*, RS0A24. https://doi.org/10.1029/2008RS004115

Rietveld, M. T., Wright, J. W., Zabotin, N., & Pitteway, M. L. V. (2008). The Tromsø dynasonde. *Polar Science*, *2*(1), 55–71. https://doi.org/10.1016/j.polar.2008.02.001

Turley, M. D. E., Heitmann, A. J., & Gardiner-Garden, R. S. (2019). Ionogram RFI rejection using an auto-regressive interpolation process. *Radio Science*, *54*. https://doi.org/10.1029/2018RS006683

Whitehead, J. D. (1989). Recent work on mid-latitude and equatorial sporadic-E. *Journal of Atmospheric and Terrestrial Physics*, *51*(5), 401–424. https://doi.org/10.1016/0021-9169(89)90122-0

Zuccheretti, E., Tutone, G., Sciacca, U., Bianchi, C., & Arokiasamy, B. J. (2003). The new AIS-INGV digital ionosonde. *Annals of Geophysics*, *46*(4), 647–659. https://doi.org/10.4401/ag-4377

F.3   Heitmann & Gardiner-Garden, *Radio Science*, 2019

The following is a reprint of Heitmann, A. J., & Gardiner-Garden, R. S. (2019), 'A robust feature extraction and parameterized fitting algorithm for bottom-side oblique and vertical incidence ionograms', *Radio Science*, 54(1), pp. 115–134.

This paper provides a detailed explanation of the ionogram feature extraction and parameterised profile fitting algorithm developed for DST Group's Digital Oblique Receiving System (DORS) OIS. The so-called DST-IIP algorithm was later included among the on-board processing for the ELOISE AoA ionosonde. The paper thus provides relevant background material to supplement the signal processing description in **Chapter 4** of this thesis.

The manuscript was written by myself, as the lead and corresponding author, and covers a program of work carried out by myself under the supervision of Dr Robert Gardiner-Garden. The paper was published during my PhD candidature.

Key Points:
- An algorithm for automatically scaling oblique and vertical incidence ionograms has been developed in conjunction with a new ionosonde
- The algorithm extracts key features describing the E, Es, and F2 layers, to which a parameterized profile of electron density is fitted
- Performance is shown to be robust under a variety of ionospheric conditions, and results compare well to other nearby sounders

A Robust Feature Extraction and Parameterized Fitting Algorithm for Bottom-Side Oblique and Vertical Incidence Ionograms

Andrew J. Heitmann[1,2] [iD] and Robert S. Gardiner-Garden[3] [iD]

[1]Defence Science and Technology (DST) Group, Edinburgh, South Australia, Australia, [2]School of Physical Sciences, University of Adelaide, Adelaide, South Australia, Australia, [3]Defence Science and Technology (DST) Group, Eveleigh, New South Wales, Australia

Abstract Automatic processing of ionograms to extract ionospheric propagation characteristics and electron density profile parameters is an important enabling step toward constructing real time data-assimilative models of the ionosphere. A robust and reliable feature extraction and fitting algorithm that supports both oblique and vertical incidence ionograms has recently been developed to enhance the regional ionospheric model for Australia's Jindalee Operational Radar Network. The so-called DST-IIP algorithm is an autoscaling technique based on a constrained quasi-parabolic segment profile, which is modified from its climatological estimate to fit a set of key peak, edge, and line features in the ionogram image, representing the E, sporadic-E (Es), and F2 layers. The F2 fitting, in particular, is performed in the ionogram domain, using analytic ray-tracing and homing to produce a synthetic trace that reproduces the image features. Following thorough testing on a large database of midlatitude Australian ionograms, the DST-IIP algorithm now runs routinely on-board DST Group's Digital Oblique Receiving System. This paper provides an overview of the technique, along with sample results and performance statistics.

Plain Language Summary An ionogram image captures the path losses and delays for radio propagation via the ionosphere, an outer charged layer of the Earth's atmosphere that refracts high-frequency radio waves. In order to relate ionograms from multiple ionospheric paths, the usual approach is to automatically convert these images to a parameterized form representing the midpoint electron density profiles. This paper describes such a technique, based on fitting the parameterized profile to a set of key peak, edge, and line features in the ionogram image. It is shown to be robust and reliable across a variety of path geometries and propagation conditions and now runs routinely on-board DST Group's own ionospheric sounder (or ionosonde), the Digital Oblique Receiving System.

1. Introduction

The ionosphere is a highly dynamic propagation medium, with multiple causes of spatial and temporal variability affecting the electron density distribution. For many high-frequency (HF) communication and radar applications, a real time data-driven approach to ionospheric modeling is essential for reliable propagation advice (including frequency selection) and, in the case of sky wave over-the-horizon radar, accurate coordinate registration of targets (i.e., the mapping from measured group delay and angle of arrival to the geographic location of the target). Such models require accurate and timely measurements of the bottomside ionospheric profile, at a large number of points in both space and time.

Networks of ground-based ionosondes are a well-established means of observing the bottomside ionosphere at low cost and with high sensitivity. The standard data product is the ionogram, representing the group delay of ionospheric propagation modes as a function of frequency (where group delay in this paper is plotted in km as a virtual range). The "echo" traces contained in the ionogram are directly related to the electron density profile (Davies, 1990). While ionogram scaling was in the past a largely manual task, advances in computing have given rise to a number of sophisticated automatic scaling algorithms in recent decades. Most run in near real time on the sounder system itself, applying a consistent set of conventions (Piggott & Rawer, 1978) to characterize the propagation characteristics and reduce the

ionogram image to a parameterized representation. These automatic techniques generally fall into one of two main categories:

1. Trace extraction or skeletonization, typically using a combination of morphological image operations, computer vision, and tracking algorithms (e.g., Chen et al., 2013; Galkin & Reinisch, 2008; Guiducci et al., 1983; Su et al., 2012; Zheng et al., 2013), followed by profile inversion (e.g., Huang & Reinisch, 1982; Phanivong et al., 1995; Reilly & Kolesar, 1989; Titheridge, 1988). A prominent example is the software package ARTIST (Galkin & Reinisch, 2008; Reinisch et al., 1988; Reinisch et al., 2005; Reinisch & Huang, 1983), developed for the Lowell Digisonde (Reinisch et al., 2009). Automatic profile inversion or true height analysis is commonly based on single or stepwise (overlapped) polynomial fitting techniques (Titheridge, 1988). Fuzzy segmentation has also been applied in place of skeletonization (Tsai & Berkey, 2000).
2. Template-matching techniques, which directly fit a synthesized trace or curve, drawn from a candidate set of parametric profiles, to extracted features or amplitude data in the ionogram image (e.g., Ding et al., 2007; Fox & Blundell, 1989; Jiang et al., 2014; Song et al., 2016; Wright et al., 1972). A prominent example is Autoscala (Pezzopane & Scotto, 2004, 2007, 2008; Scotto, 2009; Scotto & Pezzopane, 2007), developed for the AIS-INGV ionosonde (Zuccheretti et al., 2003). The use of empirical orthogonal functions has been shown to be efficient in this context (Ding et al., 2007), although with advances in computing, iterative ray tracing through realistic 3-D ionospheric representations is now also feasible (Zabotin et al., 2006).

Many such algorithms are restricted to vertical incidence (VI) geometry, although there are some examples that have specifically been developed or adapted for oblique incidence observations, without requiring an equivalent range transformation (e.g., Ippolito et al., 2015; Redding, 1996; Roughan, 1996; Song et al., 2016). In most cases, ordinary (*o*) and extraordinary (*x*) polarization modes remain unseparated in oblique incidence ionograms, which creates greater challenges for the scaling software and means magneto-ionic effects cannot be neglected. Furthermore, while it is possible to make use of directional data to fit a 3-D ionosphere (e.g., Zabotin et al., 2006), most algorithms for scaling oblique ionograms assume the ionosphere to be spherically symmetric, because directional information is not generally available.

This paper presents a robust, heuristic feature extraction and parametric curve fitting algorithm within the class of template-matching techniques described above, hereafter referred to as the DST-IIP (Ionogram Image Processing) algorithm. The general approach for extracting peak and edge features from ionograms using maximum or mean intensity arrays has similarities to the histogram-based methods of Igi (1992) and Lynn (2018), although these were designed to support solely VI ionograms, with or without amplitude information, so the maximum was not used. In the terminology of these previous methods, a mean intensity array can be interpreted as a type of weighted histogram. As well as parameterizing each ionogram image in terms of a three-layer electron density profile (with embedded sporadic-E reflectors), the DST-IIP algorithm also automatically reports levels of quality and uncertainty associated with the results. It has been developed and tested using a large database of midlatitude Australian ionograms, spanning 2002–2018 and containing tens of millions of observations from both oblique incidence sounders (OIS) and vertical incidence sounders (VIS).

For maximum compatibility with most modern sounder systems, the DST-IIP algorithm is generalized to accept either thresholded or nonthresholded ionograms, with or without polarization information (i.e., ordinary and extraordinary mode separation), and is able to cope with varying levels of input image quality. It uses solely the amplitude-scaled ionogram and does not require angle-of-arrival or Doppler measurements, which are not always available. Furthermore, as far as the algorithm is concerned, VIS ionograms are simply treated as a special case of an OIS ionogram, with transmitter and receiver colocated, so the approach is fundamentally the same for both sounder types.

The paper starts by describing the background of the DST-IIP algorithm, in the context of support for HF radar, and the ionospheric sounder instrumentation for which it was primarily developed (section 2). It then outlines the basic premise and assumptions of the algorithm (section 3) and proceeds to detail the processing steps for each layer (section 4), which are demonstrated by example. The performance and limitations over a broader set of ionospheric conditions are subsequently presented (section 5), followed by a brief summary and conclusions (section 6).
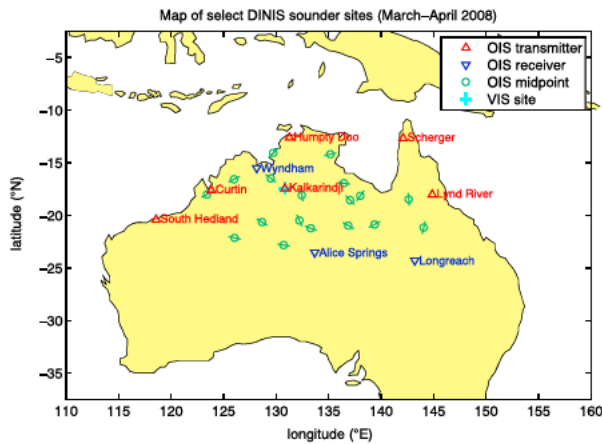
Figure 1. Map of select Australian sounder sites for the March–April 2008 period (corresponding to the DINIS trial conducted by DST Group). The 18 OIS paths, with one-hop midpoints indicated by the circular green markers, range in length from 300–2,600 km. A single VIS site at Kalkarindji (colocated with an OIS transmitter) is indicated by a cyan cross. OIS = oblique incidence sounders; VIS = vertical incidence sounder; DINIS = Dense Integrated Network of Ionospheric Sounders.

2. Background and Instrumentation

The Jindalee Operational Radar Network (JORN) is a network of three HF over-the-horizon (sky wave) radars operated by the Australian Defence Force, for monitoring aircraft and ships off the northern and western coasts of Australia. It is supported by a substantial number of remote OIS and VIS assets (currently forming 21 and 13 ionospheric sample points, respectively), which are used to accurately model radar propagation out to distances of 1,000–3,000 km. A subset of experimental JORN sounder paths from the Dense Integrated Network of Ionospheric Sounders trial of 2008 will be discussed in this paper; the sounder locations are mapped in Figure 1.

While the aforementioned database of midlatitude Australian ionograms includes data from established ionosonde systems, such as the Barry/TCI Chirpsounder and Lowell Digisonde, the main driver for the work presented in this paper has been to provide an on-board scaling capability for DST Group's Digital Oblique Receiving System (DORS; Ayliffe et al., 2018). This is a state-of-the-art direct-digital OIS receiving system, offering simultaneous collection of multiple oblique paths on a single omni-directional antenna. Along with a low-power (frequency-modulated continuous-wave) HF transmitting system, also developed by DST Group, this offers a complete point-to-point sounding capability. A customized configuration of the DST Group system for VI measurements with o/x separation has previously been reported (Harris et al., 2016). DST Group has also developed a solution for the reliable automatic separation of oblique o/x returns using twin-polarized antennas (Harris et al., 2017), although it is not part of the standard DORS configuration.

In the context of real time support for JORN, the focus of the DST-IIP algorithm is on the characterization of E and sporadic-E (Es) layers, being important propagation modes for detecting slow-moving surface targets in clutter, and the F2 layer, being important for detecting fast-moving air targets in noise. An overview of the principles of HF radar operation can be found in Headrick and Anderson (2008) and Fabrizio (2013). While climatology predicts many of the regular patterns of variation in the ionosphere, there exists a significant amount of residual day-to-day and hour-to-hour variability, particularly in the Es and F2 layers, that the algorithm seeks to capture in its parameterization of each ionogram.

Unlike the HF communication problem, HF radar demands not only accurate estimates of critical frequencies but also layer heights, in order to provide effective coordinate registration of targets. Whereas historically the quality of results from OIS observations may not have been at the same standard as VIS observations for this purpose, it is believed that the DST-IIP algorithm moves a substantial way toward bridging the gap. Indeed, it has been a key enabler for DST Group's recent integration of both VIS and OIS into a single regional ionospheric model for Australian HF radar operations.

3. Overview of the DST-IIP Algorithm

The purpose of the DST-IIP algorithm is to construct a three-layer electron density model (with embedded Es reflectors) that reproduces the key features of the ionogram image via ray tracing from transmitter to receiver site. In effect, a synthetic ionogram trace can be generated for each parameterized profile, and its ability to match the image is a measure of the algorithm's success. Essential to this technique is the use of intensity arrays of maximum or mean pixel amplitudes to draw out features in the image, and the use of equivalent range transformations to help account for the oblique geometry (although it is noted that the F2 parameters are ultimately fitted using exact analytic ray tracing (ART), not equivalent geometry approximations). This generalized approach means that the same processing can be applied to both OIS and VIS ionograms.

The algorithm is notionally split into a number of feature extraction steps, which identify key peak, edge, and line features in the original or transformed images, and fitting steps, which relate these observed features to a realistic parameterized profile of the midpoint electron density. Whereas a feature is tied directly to the image and assumes no explicit profile constraints, a fitted parameter is an element of the (no-field)
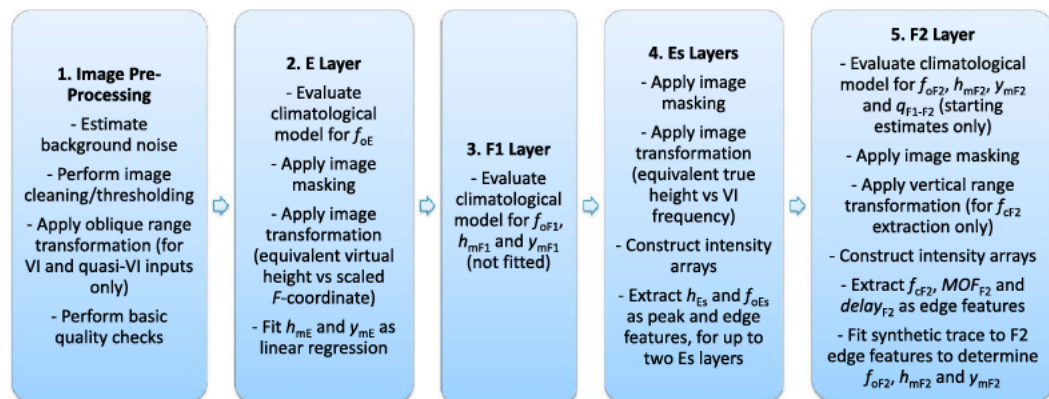
Figure 2. A top-level diagram of the DST-IIP algorithm, in terms of the processing steps taken to parameterize each of the E, F1, Es, and F2 ionospheric layers.

profile construction and must obey certain rules, as outlined in Gardiner-Garden et al. (2018). In some cases, the feature is adopted more or less directly as the profile parameter, with only minor profile constraints applied; this applies to the E and Es fitting. In other cases, there are additional corrections, to compensate for geomagnetic o/x splitting, and a full iterative fitting process; this applies to the F2 fitting. A top-level (layer-by-layer) breakdown of the individual processing steps is provided in Figure 2; a more detailed description of these steps will follow in the next section.

The processing starts with the lower ionospheric layers and progressively moves up through the ionosphere, using knowledge of the lower layers to improve the estimation of the higher-layer parameters. Established climatological (empirical) models for the E, F1, and F2 layers are used to define the initial values for the profile fitting. As most of the day-to-day and hour-to-hour variability occurs in the Es and F2 layers, the focus is on robust and reliable parameterization of these components, while the E and F1 layers remain largely based on climatological models (with the exception of E layer heights). Only one-hop returns are considered in the current implementation, which limits results to oblique paths of less than about 2,000 km for E and Es layer parameters and 3,000 km for F2 layer parameters.

The JORN sky wave radars are supported by a regional data-assimilative model of the ionospheric electron density known as the Real-Time Ionospheric Model (RTIM; Barnes et al., 2000). As the motivation for this work was primarily to provide compatible inputs for the RTIM, the 10-parameter JORN ionospheric profile specification was adopted as the standard output. This is a simple quasi-parabolic (QP) segment formulation (Croft & Hoogasian, 1968; Dyson & Bennett, 1988; Hill, 1979), consisting of three physical layers (E, F1, and F2) and their join segments (Gardiner-Garden et al., 2018). The 10 profile parameters comprise the critical frequencies, peak heights, and semithicknesses for each of the three layers, along with a single index defining the start point of the F1-F2 join. Sporadic-E layers may be inserted into this profile as separate mirror reflectors. The climatological E, F1, and F2 models, used whenever fitting is not attempted or possible, are also the same as those used by JORN in its RTIM; they are similar but not identical to those included in the International Reference Ionosphere (IRI) (Bilitza, 1990, 2001; Bilitza et al., 2014; Bilitza & Reinisch, 2008).

An ART engine, based on the integral equations in Dyson and Bennett (1988) and Bennett et al. (1991), is used throughout the algorithm to synthesize oblique incidence group delay traces from the electron density (or plasma frequency) profile that is being fitted. In this way, the solution to the profile fitting problem is obtained in the ionogram (group delay) domain, through an iterative (nonlinear) shooting method. Coupled to the trace synthesis is a first-order model for geomagnetic field effects (i.e., o and x mode splitting; Bennett et al., 1991; Chen et al., 1990, 1992). This is applied as a frequency and delay perturbation to the no-field solution from the ART engine and serves to capture the bulk of the dependence on path location and geometry in a fast, analytical approach. As such, it is well suited to the iterative fitting technique described here.

A key assumption made by the algorithm is that the observed propagation can be derived from a spherically symmetric and Earth-concentric ionosphere. This means that each one-hop oblique trace is associated with a
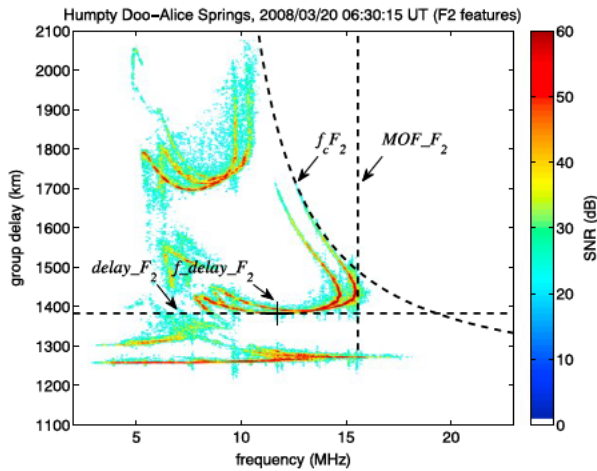
Figure 3. A daytime DORS ionogram from Humpty Doo to Alice Springs (1,234 km) in northern Australia. The DORS sounder uses an overlapped 200-kHz coherent processing window to generate high-resolution ionograms sampled at 40 kHz and 1.5 km in frequency and group delay, respectively. The key one-hop F2 features, extracted automatically, are overplotted. DORS = Digital Oblique Receiving System; SNR = signal-to-noise ratio.

single electron density profile, nominally at the path midpoint. In fact, this assumption is not quite as severe as it first appears, since horizontal structures that are antisymmetric about the midpoint (including many large-scale gradients) do not affect the measured group delay to first order (Bennett & Dyson, 2002).

As the software is designed to run on-board the DORS receiver, each ionogram is analyzed independently and has no knowledge of gradients recorded across the network of other sounder paths. Equivalent propagation geometry, based on the theory of virtual ionospheric reflection (viz., the modified secant law, Breit and Tuve's theorem, and Martyn's theorem, described in Davies, 1990, and McNamara, 1991), is used in a limited capacity to relate frequencies and group delays at different ground ranges; in most circumstances for the E and Es layer fitting only, and as a fall-back for the F2 layer fitting when convergence cannot be achieved with the QP ART engine. Otherwise, the algorithm remains in the original OIS domain as much as possible and avoids approximate image transforms that could limit its accuracy.

Note that by focusing on several key features, the limited degrees of freedom of the parameterized profile are used to capture these properties, but not necessarily other points on the one-hop traces. That said, the features are by definition among the most operationally relevant characteristics, including the critical frequency, maximum observed frequency, and minimum delay of F2 propagation, so for many applications this is of limited concern. The inherent spreading or width in the image features (whether natural or an artifact of the ionogram measurement) is used as an estimate of uncertainty in the feature extraction process. Naturally, edge features do not detect the center of the trace in the image, so this can also contribute to a slight mismatch.

In the detailed description that follows, the treatment of each layer is discussed in turn, following the process order in Figure 2 to construct a full set of electron density profile parameters from the bottom up. A single representative daytime ionogram is used to illustrate the results. The Humpty Doo to Alice Springs DORS sounding, shown with F2 features overplotted in Figure 3 and fitted (synthetic) traces overplotted in Figure 4, comes from a 1,234-km path in northern Australia at 0630 UT (1600 LT). The fitted electron density profile from which the traces were synthesized is shown in Figure 5.

All aspects of the algorithm, including a stand-alone GUI test bed for development and diagnostics, have been implemented in MATLAB®. On a modest server or laptop, the typical DST-IIP processing time is 10–20 s for a single high-resolution DORS ionogram, which is more than fast enough to keep up with real time observations, particularly if using a separate thread per sounder path. Thresholds for peak and edge detection are treated as tunable algorithm parameters, trading off the probability of detection against the probability of false alarm. As much as possible, they are adjusted automatically based on the range of signal-to-noise ratio (SNR) values in the intensity arrays of the input ionogram and, as such, are not stated explicitly in the description below.

Figure 4. The same ionogram as in Figure 3, with synthetic traces overplotted. The one-hop trace in dark blue (o-mode) and black (x-mode) has been fitted to the extracted E, Es, and F2 features, while the two-hop traces (including Es/F combinations) in light blue, light gray, red, and magenta are predicted only from the one-hop fit and may or may not match the image depending on horizontal gradients in the ionosphere.

4. Detailed Description of the DST-IIP Algorithm

4.1. Image Preprocessing

As the characteristics of the input ionogram may vary considerably depending on the source system, a number of preprocessing steps are
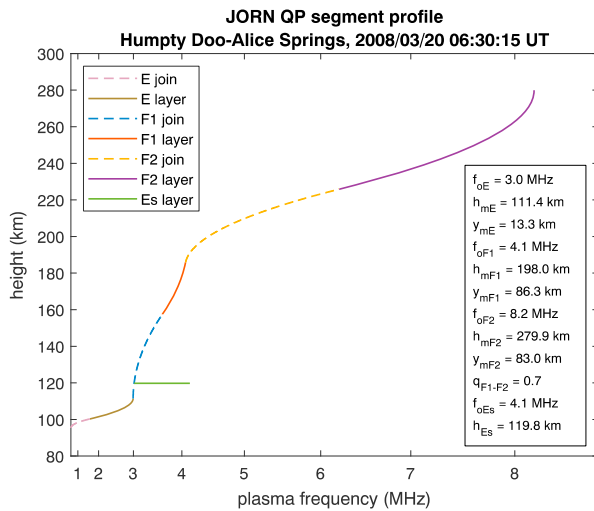
Figure 5. The electron density profile corresponding to the fitted (synthetic) traces in Figure 4. The inset text box (lower right) contains the full set of 10 QP parameters and two mirror Es parameters that specify this profile shape; refer to Gardiner-Garden et al. (2018) for a detailed description of the JORN profile specification. JORN = Jindalee Operational Radar Network; QP = quasi-parabolic.

applied to the image to condition it for the feature extraction and fitting stages that follow. These include the following:

1. Background noise estimation and rescaling of the ionogram in terms of the SNR. Wherever possible, the noise floor is estimated as a function of frequency using median values in the guard band only, that is, at group delays below the great circle ground range of the path. For typical DORS ionograms, the guard band is ~0.4 ms or ~80 delay cells wide.

2. Cleaning/thresholding to remove unwanted artifacts due to radio frequency interference (RFI) and speckle noise. Typically, a minimum SNR of 20 dB is imposed across the whole ionogram; all pixels below this level are assigned a special "no data" value and do not contribute to the subsequent feature extraction. For frequencies where the median amplitude in the guard band exceeds this minimum SNR, indicative of residual RFI, the threshold is raised accordingly, although this is rare for DORS data that have already undergone RFI rejection on the raw time samples (Turley et al., 2018). Isolated artifacts due to speckle noise are addressed by removing (i.e. setting to the "no data" value) any individual pixel for which less than 5% of the total pixels in the surrounding 1-MHz by 30-km neighborhood remain after thresholding.

3. For VI and quasi-VI systems, applying an equivalent range transformation such that the minimum effective ground range for oblique F2 feature extraction is 500 km. Although the exact value is somewhat arbitrary, a 500-km range roughly represents the transition point from a quasi-vertical trace shape to an oblique trace shape, with a maximum observed frequency that is sufficiently different from the critical frequency. Note that this does not degrade the accuracy of the fitted F2 layer parameters, as the trace synthesis at each iteration of the fit applies the same mathematical transform; it merely adjusts the shapes of the feature curves, and thus their tangent points, to provide a better characterization of the observed trace.

In addition, several checks exist to detect bad or incompatible ionograms, on the basis of insufficient trace power, abnormally strong RFI, excessive ground range, or poor ionogram dimensions. For some of these rare cases, ionogram processing is unable to proceed.

4.2. E Layer

The E layer critical frequency (f_{oE}) has a well-understood climatological dependence on latitude/longitude, hour of day, day of year, and solar activity. On oblique paths with unseparated o/x returns and marginal SNR along the E high-ray trace, it is not expected that ionogram analysis can substantially improve on the climatological f_{oE} estimates derived from monthly-median observations. The standard IRI empirical model for daytime f_{oE} (Kouris & Muggleton, 1973; Muggleton, 1975) is therefore adopted within the DST-IIP algorithm, driven by the smoothed sunspot number R_{12}. Note that to support legacy models such as this, based on the original (v1) international sunspot number, the latest R_{12} (v2) values published online are scaled down by a multiplicative factor of 0.7 (Clette et al., 2014; Clette & Lefevre, 2016). As the nighttime behavior of f_{oE} is of little consequence at the HF frequencies observed, a simple fixed residual value suffices between dusk and dawn.

While empirical models have also been developed for the E layer peak height (h_{mE}; e.g., Ivanov-Kholodny et al., 1998), these have been subject to more limited validation on a global scale. For high-resolution ionograms, such as those from DORS, the group delay resolution of ~1.5 km is more than enough to observe a large proportion of the ~10-km typical daily variation in h_{mE}, along with a corresponding scale of variation in the layer semithickness (y_{mE}). Therefore, the approach taken is to transform the ionogram image with a parabolic layer approximation and equivalent propagation geometry and fit the E heights directly as a line feature in the transformed image, as outlined below.

Assuming VI propagation through a single parabolic layer, in the absence of collisions and geomagnetic effects, it can be shown that the E layer virtual height of reflection obeys the following relationship (Davies, 1990, equation 6.11):

$$h_{\mathrm{E}}'(x_{\mathrm{vE}}) = h_{0\mathrm{E}} + y_{\mathrm{mE}}F(x_{\mathrm{vE}}), \qquad (1)$$

where

$$F(x_{\mathrm{vE}}) \equiv \frac{x_{\mathrm{vE}}}{2}\ln\left(\frac{1+x_{\mathrm{vE}}}{1-x_{\mathrm{vE}}}\right) \qquad (2)$$

and

$$0 \leq x_{\mathrm{vE}} \equiv \frac{f_{\mathrm{v}}}{f_{\mathrm{oE}}} < 1. \qquad (3)$$

Equation (1) implies that the virtual height h_{E}' is a linear function of the scaled frequency coordinate F (equation (2)), expressed in terms of the normalized wave frequency x_{vE} (equation (3)), that is, the ratio of the VI propagation frequency f_{v} to the (no-field) critical frequency f_{oE}. As such, a linear regression technique can be used to derive estimates for the base height $h_{0\mathrm{E}} = h_{\mathrm{mE}} - y_{\mathrm{mE}}$ and semithickness y_{mE}, provided f_{oE} is known; in this case, from the climatological model. For OIS observations, the ionogram pixels are range transformed to their equivalent VI frequencies and virtual heights prior to the linear regression, using the modified secant law and mirror geometry (Davies, 1990; McNamara, 1991).

To isolate the one-hop E trace from other unwanted modes, the equivalent VI ionogram must also be masked. The unmasked region only includes frequencies below 97% of the model f_{oE} estimate ($F < 2$), and virtual heights greater than the one-hop lower limit (for the equivalent raypath at zero elevation) and less than 140 km. All other pixels are set to the special "no data" value. The frequency constraint has been carefully selected to reduce sensitivity to unseparated o/x splitting, as well as errors in the model f_{oE}. Furthermore, inside the unmasked region, only pixels having SNRs within 30 dB of the peak SNR are preserved.

Although what remains is primarily one-hop E, there is the potential for low Es and other unwanted modes to remain, so to better cope with such outliers, a least trimmed squares fitting technique is used for the regression. In this case, the objective function to be minimized is the sum of squared residuals, trimmed at the ninetieth percentile and weighted by the SNR in decibels. As no closed form solution exists, the solution for $h_{0\mathrm{E}}$ and y_{mE} is found using an unconstrained nonlinear minimization function (Nelder-Mead simplex method; Lagarias et al., 1998), starting from the climatological model predictions. Nevertheless, when low Es is present, there is still an increased likelihood that the peak height and semithickness of the E layer will be underestimated.

While the line feature defining $h_{0\mathrm{E}}$ and y_{mE} can always be extracted in theory, provided two or more transformed pixels exist, in practice some of these estimates will not be of sufficient quality to be retained as a valid fit. Outliers are therefore removed by a set of filtering criteria, based on several factors including the tightness of the least trimmed squares fit (sum of squared residuals), the total number of valid pixels, the range of valid F-coordinate values, and the climatological f_{oE} estimate. This typically rejects up to 10–20% of the candidate E height fits from DORS OIS data, mostly concentrated around the dawn/dusk terminator periods ($f_{\mathrm{oE}} < 2$ MHz) and during instances of blanketing Es layers at low heights, when the E trace is naturally less distinct. The remaining fits have their quality characterized by their weighted sum of squared residuals, and estimated standard errors in $h_{0\mathrm{E}}$ and y_{mE} from the linear regression.

An illustration of the E height fitting stage is shown in Figure 6. The top panel is an enlarged region of the preprocessed DORS ionogram, while the bottom panel is with the masking and F-coordinate transformation applied (i.e., using equations (2) and (3)). After fitting (solid lines), it is seen that the synthetic traces overlaid on the top panel, and linear regression on the bottom panel, provide a good match to the image power. However, in this case, it is only a marginal improvement on the climatological model (dashed lines), mainly achieved through raising the base of the layer ($h_{0\mathrm{E}}$). It is worth noting that o and x modes are indistinguishable along the lower portion of the E trace after masking, even on this north-south oriented (quasi-longitudinal) path, which supports the underlying no-field assumption.
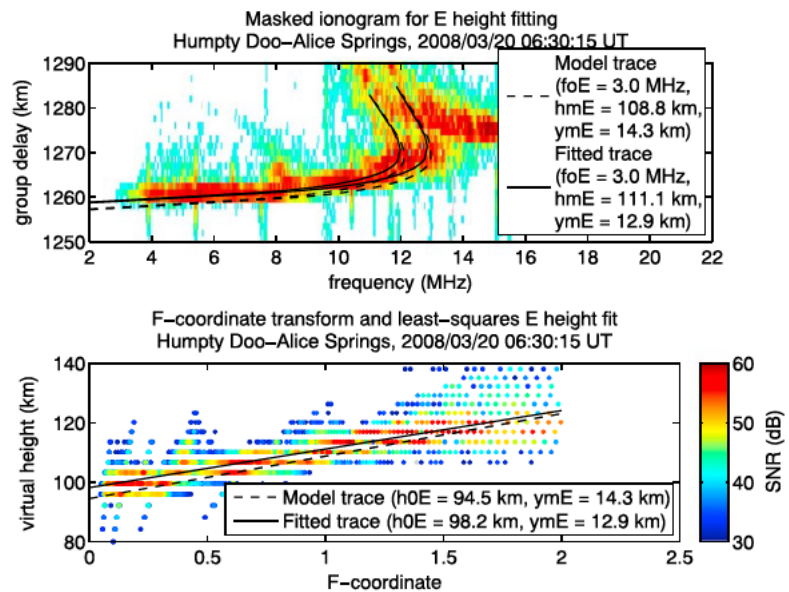
Figure 6. An illustration of the E height fitting stage for the sample ionogram in Figure 3. The unmasked pixels are transformed from the original ionogram domain (top panel) to the F-coordinate domain (bottom panel), enabling E height parameters h_{0E} and y_{mE} to be fitted directly as the intercept and slope of a least squares regression line. The fitted result (solid line) is compared directly with the JORN climatological model (dashed line) in both domains, albeit with o/x splitting neglected in the F-coordinate domain. JORN = Jindalee Operational Radar Network; SNR = signal-to-noise ratio.

4.3. F1 Layer

Like the E layer critical frequency, the F1 layer also tends to be well modeled by its climatology. The standard IRI empirical model for f_{oF1} (DuCharme et al., 1971, 1973) is adopted, although the JORN profile building rules demand a small (nonzero) nighttime value; nevertheless, this residual "pseudo" layer has only a negligible impact at HF frequencies. Using historical VIS data, a revised empirical model for F1 heights has recently been developed by DST Group to provide a reliable and accurate characterization in the Australian region; this work is yet to be published. With the F1 trace often poorly defined in OIS ionograms, it was felt that fits to individual images could not compete with the performance of the empirical model, so no such fitting is currently attempted in the DST-IIP algorithm.

4.4. Sporadic-E Layers

Sporadic-E consists of one or more additional thin layers at or around E layer heights that exhibit a high degree of variability in their occurrence and spatial structure (Mathews, 1998; Whitehead, 1989). However, when present, Es can provide excellent propagation characteristics, resembling specular reflection at times, so is a very valuable layer to detect and scale. The DST-IIP algorithm groups together all strongly ionized layers found at E region heights into the one definition of sporadic-E, without regard to their underlying physical mechanisms; this includes both descending tidal ion layers and truly sporadic layers.

Due to the layer thinness (less than a few km), each instance of Es is parameterized as a mirror reflector only, in terms of a critical frequency f_{oEs} and height h_{Es}. As described below, these profile parameters can be extracted directly as image features, so there is no separate Es fitting stage. Up to two layers are extracted by default, which supports the observed tidal pattern of Es descent and decay. Since the o and x modes are often unresolvable for Es, there is only partial compensation for geomagnetic splitting, and f_{oEs} in particular remains o/x ambiguous in the current version of the algorithm. It is important to note that the height parameter h_{Es} is the true height of the layer, not the frequency-dependent virtual height of reflection.

As with the E layer, the Es layer processing begins with the masking of unwanted modes. What remains after masking are frequencies above the synthetic E trace, and virtual heights greater than the one-hop horizon and less than 180 km. This is shown in Figure 7, with masked pixels in gray. In this case, a single Es layer
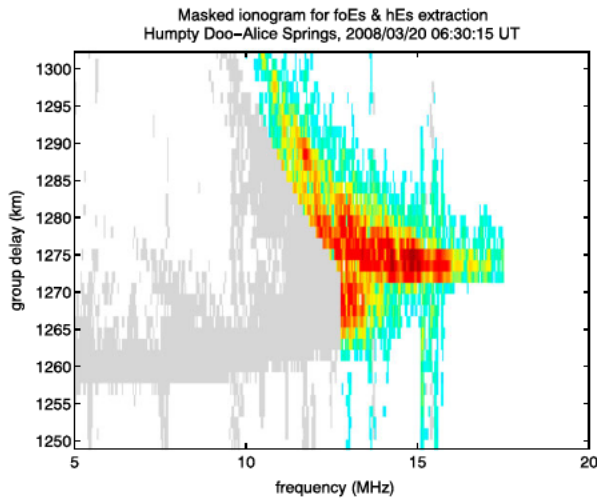
Figure 7. Masked image used for Es processing, for the sample ionogram in Figure 3. Pixels removed by the masking are shaded in gray.

is clearly apparent, along with a small amount of residual power from around the nose of the E trace; the slight mismatch of the fitted E trace in Figure 6 reveals the source of these unmasked pixels.

While there is no distinction between the true and virtual height of reflection for a thin layer in isolation, this is not true of Es embedded within a daytime E/F1 profile. To compensate for the additional refraction through the underlying E region, the masked ionogram is also transformed, using a combination of equivalent VI geometry and a first-order representation of the excess delay through the sub-Es electron density profile, in other words, accounting for the difference between the true height and virtual height. This is carried out separately for both *o* and *x* modes, and the two images are incoherently averaged to accentuate the no-field height. The resultant image of height versus equivalent VI frequency is shown in the main panel of Figure 8.

Subject to the no-field and equivalent geometry assumptions, as well as the accuracy of the sub-Es profile, each Es layer is expected to manifest as a near-horizontal line in the transformed image domain. As such, the detection problem is reduced to finding peaks in the horizontally integrated power. In general, the mean intensity array of pixel amplitudes better emphasizes line-like structure in an image (e.g., Es heights), while the maximum intensity array is better suited to edge-like features (e.g., Es critical frequencies). The left panel of Figure 8 shows this so-called height intensity array, calculated as the mean SNR across all frequencies. In this case, only one peak (Es layer) was detected, at a height of 120 km, after some mild smoothing to eliminate false alarms. A least squares parabolic fit (or Gaussian in linear units) is applied to each peak to precisely extract the layer height h_{Es} and its associated uncertainty (the 6-dB half width).

Along with the h_{Es} feature, the critical frequency f_{oEs} is also extracted from the same transformed image for each detected layer. This is based on locating the upper edge in the frequency intensity array, derived from the maximum SNR at heights within the h_{Es} uncertainty bounds. The array initially has generous smoothing applied to suppress any residual noise artifacts, and the edge, defined to be the point at which the array exceeds a specified threshold level, is gradually converged upon as the smoothing is successively reduced. The same edge detection is also carried out for higher and lower thresholds (offset by ±6 dB from the center threshold) to estimate the edge uncertainty. The bottom panel of Figure 8 shows the (unsmoothed) array and f_{oEs} edge feature at 4.1 MHz.

It is worth noting that Es layers commonly have a patchy horizontal structure (Whitehead, 1989) whose reflections cannot be fully characterized by one key frequency alone. Although f_{oEs} is arguably the simplest parameter to extract as a one-hop Es trace feature, it is not necessarily the most useful for spatial mapping purposes in ionospheric models. It represents the upper limit of partial reflection, which can be very dependent on the angle of incidence and sensitive to the background noise floor (Barnes, 1995). The blanketing frequency f_{bEs}, being the upper limit of full reflection, would contribute valuable additional information and is hoped to also be extracted in future.

4.5. F2 Layer

In the JORN ionospheric profile specification (Gardiner-Garden et al., 2018), the F2 layer is parameterized by the three QP layer parameters (f_{oF2}, h_{mF2}, and y_{mF2}), along with the F1-F2 join index (q_{F1-F2}). Under most circumstances, a unique solution for the QP layer can be fitted to three independent edge features along the trace, while the join index
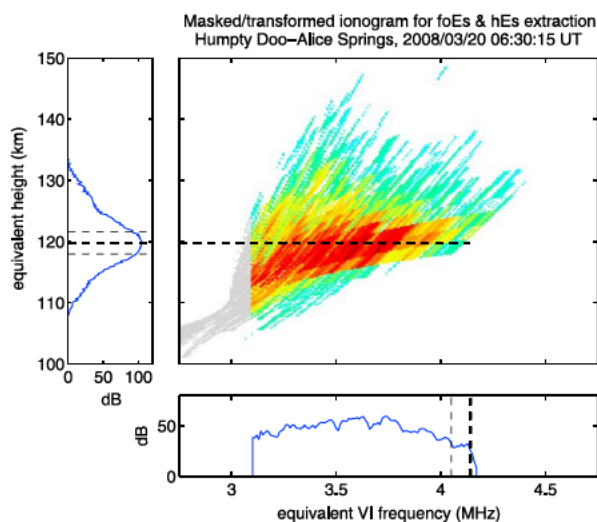
Figure 8. Masked and transformed image for Es processing (main panel), along with measures of mean and maximum signal-to-noise ratio along the height and frequency dimensions (left and bottom panels) used for h_{Es} (peak) and f_{oEs} (edge) extraction, respectively. The thick dashed lines represent the center peak and edge estimates, while the thin dashed lines capture the lower and upper uncertainty bounds.
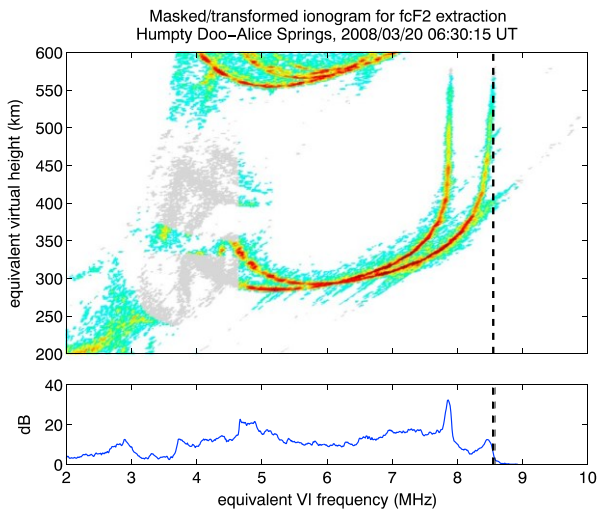
Figure 9. Masked and transformed image for F2 processing (main panel), along with a measure of mean signal-to-noise ratio along the frequency dimension (bottom panel) used for f_{cF2} (edge) extraction. The thick dashed line represents the center edge estimate, while the thin dashed lines (barely perceptible in this case) capture the lower and upper uncertainty bounds.

remains modeled by a climatological estimate. The three F2 features, each annotated on the sample ionogram in Figure 3, are as follows:

1. The F2 cusp asymptote, f_{cF2}. This resembles the critical frequency of the layer, prior to compensating for geomagnetic splitting. For iono-grams without o/x separation, it is most closely related to the equivalent f_{xF2}, although as the F2 fitting occurs in the original OIS domain, it is not interpreted directly as such in the algorithm.
2. The maximum observed F2 frequency, MOF_{F2}. Without o/x-separated ionograms, it is treated as an x-mode feature.
3. The minimum unambiguous F2 delay, $delay_{F2}$, and its associated frequency, $f(delay_{F2})$. In the absence of Es, it will correspond to reflection from the minimum virtual height h'_{F2}, but given multihop or blanketing Es, the feature may move further up the trace. By definition, $delay_{F2}$ may not be smoothly varying in time like f_{cF2} and MOF_{F2}, but by being coupled with $f(delay_{F2})$, the fitted F2 parameters (and their synthetic features) will generally still be continuous. Without o/x-separated ionograms, it is treated as an o-mode feature.

These three key features on the one-hop F2 trace are generally well separated within the image and, therefore, provide a good characterization of the trace as a whole, including the low ray, nose of penetration, and high ray.

Once again, a mask is initially applied to remove unwanted modes ahead of the F2 feature extraction, in this case targeting the removal of E, Es, and F1 traces, including multiple hops. All frequencies below the modeled f_{oE} and f_{oF1} are masked, as are virtual heights below 160 km. A more intricate multihop Es mask is also constructed based on traces synthesized from the one-hop f_{oEs} and h_{Es} parameters; this forms a band of rejected pixels about the synthetic two- and three-hop Es traces, with a width derived from the one-hop parameter uncertainties. The resultant ionogram, in the original OIS domain, is pictured in the main panel of Figure 10; this is the basis for MOF_{F2} and $delay_{F2}$ extraction. After transforming the oblique ionogram to its equivalent vertical representation (Davies, 1990; McNamara, 1991), though without the E and F1 masking, the result is the image in Figure 9; this is the basis for f_{cF2} extraction.

The F2 feature extraction proceeds in much the same way as for the previously described edge feature f_{oEs}. That is, for each of f_{cF2}, MOF_{F2}, and $delay_{F2}$, a frequency or delay intensity array is constructed from the maximum or mean SNR in the appropriate image, and the edge is found at three different thresholds (i.e., high, center, and low, separated by 3–6 dB depending on the feature). The bottom panel of Figure 9 illustrates this for f_{cF2}; the vertical dashed line through the corresponding image transforms back to the curved high-ray asymptote in the original (oblique) ionogram domain of Figure 3. Note that while the E/F1 mask is still used to constrain the edge detection, it has not been applied to the image in this case, as this can reduce f_{cF2} detectability when the F2 critical frequency is depressed, due to smoothing of the frequency intensity array. The bottom and left panels of Figure 10 show the same thing for MOF_{F2} and $delay_{F2}$, respectively. The $f(delay_{F2})$ value is simply the frequency with maximum SNR at the group delay slice equal to $delay_{F2}$ (with a small degree of smoothing applied).

Figure 10. Masked image for F2 processing (main panel), along with measures of maximum and mean signal-to-noise ratio along the frequency and delay dimensions (bottom and left panels) used for MOF_{F2} (edge) and $delay_{F2}$ (edge) extraction, respectively. The thick dashed lines represent the center edge estimates, while the thin dashed lines capture the lower and upper uncertainty bounds.

Given the three F2 features, a QP layer is fitted using a shooting technique and numerical solve routine (Broyden's method; Press et al., 1992), that is, by synthesizing traces through a spherically symmetric ionosphere until

the F2 electron density parameters (f_{oF2}, h_{mF2}, and y_{mF2}) match the observed F2 features. With three (known) features and three (unknown) parameters in this multidimensional nonlinear problem space, a unique and exact fit (to within numerical tolerance) is possible under most circumstances. In the small number of cases when one or more of the extracted features are wrong, the F2 layer fit will either be poor (for modest feature errors), or no solution will exist (for gross feature errors). Upper limits on the feature uncertainties and outlier filtering across adjacent times may be used to identify and remove these bad fits prior to subsequent ingestion into data-assimilative models such as the JORN RTIM.

During the fitting process, the F2 parameters remain constrained by the JORN profile building rules to avoid nonphysical solutions. The convergence criteria take into account the coarseness of trace sampling, as well as numerical noise in the QP ART integrals and ray interpolation, with default values being 20 kHz for f_{cF2} and MOF_{F2}, and 200 m for $delay_{F2}$. The starting estimates for the F2 parameters are derived from the JORN monthly-median maps, based on historical sounder data from the Australian Ionospheric Prediction Service, which are interpolated in latitude, longitude, and R_{12} index. If a complete fit is not possible, a subset of the F2 features may still yield a partial fit, by falling back to climatology to populate one or both of the peak height (h_{mF2}) and semithickness (y_{mF2}) parameters.

The F1-F2 join index is a dimensionless parameter controlling the shape of the F1-F2 cusp (Gardiner-Garden et al., 2018). By default, a climatological model for this parameter forces a stronger cusp (i.e., sharper join nearer the peak of the F1 layer) during the day, and a smoother transition from the (pseudo) F1 layer at night. However, under certain conditions, this may prohibit a complete fit; an example is depressed f_{oF2} conditions, during which the F2 trace is more sensitive to underlying ionization and the join may need to be reduced. As such, the algorithm is permitted to make small adjustments to the F1-F2 join index if necessary; that is, the $q_{F1\text{-}F2}$ parameter is sometimes indirectly fitted and may deviate from its model value. Note that in order to produce a valid set of QP segments, the JORN profile building rules may separately demand modifications to the underlying climatological F1 model (including the peak height), which will also affect the shape of the F1-F2 join.

The synthetic trace for a successful (complete) F2 fit is shown in Figure 4. The profile also includes the E height and Es layer fits from Figures 6 and 8, respectively. For a typical daytime ionogram such as this, the DST-IIP algorithm has no problem capturing the key one-hop modes with excellent fidelity, and indeed many of the multimode components too (suggesting no strong horizontal gradients).

More generally, one of the factors that can undermine confidence in the F2 fit is the influence of multihop Es. To address this, the $delay_{F2}$ extraction is repeated without multihop Es masking, and the same iterative F2 fitting algorithm is applied using this alternative feature estimate. If the $delay_{F2}$ features and/or minimum delays of the fitted (synthetic) F2 traces are comparable, with and without the extra masking, the multihop Es influence is declared to be negligible or, at least, unambiguous. If, on the other hand, a substantially different F2 fit is found, these results are flagged as ambiguous. Either way, the original F2 fit, obtained with multihop Es masking, is trusted to be the more correct solution. Note that for OIS path lengths greater than 2,000 km, without fitted Es parameters, this multihop Es influence cannot be assessed without the use of an external Es model.

5. Results and Performance

The DST-IIP algorithm has been in use experimentally since 2006, operationally as part of the JORN RTIM since 2011, and on-board DORS receivers since 2012. Over this time, the algorithm has been exposed to tens of millions of ionograms, from dozens of midlatitude sounder paths, some spanning more than a solar cycle. This provides considerable confidence in the quality and robustness of the results, and the internally reported measures of uncertainty. As deficiencies have been revealed, revisions have been made to the algorithm, so in its current form, it has evolved to better deal with the wide variety of ionospheric propagation conditions commonly observed.

In this section, results are presented for a set of 18 DORS OIS paths and one Lowell Digisonde VIS from April 2008; the path geometries are shown in Figure 1. This month is representative of equinox near sunspot minimum and was collected during the Dense Integrated Network of Ionospheric Sounders trial
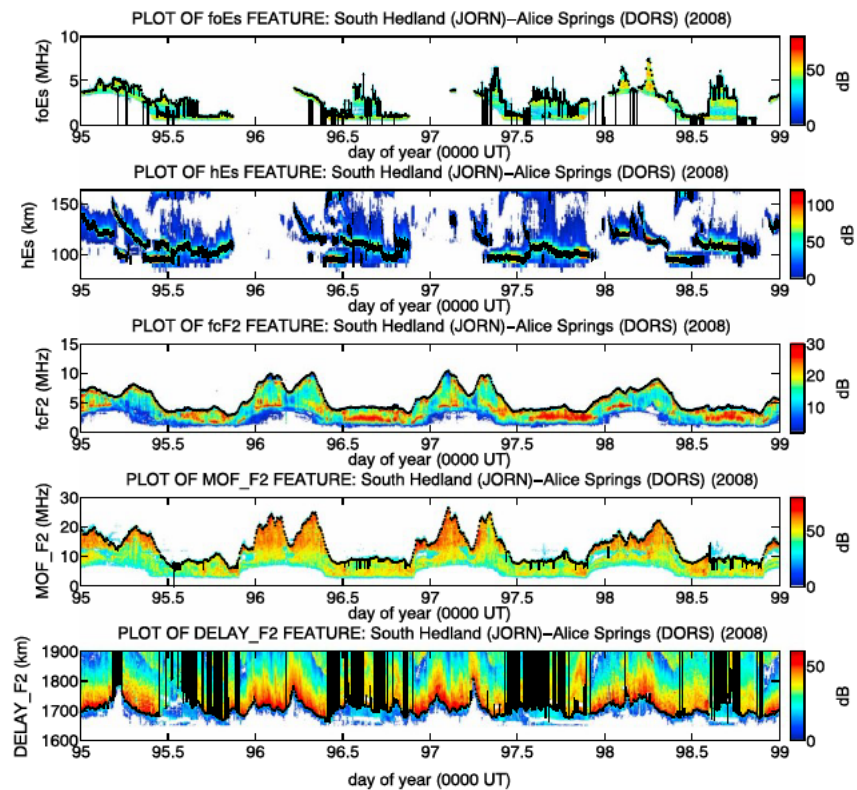
Figure 11. A 4-day time series of Es and F2 feature estimates (4–7 April 2008), overlaid on images of the frequency and delay intensity arrays from which the features were derived. The path, South Hedland to Alice Springs, is 1,597 km in length. In all panels, the black markers represent the center estimate, while the vertical lines show the uncertainties bounded by the lower and upper estimates. JORN = Jindalee Operational Radar Network; DORS = Digital Oblique Receiving System.

(Gardiner-Garden et al., 2011). Ionograms were created every 7.5 min on the OIS paths, and 3.75 min on the VIS. Local times are between 8 and 10 hr ahead of UT in this region.

5.1. Sample of Extracted Feature Estimates

In the absence of manually scaled truth data, the primary mechanism used to validate very large sets of extracted feature estimates is to overplot or difference the results against the corresponding (maximum/ mean) intensity arrays, fixed-frequency ionogram slices, or fitted trace peaks. In Figure 11, 4 days of Es and F2 features from a medium-length DORS path (South Hedland to Alice Springs, 1,597 km) are plotted on their respective frequency or delay intensity arrays, as described in section 4. The vertical uncertainty bars attached to each feature represent, for edge features, the spread in estimates across the three thresholds, or for peak features, the 6-dB half width of the fitted parabola.

Overall, the level of temporal smoothness in the features is good, despite the fact that each ionogram has been analyzed independently. The relatively small numbers of outliers are usually identifiable as having larger uncertainties. At other times, the uncertainty can be raised when the feature is spread or weak; for example, in the case of edge detections, the highest threshold may exceed all values in the intensity array, causing the uncertainty to become unbounded. This is most evident in the $delay_{F2}$ feature, which becomes poorly defined on some nights when spread-Es or spread-F conditions (Davies, 1990) are present.

The top two panels of Figure 11 show the Es features, f_{oEs} and h_{Es}, for up to two simultaneous layers. Whereas f_{oEs}, representing the upper equivalent VI frequency of partial Es reflection, is a weaker feature and consequently tends to be quite noisy, h_{Es} exhibits very clean time-varying characteristics. The
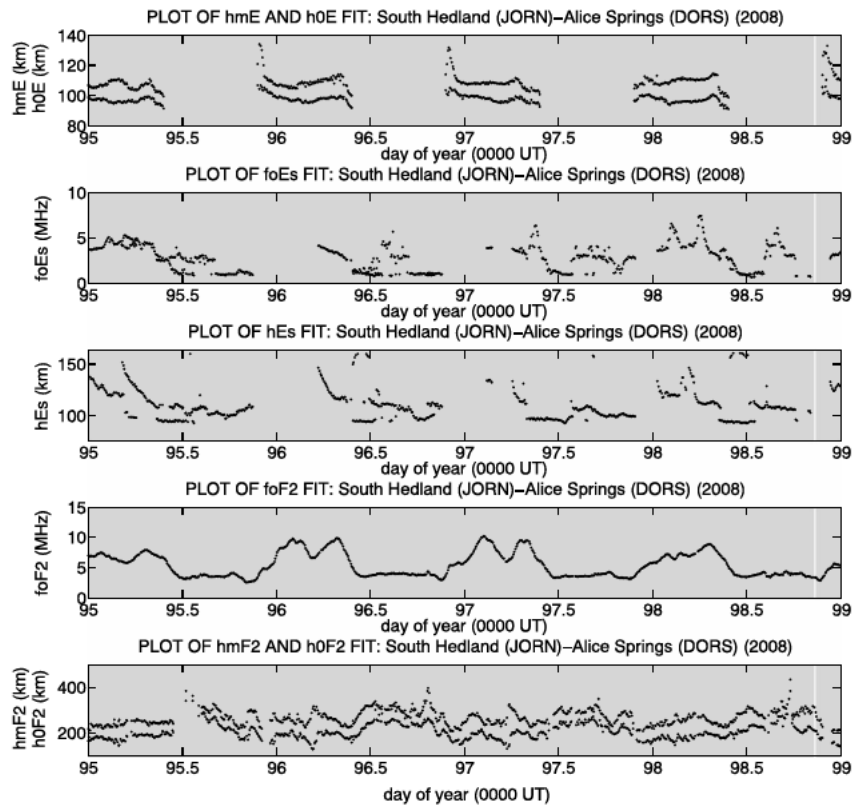
Figure 12. A 4-day time series of E, Es, and F2 fitted parameters (4–7 April 2008). Parameter estimates derived from climatology, due to incomplete fits, have been removed for clarity. All times with valid ionograms are shaded with a gray background, while those missing (of which there was only one ionogram, on day 98) are shaded white. JORN = Jindalee Operational Radar Network; DORS = Digital Oblique Receiving System.

patterns of Es layer descent in the second panel, and in such height-time intensity arrays more broadly, hint at the wind-shear forcing mechanism and resultant tidal periodicities. They are similar to the plots presented by Haldoupis et al. (2006) using VIS ionograms, albeit with the additional image masking and transformation steps, and incorporating all HF frequencies rather than subbands.

The F2 features, shown in the bottom three panels of Figure 11, tend to track the typical day-to-day and hour-to-hour variability well in the underlying intensity arrays. Apart from the occasional o/x ambiguities (on the scale of ~1 MHz), f_{cF2} and MOF_{F2} in particular are exceptionally reliable. The $delay_{F2}$ feature, by definition, is not necessarily smoothly varying over time, although it nevertheless remains fairly continuous in the absence of multihop and blanketing Es. However, when one-hop Es is detected, the multihop Es mask can overlap the one-hop F2 trace and shift the $delay_{F2}$ feature further up the trace.

5.2. Sample of Fitted Parameter Estimates

Figure 12 plots the same 4-day data set for the E, Es, and F2 fitted parameters. Again, it is emphasized that no temporal smoothing has been applied to these results; each ionogram fit has been obtained independently. For sake of clarity, no partially fitted parameters are included in any of these plots. Note also that the base heights of the E and F2 layers (h_{0E} and h_{0F2}) are plotted in place of the semithicknesses (y_{mE} and y_{mF2}), on the same set of axes as the peak heights (h_{mE} and h_{mF2}). As the Es parameters in the second and third panels are taken directly from the Es features, they are not discussed any further in this subsection.

The E height fitting offers robust performance throughout the day, and although estimates may get a little noisier around the dawn and dusk terminators, there are very few gross outliers in the top panel of

Table 1
Feature Extraction and Fit Statistics for Three Short-Range (300–600 km) DORS Oblique Paths in April 2008

Layer		Image features			Fitted parameters		
Name	Number of observations	Name	Number of valid est.	Median uncert.	Names	Complete fits	Partial fits
E	7909 (day only)	h0E	7346 (92.9%)	0.4 km	hmE, ymE	79.0%	-
		ymE	7346 (92.9%)	0.4 km			
Es	15782	hEs(1)	13395 (84.9%)	1.6 km	foEs(1), hEs(1)	84.9%	-
		hEs(2)	6349 (40.2%)				
		foEs(1)	13395 (84.9%)	60 kHz	foEs(2), hEs(2)	40.2%	-
		foEs(2)	6349 (40.2%)				
F2	15782	fcF2	15751 (99.8%)	20 kHz	foF2, hmF2, ymF2	73.5%	17.2%
		MOF_F2	15318 (97.1%)	40 kHz			
		delay_F2	13352 (84.6%)	0.8 km			

Figure 12. Furthermore, the general diurnal variation in h_{mE} appears to be consistent with modeling based on the simple Chapman theory of layer production (Titheridge, 2000). While periods of blanketing Es at heights below the E layer peak (often shortly before dusk) have the potential to cause h_{mE} to be underestimated, most of these fits are automatically rejected on the basis of their unrealistic layer thickness.

The F2 fitted parameters, in the bottom two panels of Figure 12, are also well behaved on the whole, with a strong sense of continuity over time. When f_{oF2} is at its most depressed levels, there is an increased likelihood of the h_{mF2} and y_{mF2} (or h_{oF2}) parameters remaining unfitted, but this only occurs in a relatively small number of cases. The performance around the middle of UT day 95 was made worse by blanketing Es. Characteristic F2 signatures, such as the midday f_{oF2} bite-out (Kohl & King, 1967; Lynn et al., 2014) and quasiperiodic height fluctuations, in which h_{oF2} commonly lags h_{mF2} due to the downward phase progression of traveling ionospheric disturbances (Cabrera et al., 2010; Pederick et al., 2017; Reinisch et al., 1997), are quite distinct in this data set.

5.3. Feature and Fit Statistics

Considering now a broader set of 18 DORS paths, and the whole of April 2008, the general performance of the algorithm can be evaluated with some simple statistics. The tables below summarize the extracted features in terms of their detection rates (number of valid estimates) and median uncertainties, and the fitted parameters in terms of their complete and partial fit rates. In this context, Es and F2 feature uncertainties are defined as half of the difference between the upper and lower bounds, although these may not be strictly symmetric about the center estimate. The E height uncertainties are the standard errors returned by the linear regression. Such uncertainties represent just one component of the total error, which may also include mode misclassifications and instrumental errors. Poor feature estimates, including those encountered during disturbed propagation conditions, can often be identified as having uncertainties many times larger than the reported medians.

Be aware that a successful fit is not necessarily a good fit, as judged by manual inspection; it just indicates that the features have been matched to the synthetic trace. As performance is somewhat dependent on path length, the OIS paths have been split into short (300–600 km), medium (600–2,000 km), and long (2,000–3,000 km) ground ranges, which are presented separately in Tables 1, 2, and 3, respectively. Note that partial fits only apply to the F2 layer (hence the blank table entries for E/Es fitted parameters), and E/Es processing is not possible on the long paths due to one-hop obscuration.

As described in section 4.2, the E layer heights are fitted directly to a line feature in the transformed ionogram. As such, the only difference between the feature and fit statistics in this row of the tables is that the latter has quality checks applied. From the subset of daytime ionograms only, the fit success rate is 85.1% for medium-length paths after applying the quality checks, although this falls to 79.0% on shorter paths. Most of the failed or rejected fits occur near the dawn and dusk terminators, when f_{oE} is reduced, or in the presence of blanketing Es at low heights. On shorter paths, the geometry means there are often not enough valid trace pixels observed above the 2 MHz start frequency near dawn and dusk.

For Es extraction, the number of valid estimates has a significant seasonal dependence, which scales according to the occurrence rates of the layer itself. As such, Es detection rates cannot be interpreted as a direct

Table 2
Feature Extraction and Fit Statistics for 13 Midrange (600–2000 km) DORS Oblique Paths in April 2008

Layer		Image features			Fitted parameters		
Name	Number of observations	Name	Number of valid est.	Median uncert.	Names	Complete fits	Partial fits
E	35984 (day only)	h0E	35614 (99.0%)	0.3 km	hmE, ymE	85.1%	-
		ymE	35614 (99.0%)	0.3 km			
Es	72383	hEs(1)	59971 (82.9%)	3 km	foEs(1), hEs(1)	82.8%	-
		hEs(2)	14290 (19.7%)				
		foEs(1)	59959 (82.8%)	30 kHz	foEs(2), hEs(2)	19.7%	-
		foEs(2)	14284 (19.7%)				
F2	72383	fcF2	72355 (100.0%)	30 kHz	foF2, hmF2, ymF2	86.9%	11.3%
		MOF_F2	72317 (99.9%)	60 kHz			
		delay_F2	69196 (95.6%)	1.5 km			

measure of the algorithm's success when the layer is not always present. Given the lower Es occurrence rates around equinox, feature extraction rates of 80–85% for the first layer and 20–40% for the second layer are within expectations. The second layer is the weaker of the two layers, as measured from the delay intensity array, and may be higher or lower in height than the first layer.

With the DORS ionogram resolution being fixed, shorter paths benefit from better equivalent height resolution. Along with their typically higher SNR levels, this translates to improved Es performance statistics, both in terms of Es detection rates (particularly for the second layer) and h_{Es} uncertainties, although the reduced equivalent VI frequency resolution means that f_{oEs} uncertainties tend to be slightly larger. After some minor quality checks, what remains are deemed to be the fitted Es parameters, with success rates comparable to the detection rates quoted above.

The F2 feature extraction rates are all consistently greater than 95% on medium-length paths, which make up the majority of the data set. While the numbers for f_{cF2} and MOF_{F2} remain similarly good on the shorter and longer paths, the $delay_{F2}$ performance does degrade. The cause is thought to be the increased propagation loss and poorer antenna sensitivity (Ayliffe et al., 2018) on long paths (with the $delay_{F2}$ feature arriving below 10° elevation), and the very limited support for nighttime propagation on short paths (often with only the faintest ionogram trace left below the predawn MOF). On all but the longest paths, median uncertainties on the scale of a single frequency or group delay cell are typical: 20–30 kHz in f_{cF2}, 40–60 kHz in MOF_{F2}, and 0.8–1.5 km in $delay_{F2}$ (cf. 40 kHz and 1.5-km resolution for DORS ionograms).

Complete F2 fits to all three features are achieved in 86.9% of medium-length ionograms, with a further 11.3% of partial fits boosting the total fit rate to over 98%. At both shorter and longer path lengths, the overall F2 fit rate remains above 90%, but a greater proportion of these are partial fits, largely due to the poorer $delay_{F2}$ detection rates.

5.4. Comparisons Between Closely Spaced Paths

The JORN RTIM combines a large number of both VIS and OIS observations to produce real time spatial maps of the electron density over Australia. The underlying philosophy is that ionograms from different path lengths and orientations provide complementary and consistent characterizations of the one-hop midpoint

Table 3
Feature Extraction and Fit Statistics for Two Long-Range (2000–3000 km) DORS Oblique Paths in April 2008

Layer		Image features			Fitted parameters		
Name	Number of observations	Name	Number of valid est.	Median uncert.	Names	Complete fits	Partial fits
F2	10454	fcF2	10449 (100.0%)	40 kHz	foF2, hmF2, ymF2	67.2%	28.0%
		MOF_F2	10435 (99.8%)	140 kHz			
		delay_F2	7926 (75.8%)	15 km			

Note. The geometry obscures one-hop E and Es propagation on the long-range paths; as such, these layers are not fitted.
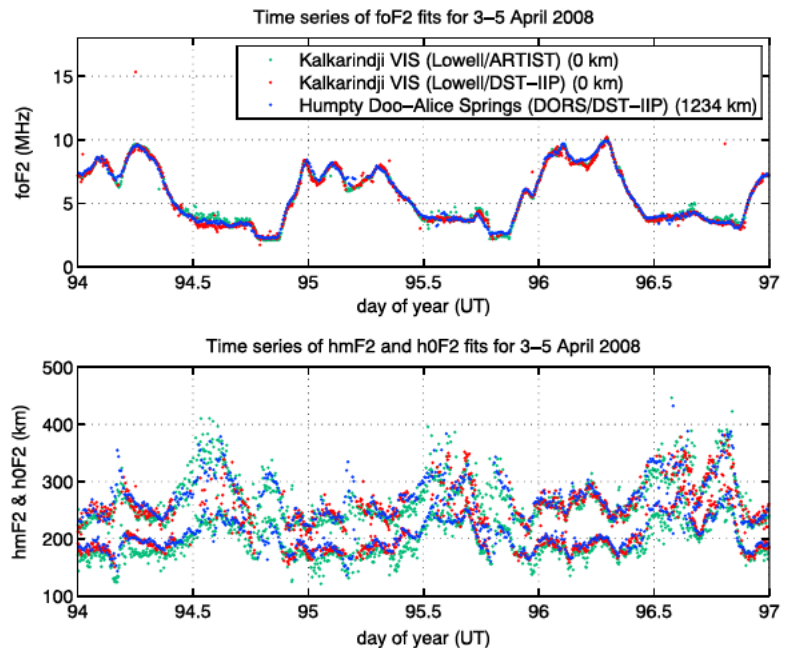
Figure 13. Comparison of F2 critical frequency and peak/base height parameters from both a DORS OIS and a Lowell Digisonde DPS-1 VIS, about a similar one-hop midpoint. DORS = Digital Oblique Receiving System; OIS = oblique incidence sounders; VIS = vertical incidence sounder.

profiles. Furthermore, it is important that DORS OIS systems running the DST-IIP algorithm do not output results that conflict with those from legacy systems, such as the Lowell Digisonde DPS-1 network still operating as part of JORN. It is noted that this does not represent the latest generation of the Digisonde in either hardware or software.

To test the validity of these assumptions, Figure 13 shows 3 days of DORS/DST-IIP fitted F2 parameters, from a 1,234-km oblique path (Humpty Doo to Alice Springs, in blue), alongside those from a Lowell Digisonde (Kalkarindji, in teal) running ARTIST-4 (Reinisch et al., 2005). For comparison, the Digisonde ionograms have also been processed with the DST-IIP algorithm (in red). As mapped in Figure 1, the OIS and VIS one-hop midpoints are separated by 180 km, well within the typical correlation distances of large-scale midlatitude F2 variability, and in general, the results are in very good agreement; this applies even through the dawn and dusk terminator periods when stronger gradients may be present. However, there is a slight systematic offset in the F2 heights from the DST-IIP and ARTIST algorithms (the latter being ~10 km lower on average) that could be due to differences in the underlying E/F1 profile parameterization or the position of the trace edge with respect to the peak of the image power.

Figure 14 shows another example, this time for three DORS OIS paths with 100- to 200-km midpoint separation. The paths, of various lengths and very different orientations, observe a similar region of the ionosphere, just south of the Gulf of Carpentaria (centered about −17.5°N, 137.5°E). Other than the nighttime period on day 95 (12–18 UT), when ionospheric disturbances were particularly prevalent in the ionograms, the fitted f_{oF2} parameters are largely consistent; this is despite the paths having different aspects to the geomagnetic field, and therefore differing amounts of o/x splitting. The fitted heights are naturally noisier at night, due to a combination of depressed f_{oF2} and spread-F conditions, both of which were elevated in severity during this period near sunspot minimum. Nevertheless, much of the height variability appears to represent real geophysical effects, strongly correlated between closely spaced paths.

5.5. Known Limitations

As with all autoscaling algorithms, the challenge is achieving robust performance under a variety of disturbed propagation conditions. The presence of post-midnight spread-F, intense multihop/blanketing Es,
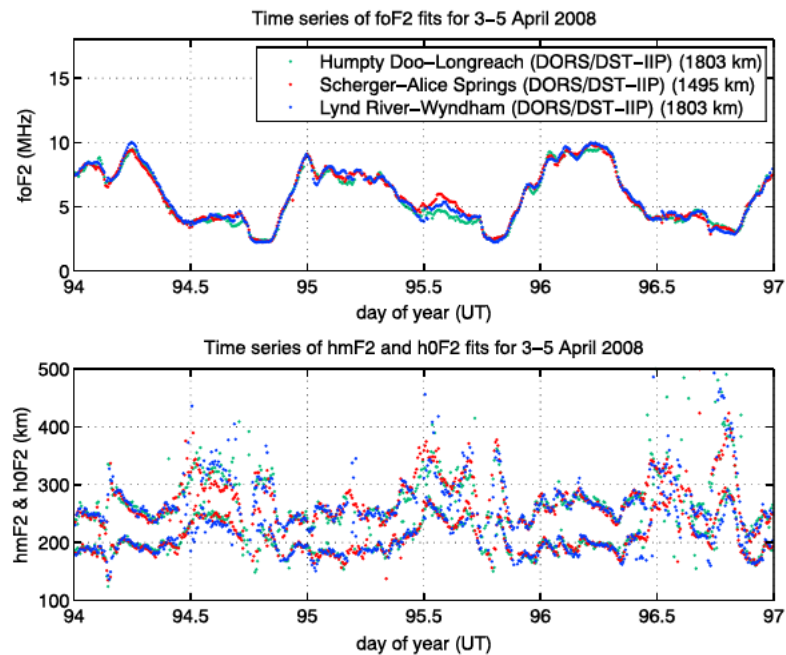
Figure 14. Comparison of F2 critical frequency and peak/base height parameters from three DORS OIS paths of various lengths and orientations, about a similar one-hop midpoint. DORS = Digital Oblique Receiving System.

depressed f_{oF2} following the onset of geomagnetic storms, and trace perturbations caused by medium-scale traveling ionospheric disturbances are among the most common adverse conditions observed at midlatitudes. However, this list is by no means exhaustive, and all manner of anomalous ionograms can of course present themselves to the algorithm.

The feature uncertainty estimates, being measures of peak width or edge steepness, go some of the way toward characterizing conditions like spread-F and multihop Es, which lead to poorly defined or ambiguous features. In many of these cases, a fit may still be possible, but it will come with a reduced level of confidence. Other conditions, like depressed f_{oF2} and medium-scale traveling ionospheric disturbances, may perturb the trace in such a way as to render the profile parameterization inadequate and unusable, that is, given errors in the underlying (model) E and F1 ionization, the lack of electron density valleys, and the assumption of no off-angle reflections. A failed fit is therefore considerably more likely in such cases.

Mode misclassifications, such as one-hop F2 being seduced by multihop returns and Es by transient meteor echoes, are comparatively rare; typically <1% of ionograms (but slightly higher if o/x mismatches are also included). The special multihop Es influence flag, described earlier, aims to detect some of these potential ambiguities. However, on paths of more than about 2,000 km, one-hop Es extraction and, therefore, multihop Es masking are not possible, so the F2 fitting can be more susceptible to unknown errors. While the E height fitting remains fairly robust to misclassification, low Es layers may occasionally remain in the ionogram after masking, leading to artificially thin E layer fits. These same low Es layers may escape detection altogether.

There are also system and environmental limitations that can cause reduced SNR or trace gaps in the ionograms, making feature extraction more difficult. Longer paths (>2,000 km), for which path losses are much greater and antenna gain patterns less favorable, as well as short and quasi-vertical paths (<600 km) at night, for which propagation is often limited to the congested band below 5 MHz, are geometries that regularly suffer from low SNR in the F2 trace. The DST-IIP algorithm includes a low SNR mode to reduce the feature extraction thresholds in such marginal cases, at the cost of higher misdetection rates. Gaps in the trace may also result from strong RFI and barred frequencies, although newer-generation sounders such as DORS are better equipped to handle these effects.

The impact of horizontal gradients and inhomogeneities on the spherically symmetric ionosphere assumption is harder to quantify, although the level of consistency in the two-hop synthetic trace is a good indicator (see Figure 4, for example). Around the dawn and dusk terminators, large-scale gradients become more likely, and nighttime irregularities can lead to significant out-of-plane propagation, in the form of "satellite" traces and post-midnight spread-F. More work is being undertaken to extend the automatic analysis to two-hop returns, both for characterization of gradients and support for very long (>3,000 km) oblique paths, on which one-hop F2 returns are obscured.

Finally, it is noted that the DST-IIP algorithm has only been tuned for midlatitude ionograms thus far. Equatorial and auroral/polar regions are expected to present additional challenges, such as trans-equatorial (ducted) modes and extreme spreading due to small-scale high-latitude irregularities, that have not been encountered in the midlatitude test data set presented here.

6. Conclusions

An algorithm has been described that automatically extracts trace features and fits electron density profile parameters for a general amplitude-scaled OIS or VIS ionogram. The parameterized results, describing the E, Es, and F2 layers, are specifically designed for direct assimilation into the JORN RTIM. In this context, feature uncertainty estimates based on real time assessment (rather than historical experience) are an important output for quality evaluation and removal of outliers. The DST-IIP algorithm has been extensively tested on high-resolution OIS ionograms from DST Group's DORS sounder, for which it is now part of the on-board processing, and also a broader set of legacy sounders including the Australian Lowell Digisonde network.

The algorithm performance is found to be robust under a wide variety of midlatitude propagation conditions, although this is of course dependent on the quality of ionograms. The main mechanism for evaluating performance, as described in this paper, is the inspection of features and synthetic trace slices overlaid on intensity arrays and slices of the input ionogram images. Medium-length oblique paths (nominally between 600 and 2,000 km) are found to produce the most reliable results across both day and night.

For the April 2008 DORS data set, typical F2 feature uncertainties were on the scale of a single image cell (20–60 kHz in frequency and 0.8–1.5 km in group delay), and F2 fit success rates were well over 90% including partial fits. These statistics degraded slightly on longer paths (>2,000 km). After fit quality checks, E layer heights were returned for around 80–85% of daytime ionograms on short- to medium-length paths. The Es detection rates were likewise very strong, with at least one layer found in more than 80% of ionograms. It has also been demonstrated that consistent fits can be achieved on multiple OIS and VIS paths with similar one-hop midpoints, which underpins a key assumption of the JORN RTIM.

While further development and validation is required to address some known limitations, such as characterization of disturbed propagation conditions, mode misclassifications, and multihop extraction, these are not seen to be critical issues. Indeed, the outputs are already in routine use by the JORN RTIM, providing improved frequency advice and coordinate registration to Australian HF radar operations.

References

Ayliffe, J. K., Durbridge, L. J., Frazer, G. J., Gardiner-Garden, R. S., Heitmann, A. J., Praschifka, J., et al. (2018). The DST Group high-fidelity, multi-channel oblique incidence Ionosonde. *Radio Science*, *53*. https://doi.org/10.1029/2018RS006681

Barnes, R. (1995). Analysis of E_s traces from a calibrated oblique ionosonde. *Journal of Atmospheric and Terrestrial Physics*, *57*(14), 1753–1761. https://doi.org/10.1016/0021-9169(95)00095-J

Barnes, R. I., Gardiner-Garden, R. S., & Harris, T. J. (2000). Real time ionospheric models for the Australian Defence Force. Paper presented at the Workshop on the Applications of Radio Science WARS-2000, 122–135. Retrieved from http://www.sws.bom.gov.au/IPSHosted/NCRS/wars/wars2000/

Bennett, J. A., Chen, J., & Dyson, P. L. (1991). Analytic ray tracing for the study of HF magneto-ionic radio propagation in the ionosphere. *Applied Computational Electromagnetics Society Journal*, *6*(1), 192–210.

Bennett, J. A., & Dyson, P. L. (2002). Effect of horizontal ionospheric structure on HF radio systems. Paper presented at the 27th URSI General Assembly, Maastricht, The Netherlands.

Bilitza, D. (1990). International Reference Ionosphere 1990 (NSSDC/WDC-A-R&S Report 90–22). Greenbelt, MD: National Space Science Data Center/World Data Center A for Rockets and Satellites.

Bilitza, D. (2001). International Reference Ionosphere 2000. *Radio Science*, *36*(2), 261–275. https://doi.org/10.1029/2000RS002432

Bilitza, D., Altadill, D., Zhang, Y., Mertens, C., Truhlik, V., Richards, P., et al. (2014). The International Reference Ionosphere 2012—A model of international collaboration. *Journal of Space Weather and Space Climate*, *4*, A07. https://doi.org/10.1051/swsc/2014004

Bilitza, D., & Reinisch, B. W. (2008). International Reference Ionosphere 2007: Improvements and new parameters. *Advances in Space Research*, *42*(4), 599–609. https://doi.org/10.1016/j.asr.2007.07.048

Cabrera, M. A., Pezzopane, M., Zuccheretti, E., & Ezquer, R. G. (2010). Satellite traces, range spread-F occurrence, and gravity wave propagation at the southern anomaly crest. *Annales Geophysicae*, *28*(5), 1133–1140. https://doi.org/10.5194/angeo-28-1133-2010

Chen, J., Bennett, J. A., & Dyson, P. L. (1990). Automatic fitting of quasi-parabolic segments to ionospheric profiles with application to ground range estimation for single-station location. *Journal of Atmospheric and Terrestrial Physics*, *52*(4), 277–288. https://doi.org/10.1016/0021-9169(90)90095-5

Chen, J., Bennett, J. A., & Dyson, P. L. (1992). Synthesis of oblique ionograms from vertical ionograms using quasi-parabolic segment models of the ionosphere. *Journal of Atmospheric and Terrestrial Physics*, *54*(3–4), 323–331. https://doi.org/10.1016/0021-9169(92)90011-9

Chen, Z., Wang, S., Zhang, S., Fang, G., & Wang, J. (2013). Automatic scaling of F layer from ionograms. *Radio Science*, *48*, 334–343. https://doi.org/10.1002/rds.20038

Clette, F., & Lefevre, L. (2016). The new sunspot number: Assembling all corrections. *Solar Physics*, *291*(9-10), 2629–2651. https://doi.org/10.1007/s11207-016-1014-y

Clette, F., Svalgaard, L., Vaquero, J. M., & Cliver, E. W. (2014). Revisiting the sunspot number: A 400-year perspective on the solar cycle. *Space Science Reviews*, *186*(1–4), 35–103. https://doi.org/10.1007/s11214-014-0074-2

Croft, T. A., & Hoogasian, H. (1968). Exact ray calculations in a quasi-parabolic ionosphere with no magnetic field. *Radio Science*, *3*(1), 69–74. https://doi.org/10.1002/rds19683169

Davies, K. (1990). *Ionospheric Radio, IEEE Electromagnetic Waves Series* (Vol. 31). London, UK: Peter Peregrinus Ltd.

Ding, Z., Ning, B., Wan, W., & Liu, L. (2007). Automatic scaling of F2-layer parameters from ionograms based on the empirical orthogonal function (EOF) analysis of ionospheric electron density. *Earth, Planets and Space*, *59*(1), 51–58. https://doi.org/10.1186/BF03352022

DuCharme, E. D., Petrie, L. E., & Eyfrig, R. (1971). A method for predicting the F_1 layer critical frequency. *Radio Science*, *6*(3), 369–378. https://doi.org/10.1029/RS006i003p00369

DuCharme, E. D., Petrie, L. E., & Eyfrig, R. (1973). A method for predicting the F_1 layer critical frequency based on the Zurich smoothed sunspot number. *Radio Science*, *8*(10), 837–839. https://doi.org/10.1029/RS008i010p00837

Dyson, P. L., & Bennett, J. A. (1988). A model of the vertical distribution of the electron concentration in the ionosphere and its application to oblique propagation studies. *Journal of Atmospheric and Terrestrial Physics*, *50*(3), 251–262. https://doi.org/10.1016/0021-9169(88)90074-8

Fabrizio, G. A. (2013). *High frequency over-the-horizon radar*. New York: McGraw-Hill.

Fox, M. W., & Blundell, C. (1989). Automatic scaling of digital ionograms. *Radio Science*, *24*(6), 747–761. https://doi.org/10.1029/RS024i006p00747

Galkin, I. A., & Reinisch, B. W. (2008). *The new ARTIST 5 for all Digisondes, Ionosonde Network Advisory Group Bulletin* (Vol. 69). Sydney, New South Wales, Australia: Ionospheric Prediction Service.

Gardiner-Garden, R., Heitmann, A., & Frazer, G. (2011). Variability observed in a high-fidelity model of the ionosphere based on a dense network of vertical and oblique incidence sounders. Paper presented at the Ionospheric Effects Symposium (IES2011), Alexandria, VA.

Gardiner-Garden, R., Heitmann, A., & Northey, B. (2018). A parametric model of the ionospheric electron density profile for JORN (DST Group Technical Note TN-1722). Edinburgh, South Australia, Australia: National Security and ISR Division, Defence Science and Technology Group. Retrieved from https://www.dst.defence.gov.au/publication/parametric-model-ionospheric-electron-density-profile-jorn

Guiducci, A., Melen, R., Nardi, D., Neri, F., Nesti, F., & Quaglia, G. (1983). Automatic scaling of ionograms by the method of structural description. *Pattern Recognition*, *16*(5), 489–499. https://doi.org/10.1016/0031-3203(83)90053-5

Haldoupis, C., Meek, C., Christakis, N., Pancheva, D., & Bourdillon, A. (2006). Ionogram height-time-intensity observations of descending sporadic E layers at mid-latitude. *Journal of Atmospheric and Solar-Terrestrial Physics*, *68*(3–5), 539–557. https://doi.org/10.1016/j.jastp.2005.03.020

Harris, T. J., Cervera, M. A., Pederick, L. H., & Quinn, A. D. (2017). Separation of O/X polarization modes on oblique ionospheric soundings. *Radio Science*, *52*, 1522–1533. https://doi.org/10.1002/2017RS006280

Harris, T. J., Quinn, A. D., & Pederick, L. H. (2016). The DST group ionospheric sounder replacement for JORN. *Radio Science*, *51*, 563–572. https://doi.org/10.1002/2015RS005881

Headrick, J. M., & Anderson, S. J. (2008). HF over-the-horizon radar. In M. Skolnik (Ed.), *Radar handbook* (3rd ed., Chap. 20, pp. 20.1–20.83). New York: McGraw-Hill.

Hill, J. R. (1979). Exact ray paths in a multisegment quasi-parabolic ionosphere. *Radio Science*, *14*(5), 855–861. https://doi.org/10.1029/RS014i005p00855

Huang, X., & Reinisch, B. W. (1982). Automatic calculation of electron density profiles from digital ionograms. 2. True height inversion of topside ionograms with the profile-fitting method. *Radio Science*, *17*(4), 837–844. https://doi.org/10.1029/RS017i004p00837

Igi, S. (1992). Automatic ionogram processing system. 3. A new method of displaying ionospheric characteristics. *Journal of the Communications Research Laboratory*, *39*(2), 381–402.

Ippolito, A., Scotto, C., Francis, M., Settimi, A., & Cesaroni, C. (2015). Automatic interpretation of oblique ionograms. *Advances in Space Research*, *55*(6), 1624–1629. https://doi.org/10.1016/j.asr.2014.12.025

Ivanov-Kholodny, G. S., Kishcha, P. V., & Zhivolup, T. G. (1998). A mid-latitude E-layer peak height model. *Advances in Space Research*, *22*(6), 767–770. https://doi.org/10.1016/S0273-1177(98)00096-9

Jiang, C., Yang, G., Zhao, Z., Zhang, Y., Zhu, P., Sun, H., & Zhou, C. (2014). A method for the automatic calculation of electron density profiles from vertical incidence ionograms. *Journal of Atmospheric and Solar-Terrestrial Physics*, *107*, 20–29. https://doi.org/10.1016/j.jastp.2013.10.012

Kohl, H., & King, J. W. (1967). Atmospheric winds between 100 and 700 km and their effects on the ionosphere. *Journal of Atmospheric and Terrestrial Physics*, *29*(9), 1045–1062. https://doi.org/10.1016/0021-9169(67)90139-0

Kouris, S. S., & Muggleton, L. M. (1973). Diurnal variation in the E-layer ionization. *Journal of Atmospheric and Terrestrial Physics*, *35*(1), 133–139. https://doi.org/10.1016/0021-9169(73)90221-3

Lagarias, J. C., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM Journal of Optimization*, *9*(1), 112–147. https://doi.org/10.1137/S1052623496303470

Lynn, K. J. W. (2018). Histogram-based ionogram displays and their application to autoscaling. *Advances in Space Research*, *61*(5), 1220–1229. https://doi.org/10.1016/j.asr.2017.12.019

Lynn, K. J. W., Gardiner-Garden, R. S., & Heitmann, A. (2014). The spatial and temporal structure of twin peaks and midday bite out in f_oF_2 (with associated height changes) in the Australian and South Pacific midlatitude ionosphere. *Journal of Geophysical Research: Space Physics*, *119*, 10,294–10,304. https://doi.org/10.1002/2014JA020617

Mathews, J. D. (1998). Sporadic E: Current views and recent progress. *Journal of Atmospheric and Solar-Terrestrial Physics*, *60*(4), 413–435. https://doi.org/10.1016/S1364-6826(97)00043-6

McNamara, L. F. (1991). *The ionosphere: Communications, surveillance, and direction finding*. Malabar, FL: Krieger Publishing Company.

Muggleton, L. M. (1975). A method of predicting foE at any time and place. *ITU Telecommunication Journal*, *42*(7), 413–418.

Pederick, L. H., Cervera, M. A., & Harris, T. J. (2017). Interpreting observations of large-scale traveling ionospheric disturbances by ionospheric sounders. *Journal of Geophysical Research: Space Physics*, *122*, 12,556–12,569. https://doi.org/10.1002/2017JA024337

Pezzopane, M., & Scotto, C. (2004). Software for the automatic scaling of critical frequency *foF2* and *MUF(3000)F2* from ionograms applied at the Ionospheric Observatory of Gibilmanna. *Annals of Geophysics*, *47*(6), 1783–1790. https://doi.org/10.4401/ag-3375

Pezzopane, M., & Scotto, C. (2007). Automatic scaling of critical frequency *foF2* and *MUF(3000)F2*: A performance comparison between Autoscala and ARTIST 4.5 on Rome data. *Radio Science*, *42*, RS4003. https://doi.org/10.1029/2006RS003581

Pezzopane, M., & Scotto, C. (2008). A method for automatic scaling of F1 critical frequencies from ionograms. *Radio Science*, *43*, RS2S91. https://doi.org/10.1029/2007RS003723

Phanivong, B., Chen, J., Dyson, P. L., & Bennett, J. A. (1995). Inversion of oblique ionograms including the Earth's magnetic field. *Journal of Atmospheric and Terrestrial Physics*, *57*(14), 1715–1721. https://doi.org/10.1016/0021-9169(95)00092-G

Piggott, W. R., & Rawer, K. (1978). U.R.S.I. handbook of ionogram interpretation and reduction (Report UAG-23A, 2nd ed., revised). Boulder, CO: World Data Centre A for Solar-Terrestrial Physics, NOAA.

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1992). *Numerical recipes in C: The art of scientific computing* (2nd ed.). Cambridge, UK: Cambridge University Press.

Redding, N. J. (1996). The autoscaling of oblique ionograms (DSTO Research Report RR-074). Edinburgh, South Australia, Australia: Information Technology Division, Defence Science and Technology Organisation.

Reilly, M. H., & Kolesar, J. D. (1989). A method for real height analysis of oblique ionograms. *Radio Science*, *24*(4), 575–583. https://doi.org/10.1029/RS024i004p00575

Reinisch, B. W., Galkin, I. A., Khmyrov, G. M., Kozlov, A. V., Bibl, K., Lisysyan, I. A., et al. (2009). New Digisonde for research and monitoring applications. *Radio Science*, *44*, RS0A24. https://doi.org/10.1029/2008RS004115

Reinisch, B. W., Gamache, R. R., Huang, X., & McNamara, L. F. (1988). Real-time electron density profiles from ionograms. *Advances in Space Research*, *8*(4), 63–72. https://doi.org/10.1016/0273-1177(88)90208-6

Reinisch, B. W., Haines, D. M., Bibl, K., Galkin, I., Huang, X., Kitrosser, D. F., et al. (1997). Ionospheric sounding in support of over-the-horizon radar. *Radio Science*, *32*(4), 1681–1694. https://doi.org/10.1029/97RS00841

Reinisch, B. W., & Huang, X. (1983). Automatic calculation of electron density profiles from digital ionograms. 3. Processing of bottomside ionograms. *Radio Science*, *18*(3), 477–492. https://doi.org/10.1029/RS018i003p00477

Reinisch, B. W., Huang, X., Galkin, I. A., Paznukhov, V., & 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. https://doi.org/10.1016/j.jastp.2005.01.009

Roughan, M. (1996). Trace extraction from oblique ionograms. Paper presented at the SME Applied Machine Vision '96 Conference, Cincinnati, OH.

Scotto, C. (2009). Electron density profile calculation technique for Autoscala ionogram analysis. *Advances in Space Research*, *44*(6), 756–766. https://doi.org/10.1016/j.asr.2009.04.037

Scotto, C., & Pezzopane, M. (2007). A method for automatic scaling of sporadic *E* layers from ionograms. *Radio Science*, *42*, RS2012. https://doi.org/10.1029/2006RS003461

Song, H., Hu, Y., Jiang, C., Zhou, C., Zhao, Z., & Zou, X. (2016). An automatic scaling method for obtaining the trace and parameters from oblique ionogram based on hybrid genetic algorithm. *Radio Science*, *51*, 1838–1854. https://doi.org/10.1002/2016RS005987

Su, F., Zhao, Z., Li, S., Yao, M., Chen, G., & Zhou, Y. (2012). Signal identification and trace extraction for the vertical ionogram. *IEEE Geoscience and Remote Sensing Letters*, *9*(6), 1031–1035. https://doi.org/10.1109/LGRS.2012.2189350

Titheridge, J. E. (1988). The real height analysis of ionograms: A generalized formulation. *Radio Science*, *23*(5), 831–849. https://doi.org/10.1029/RS023i005p00831

Titheridge, J. E. (2000). Modelling the peak of the ionospheric E-layer. *Journal of Atmospheric and Solar-Terrestrial Physics*, *62*(2), 93–114. https://doi.org/10.1016/S1364-6826(99)00102-9

Tsai, L.-C., & Berkey, F. T. (2000). Ionogram analysis using fuzzy segmentation and connectedness techniques. *Radio Science*, *35*(5), 1173–1186. https://doi.org/10.1029/1999RS002170

Turley, M. D. E., Heitmann, A. J., & Gardiner-Garden, R. S. (2018). Ionogram RFI rejection using an auto-regressive interpolation process. *Radio Science*, *53*. https://doi.org/10.1029/2018RS006683

Whitehead, J. D. (1989). Recent work on mid-latitude and equatorial sporadic-E. *Journal of Atmospheric and Terrestrial Physics*, *51*(5), 401–424. https://doi.org/10.1016/0021-9169(89)90122-0

Wright, J. W., Laird, A. R., Obitts, D., Violette, E. J., & McKinnis, D. (1972). Automatic N (h, t) profiles of the ionosphere with a digital ionosonde. *Radio Science*, *7*(11), 1033–1043. https://doi.org/10.1029/RS007i011p01033

Zabotin, N. A., Wright, J. W., & Zhbankov, G. A. (2006). NeXtYZ: Three-dimensional electron density inversion for dynasonde ionograms. *Radio Science*, *41*, RS6S32. https://doi.org/10.1029/2005RS003352

Zheng, H., Ji, G., Wang, G., Zhao, Z., & He, S. (2013). Automatic scaling of F layer from ionograms based on image processing and analysis. *Journal of Atmospheric and Solar-Terrestrial Physics*, *105-106*, 110–118. https://doi.org/10.1016/j.jastp.2013.09.007

Zuccheretti, E., Tutone, G., Sciacca, U., Bianchi, C., & Arokiasamy, B. J. (2003). The new AIS-INGV digital ionosonde. *Annals of Geophysics*, *46*(4), 647–659. https://doi.org/10.4401/ag-4377

F.4 Turley et al., *Radio Science*, 2019

The following is a reprint of Turley, M. D. E., Heitmann, A. J., & Gardiner-Garden, R. S. (2019), 'Ionogram RFI rejection using an auto-regressive interpolation process', *Radio Science*, 54(1), pp. 135–150.

This paper provides a detailed explanation of the radio frequency interference (RFI) mitigation algorithm developed for DST Group's Digital Oblique Receiving System (DORS) OIS. The algorithm, based on autoregressive detection and interpolation of impulsive signals, was later included among the on-board processing for the ELOISE AoA ionosonde. The paper thus provides relevant background material to supplement the signal processing description in **Chapter 4** of this thesis.

The manuscript was written jointly by myself and Dr Mike Turley, and I served as the corresponding author. My own contributions to the work were mainly in the validation and testing of the algorithm, developed by Dr Turley, and its deployment on-board the sounder, which was carried out under the supervision of Dr Robert Gardiner-Garden. The paper was published during my PhD candidature.

Ionogram RFI Rejection Using an Autoregressive Interpolation Process

Mike D. E. Turley[1], Andrew J. Heitmann[1,2] (iD)**, and Robert S. Gardiner-Garden[3]** (iD)

[1]Defence Science and Technology Group, Edinburgh, South Australia, Australia, [2]School of Physical Sciences, University of Adelaide, Adelaide, South Australia, Australia, [3]Defence Science and Technology Group, Eveleigh, New South Wales, Australia

Abstract Defence Science and Technology Group's Digital Oblique Receiving System (DORS) is a direct-digital high-frequency ionosonde receiver, capable of collecting high-quality ionograms simultaneously on multiple oblique ionospheric paths. One of the key signal processing steps that run on board the receiver is a novel technique for the detection and removal of radio frequency interference (RFI). In the down-converted narrow-band time series, external RFI sources manifest as impulsive noise, often many orders of magnitude stronger than the ionosonde signal of interest. The RFI rejection technique applies a threshold detector to the whitened data, to locate the corrupt samples, and then replaces the bad data using forward and backward linear prediction, based on an autoregressive model of the desired ionosonde signal. Algorithm performance is evaluated using a set of simple statistics and ~6,000 DORS ionograms from central Australia. Compared to a more traditional approach of clipping the strong RFI impulses, the new technique significantly reduces instances of undersuppression and recovers more of the weaker propagation mode content in the ionogram. The result is a much cleaner image for the DORS automatic feature extraction and parameterized fitting technique.

Plain Language Summary Ionospheric sounders (or ionosondes) are a widely used instrument for measuring the vertical electron density structure of the ionosphere, an outer charged layer of the Earth's atmosphere that refracts high-frequency (HF) radio waves. An ionogram is a record of the path losses and delays for HF propagation via the ionosphere. However, ionograms can be potentially corrupted by the transmissions from other HF users, many of which are stronger than the ionosonde signal itself. This paper presents a novel signal processing technique for removing the influence of radio frequency interference in ionogram observations by Defence Science and Technology Group's new digital ionosonde. The result is a much cleaner standard of ionogram than traditionally seen, with improved sensitivity for detecting weaker propagation modes and fewer artifacts to degrade the automatic extraction of ionospheric parameters.

1. Introduction

Modern ionospheric sounders, such as Defence Science and Technology Group's Digital Oblique Receiving System (DORS) (Ayliffe et al., 2019) and the related Portable Remote Ionospheric Monitoring Equipment (PRIME) (Harris et al., 2016), provide high-fidelity observations of the ionosphere using direct-digital sampling of the high-frequency (HF) band. This design philosophy enables systems to be produced at relatively low cost and with excellent reliability and allows for multiple (simultaneous) ionosonde channels to be collected on one receiver. The raw data acquired by such ionosondes are processed to form ionograms that may be analyzed to extract image features using either human visual or computational interpretation systems. Improved clarity of the ionospheric trace features in the ionogram will in turn improve the performance of autoscaling algorithms such as the feature extraction and profile fitting technique described in Heitmann and Gardiner-Garden (2019). In this article we are concerned with improving the ionogram processing by reducing ionogram corruption caused by radio frequency interference (RFI).

For chirp sounders such as DORS and PRIME (configured for oblique incidence [OI] and vertical incidence [VI] reception, respectively), the signal of interest is a wide-band linear frequency-modulated continuous-wave (LFMCW) transmission of roughly 20 W. After propagation through the ionosphere, these signals arrive at the receiver with a power of 1 nW or less. The soundings are typically received on a single wide-band, omnidirectional antenna, with no directional information. Even for very remote receiving

sites, with few sources of local interference, there are a large number of RFI sources propagated via sky or surface wave that are orders of magnitude stronger than the signal of interest. These have the effect of (1) restricting the dynamic range allocation of the receiver (i.e., to prevent digital clipping of the strongest RFI source), and (2) introducing impulsive interference bursts in the down-converted channel. The latter occurs as the chirp sounder waveform scans momentarily through the frequency bands of the RFI sources and represents the RFI rejection problem that is the subject of this paper.

The typical approach to rejecting impulsive interference is to detect corrupted samples and then treat the problem as that of replacing missing samples. The effectiveness of such an approach relates to the fact that the impulses occupy only a small number of the total samples, which have a disproportionate effect on the frequency domain representation (in this case, the ionogram) due to their large amplitude compared to the signals of interest. The simplest replacement technique is to substitute the missing samples with zeros; this rejects the impulse but unfortunately also amplitude-modulates the desired signals, causing unwanted spectral domain leakage. Thus, the aim is to reduce or eliminate the deleterious effects of RFI in the ionogram while maintaining the ionospherically propagated sounder signals of interest (typically received across multiple modes). Other rudimentary techniques for impulsive noise mitigation include local median, nearest neighbor, or amplitude thresholding (Vaseghi, 2008), but these suffer from incomplete rejection and poor control of spectral leakage and are not robust to the various spectral characteristics of the signals of interest. The amplitude thresholding technique, also known as "clipping," is used throughout this paper as the ionosonde baseline technique for comparative purposes.

A more sophisticated approach to the missing samples problem is to use interpolation techniques to reconstruct the missing samples from the neighboring uncorrupted samples (Vaseghi, 2008). Common techniques include polynomial interpolators and model-based methods, such as linear prediction (Makhoul, 1975) using an autoregressive (AR) model (Burg, 1967). Nonparametric techniques using iterative time-frequency domain schemes with assumed spectral sparsity (Sacchi et al., 1998) and the Papoulis-Gerchberg technique (Gerchberg, 1974; Papoulis, 1975) are relatively computationally expensive.

AR interpolation has previously been applied to HF radar spectral estimation (Turley & Netherway, 1990), which was based on a beamforming method (Swingler & Walker, 1989). Similarly, a generalized expectation-maximization iterative procedure was demonstrated (Abramovich et al., 1993). Modifications to these iterative techniques have achieved very high convergence rates, excellent interference rejection, and robustness to model parameterization order and large missing data gaps (Turley, 2003). This methodology is well suited to the ionogram RFI rejection problem.

This paper is structured as follows. Section 2 presents a brief overview of ionogram spectral processing and the impact of RFI. The RFI rejection algorithm consists of two distinct stages: (1) RFI detection and (2) signal interpolation, both of which will be described in section 3. The detection stage breaks the acquired data into segments and flags bad samples whose amplitude grossly exceeds a low-order AR model of the segment. The interpolation stage then replaces blocks of bad samples with linear predictions from a higher-order AR-model-based filter. The performance of the algorithm, as implemented in MATLAB®, is evaluated against a set of key metrics in section 4, and finally conclusions are outlined in section 5. While equally applicable to OI and VI sounders, the examples and performance statistics presented here make use of DORS OI soundings. Examples of PRIME VI soundings, which adopt the same RFI rejection algorithm from this paper, can be found in Harris et al. (2016).

2. Effect of RFI on Ionogram Processing

For each chirp waveform ionosonde channel, being a time-frequency lane observed through the HF band, the DORS receiver applies stretch processing (deramping) to the raw samples from the analog-to-digital converter (Fabrizio, 2013) and decimates to produce a 1,562.5- or 3,125-Hz time series of in-phase and quadrature samples (Ayliffe et al., 2019). Ionograms are generated by segmenting the acquired digital samples and performing spectral analysis, using a windowed Fast Fourier Transform (FFT), to reveal the difference tones (between the received signal and local sweeping oscillator) and their frequencies that correspond to ionospheric propagation modes and their respective group delays. Each segment in time represents data captured from a particular subset of the HF band and contributes a single vertical slice to the ionogram. If desired, the segments may be temporally overlapped to reduce data exclusion due to Fourier transform

window tapering (Harris, 1978). For high dynamic range windows, little correlation occurs with up to 50% overlap.

As the receiver scans the HF spectrum, additive RFI power is observed as each cochannel interferer is passed by the swept sounder waveform. In the fundamental case of the RFI source being an unmodulated continuous-wave (CW) interferer, the demodulation of the chirp sounder waveform will transform the CW RFI into a chirp with sweep rate exactly opposite to that of the sounder. A typically modulated RFI source can be interpreted as a superposition of chirps. If the bandwidth of the RFI source is small (typically of the order of 10 kHz) compared to the coherent processing bandwidth of the ionogram (typically 200 kHz) and sweep rate large compared to the decimated sampling rate, then the RFI will resemble an impulse, that is, a short-lived, wide-bandwidth signal in the acquired data.

Upon ionogram processing, as described above, the impulsive RFI will thus manifest as a vertical stripe of raised power on the ionogram image. For the CW RFI source, this will be the spectrum of a chirp signal (hence its characteristic spread across the ionogram's group delay dimension). In general, a single ionogram processing segment will combine a number of different RFI sources in the raw data, leaving them inseparable in the amplitude-only image.

An example of RFI on a DORS OI ionogram (prior to RFI rejection) is shown in the top panel of Figure 1. This represents a 1,234-km path from Humpty Doo (−12.6°N, 131.3°E, near Darwin) to Alice Springs (−23.5°N, 133.7°E), a quiet midlatitude site in central Australia; all sources of RFI are expected to be remote (and thus ionospherically propagated).

In Figure 2, the probability distribution of the demodulated time series (corresponding to the top ionogram from Figure 1) provides an example of the amplitude characterization of impulsive interference. For this purpose, the time series was whitened using a third-order AR-based finite impulse response filter to exclude signals of interest and obtain a measure of impulsive interference (a low model order provides little transient distortion while rejecting strong narrow-band stationary signals). A reference exponential distribution is overlaid, based on independent in-phase and quadrature components of Gaussian noise, indicating that about 5% of the population are large-amplitude samples (caused by the impulses) that depart significantly from Gaussianity. For the same example, the temporal distribution of the impulsive interference for samples exceeding a detection threshold of 20 dB is plotted in Figure 3. This shows most narrow-band RFI only occupies the order of 10 samples and is sufficiently close in character to be treated as pseudoimpulsive noise.

3. Description of the RFI Rejection Algorithm

In this section we describe the RFI rejection algorithm as a two-stage process. In the first stage, the RFI is detected by exploiting its impulsive nature, and in the second stage, the desired signal is modeled then used to replace corrupt samples using AR interpolation. This structure is illustrated as a block diagram in Figure 4.

3.1. RFI Detection

The RFI detection stage operates on independent time-stepped segments of the acquired data spanning several hundred kilohertz of the sweep (typically 320 kHz). In this RFI detection stage, the time segments are longer than the ionogram processing segments and do not overlap. For each segment, a low-order AR(3) model is calculated using Burg's maximum entropy technique (Andersen, 1974; Burg, 1967, 1975). The AR coefficients are then used as a finite impulse response filter to whiten the data. In effect, stronger narrow-band signals (such as the ionospherically propagated signals from the sounder transmitter) are attenuated, with the residual output exposing impulsive noise. The low-order model was selected to reduce temporal distortion of impulses by the filter transients.

A threshold detector is subsequently applied to the whitened data samples (after median normalization) to identify those that are excessively large (>10 dB) and therefore deemed to be corrupted by impulsive noise (RFI). The threshold itself was manually tuned based on a range of typical DORS data. The whole detection process (i.e., AR model estimation, whitening, normalization, and thresholding) is repeated a second time using a higher threshold than the original (20 vs. 10 dB), with bad samples from the first attempt excluded in the AR estimation step to give an improved whitening filter. A small amount of dilation is applied to the
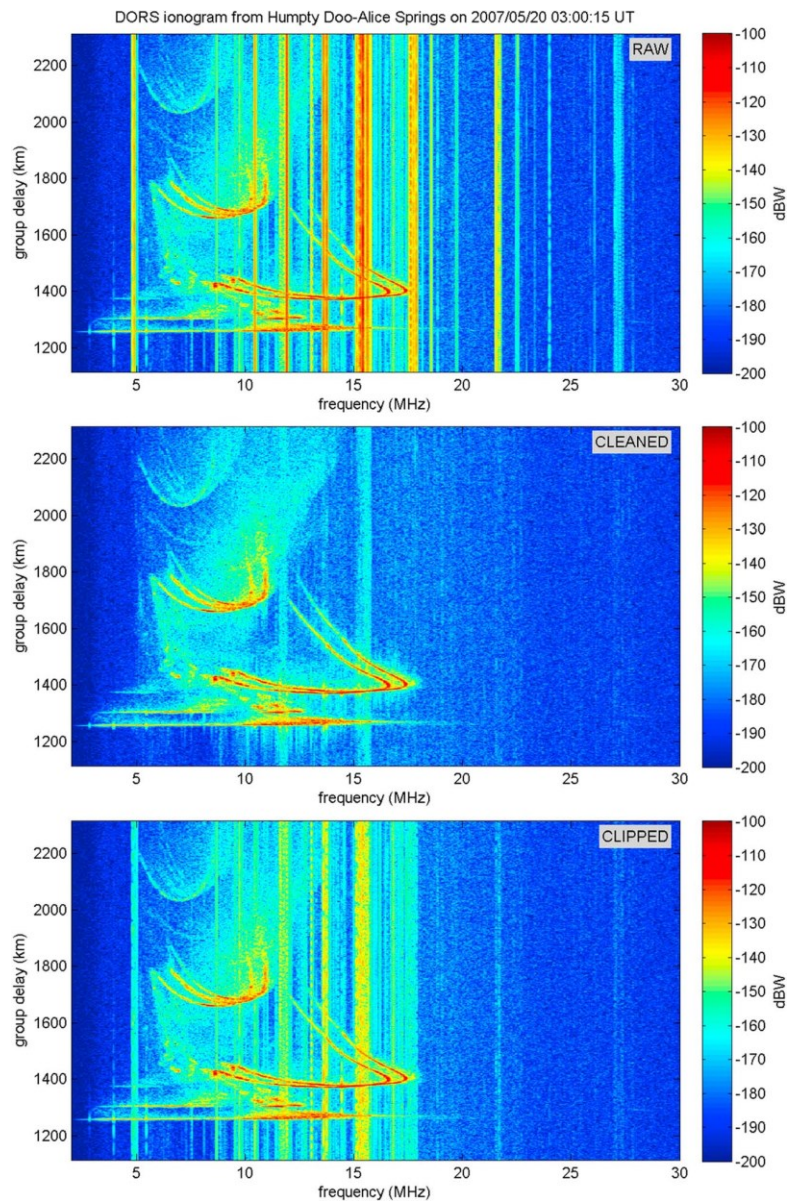
Figure 1. Comparison of raw (top panel), cleaned (middle panel), and clipped (bottom panel) ionograms from a single DORS OI sounding. DORS = Digital Oblique Receiving System; OI = oblique incidence.

bad samples output by the second-phase threshold detector to eliminate residual impulse edges. This defines the bad data map used in the signal interpolation stage.

For certain classes of sounder transmitters, there may be gaps in the waveform (e.g., due to mechanically switched filters in the power amplifier stage or regulatory barring of certain frequencies) that may be known or can be detected by a similar filtering and amplitude thresholding technique. These can be substantially longer than the typical span of ionospheric group delays and so are relatively straightforward to detect in the data series. Without correction, the integrated signal power is reduced in such a case, and, as the signal is not continuous across the entire ionogram segment (coherent processing interval), spectral leakage occurs in group delay. The effect in the ionogram domain is a vertical stripe of raised side lobe power, not unlike the manifestation of some RFI sources. Examples of this type of spectral leakage can be observed emanating
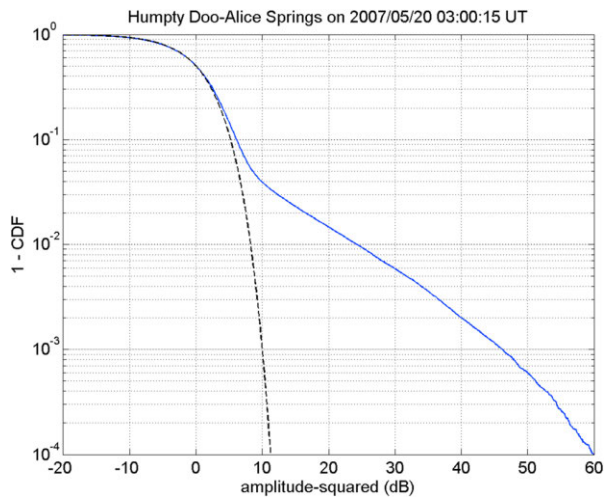
Figure 2. Amplitude characterization of impulsive interference in the whitened time series, plotted (blue curve) as the complementary CDF of squared sample amplitudes. A reference theoretical distribution for Gaussian in-phase and quadrature samples is plotted as the black dashed curve. CDF = cumulative distribution function.

from the E layer trace (centered on 1,260-km delay) at approximately 2.9, 4.0, and 5.5 MHz in the top panel of Figure 1.

Unknown waveform transmission nulls may be detected by applying a median filter to the normalized samples and flagging anomalously small-amplitude samples (i.e., below a threshold of −8 dB) as bad data that require signal interpolation. Alternatively, for known transmitters, a fixed set of nulls may be specified. As the waveform may not be coherent on either side of the gaps, the null positions are dilated rather generously prior to the signal interpolation stage to allow for a smoother transition. The resulting null data map is combined with the bad data map for the signal interpolation stage.

A portion of a DORS data segment in which RFI has been detected is shown in Figure 5. The 625 samples correspond to 400 ms or 50 kHz of the sounder sweep (at 125 kHz/s), for frequencies just below 12 MHz. Only the real part of the complex data representation (without calibration) is plotted. In the shaded bad data regions, impulsive noise is clearly evident from the large-amplitude excursions in the raw data (red). Figure 6 shows the two-phase threshold detector used to identify these corrupted regions. Elsewhere in Figure 5, the raw data are identical to the interpolated data (blue) and are therefore hidden.

3.2. Signal Interpolation

The signal interpolation stage of the algorithm is designed to replace samples in the bad (or null) data map with samples calculated using a weighted combination of forward and backward AR model predictions. This AR interpolation process requires that a model of at least the strongest desired signals, propagated via one or more ionospheric modes from the transmitter, is obtained. To ensure capture of the desired signals, the model order used here ($p = 45$) is substantially higher than in the RFI detection stage ($p = 3$). For robust statistical calculation of the AR coefficients, the segment must contain sufficient interference-free samples, and the signal should be stationary (otherwise a time-varying AR model would be more appropriate). While the selected segment length is variable, according to the actual positions of the bad data, it typically spans at least 160 kHz of the sweep, corresponding to about 2,000 samples at the decimated sampling rate. The sample support is thus an order of magnitude greater than the AR model order, and these statistical requirements should be well met.

The higher-order AR(45) model calculation again uses Burg's technique on each time-stepped segment, surrounding an interval of bad data, while excluding all corrupted samples (as in Nuttall, 1976). This technique achieves generation of stable AR coefficients, which is desirable for the subsequent interpolation process. Note that the segments are not necessarily the same as those used in the RFI detection stage or later ionogram processing and may contain some overlapping good data at the ends. Using the AR coefficients as linear prediction coefficients, each bad data sample is replaced by a linear sum of its AR coefficient weighted neighbors. The linear prediction is performed in both the forward and backward directions from neighboring good samples, with the forward and backward predictions further weighted using a sigmoid function (as described in Turley, 2003). The described technique uses an iterative process that achieves a very high convergence rate, excellent interference rejection, and robustness to model parameterization order and large missing data gaps.

Figure 3. Temporal characterization of impulsive interference in the whitened time series, plotted as the complementary CDF of the lengths of contiguous samples (blocks) exceeding a 20-dB threshold. CDF = cumulative distribution function.

In the presence of a very congested band of multiple interferers (e.g., in the lower HF band at night), or a single wide-band source, the proportion of bad data samples may on occasion overwhelm the good data samples

AGU
100
ADVANCING EARTH
AND SPACE SCIENCE

Radio Science

10.1029/2018RS006683
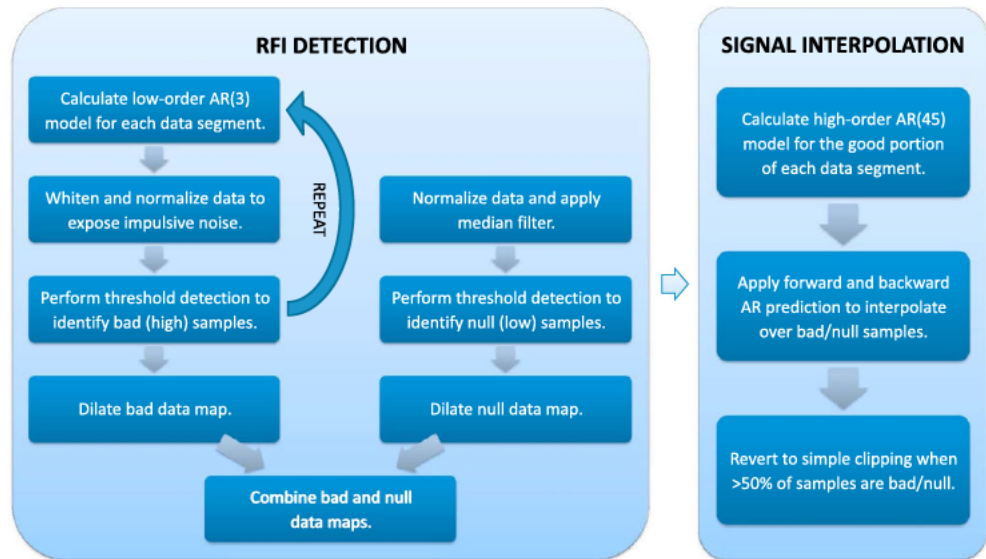
Figure 4. A block diagram showing the two main stages of the RFI rejection algorithm and their processing steps. Improved levels of rejection may be achieved with multiple passes through the entire algorithm. The results presented in this paper use three complete passes, except where otherwise indicated. RFI = radio frequency interference.

(e.g., with many 3-kHz RFI bands each corrupting ~2% of the 160-kHz segment). This can lead to a poor AR model representation and likely oversuppression of the signal, as the linear prediction decays exponentially over time according to the impulse response of the AR model. The approach taken for such a segment is to revert to a simple clipping algorithm, in which amplitudes are clipped at a fixed level above the median when more than 50% of the samples in a segment are bad. Although rudimentary, this still provides some level of RFI rejection, as evident in the bottom panel of Figure 1 (with clipping applied over the entire data series), and makes the overall algorithm more robust against a wide variety of RFI sources. In less extreme cases, there may be sufficient samples overall, but there exist short contiguous blocks of good data that are less than the AR model order. These short blocks are flagged as bad data prior to the model calculation.
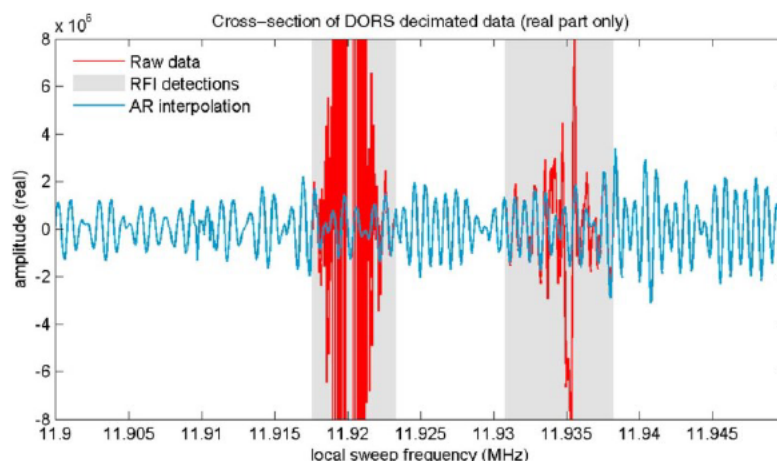
Figure 5. An example of a DORS data series, before (red) and after (blue) the first pass of the RFI rejection algorithm. In this case, impulsive RFI was detected in two sections of the raw data, across which AR interpolation was used to restore the underlying signal. DORS = Digital Oblique Receiving System; RFI = radio frequency interference; AR = autoregressive.
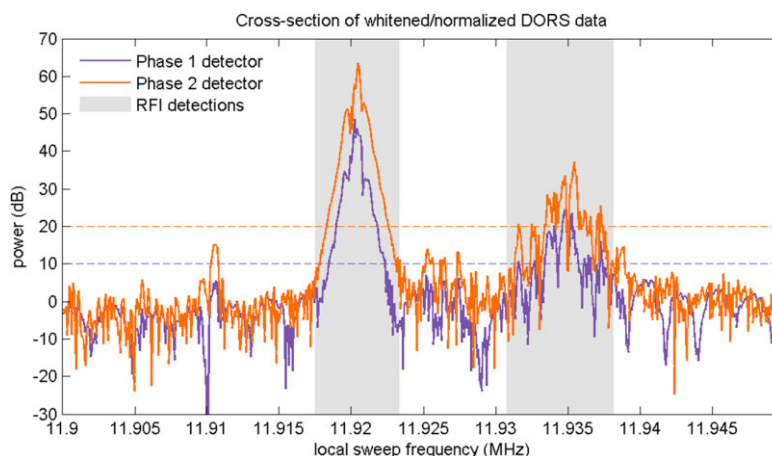
Figure 6. The same example as in Figure 5, after applying the whitening and normalization filter to detect impulsive noise within the first pass of the algorithm. The RFI detection process is carried out over two sequential phases (blue and orange data points), with different predefined thresholds for each (10 and 20 dB, respectively, as indicated by the dashed horizontal lines). The shaded bad data regions are shown after dilation. DORS = Digital Oblique Receiving System; RFI = radio frequency interference.

The blue line in Figure 5 shows the output of the AR interpolation process. Again, note that this curve is identical to the original raw data (red line) outside the bad data regions (shaded), whereas inside these regions an approximation of the underlying signal has been restored. The reduction in signal amplitude is negligible for short interpolation spans, such as in this example, noting the assumption of stationarity. The problem of oversuppression, described above, only becomes evident for much longer sections of corrupted data. In general, for the best spectral integrity, it is more important to preserve the coherency of the multiharmonic signal across the segment than to strictly maintain constant amplitude.

Note also that the number of bad samples in this example is considerably less than 50% of the total number of samples, and in most cases this remains true over the coherent processing bandwidth of the ionogram. Variance in the estimation of the AR coefficients is reduced with increased sample support size (Porat & Friedlander, 1987); here we aim to have the number of independent uncorrupted samples at least twice the sample covariance matrix dimensionality (in this case, the AR model order).

In more marginal cases of impulsive RFI, additional passes of both the RFI detection and signal interpolation stages may help to further improve the data. The standard mode used on the DORS sounder is a three-pass approach, and this is captured in the performance statistics presented in section 4. With the algorithm running as a MEX file (compiled C code) in MATLAB®, the three passes are completed within a few seconds on a typical desktop PC. This allows the RFI rejection algorithm to be executed routinely in real time, as part of the DORS onboard sounder processing.

Applying conventional ionogram processing to such "cleaned" data (after three complete passes) gives the middle panel in Figure 1. Compared to the raw ionogram (top panel), most of the vertical RFI stripes are now coherently rejected and the underlying trace is better defined, particularly in the weaker modes. Small residual spurs (i.e., the short vertical streaks emanating from parts of the ionogram trace) are still visible on some of the stronger modes (e.g., the F_2 low mode around 14–15 MHz and 1,400 km); this residual spectral leakage is an indication of imperfect signal reconstruction. As described in section 3.1, some spurs are caused by time gaps in the transmitted waveform; they may be observed in the original data and are not due to RFI events. In zones of very dense bands of interferers (e.g., just above 15 MHz), high rejection rates are achieved, but residual interference may remain. However, these artifacts are sufficiently low in power so as not to greatly affect ionogram analysis such as autoscaling. It is worth also noting that even distributed ionospheric mode features, such as the returns attributed to side scatter off the ground extending out in frequency and delay from the main two-hop trace at 1,700 km, are generally well preserved, despite potentially soaking up extra degrees of freedom in the AR model.
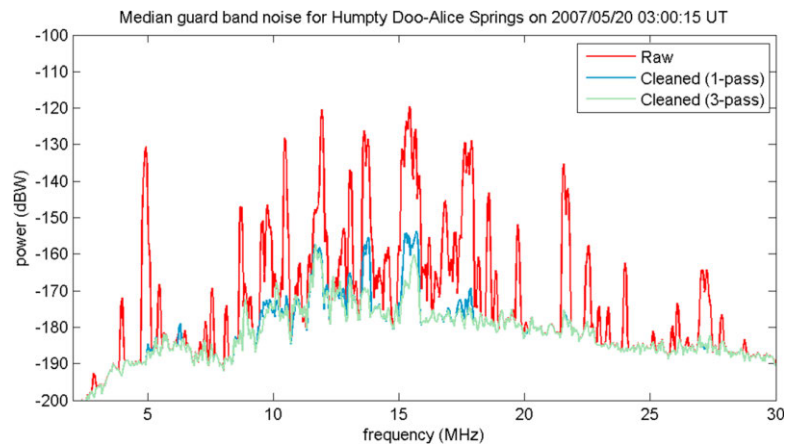
Figure 7. Median guard band noise in the raw and cleaned ionograms (i.e., for group delays less than the great circle ground range).

Figure 7 plots the median guard band noise in the raw (red) and cleaned (blue and green) ionograms of Figure 1. This is defined as the median power level at each frequency for group delays less than the great circle path from the transmitter (i.e., containing signal-free samples for interference and noise power estimation). Since the ionogram has been normalized for coherent narrow-band sounder signals, assumed to persist throughout the entire coherent processing bandwidth, the median guard band noise is an apparent noise level only in this domain. As such, it does not represent the actual peak power of any (impulsive) interferer, nor does it measure the power spectral density of truly incoherent background noise. The Figure 7 results show that AR interpolation using only a single pass (blue) commonly achieves RFI rejection ratios of 40–60 dB in the guard band. With additional iterations, the three-pass (green) results show further rejection of residual RFI. The performance statistics reported in section 4 will demonstrate that these levels of rejection are also achieved outside the guard band, in the vicinity of the ionogram trace itself.

Other interpolation schemes were also tested on the same data, including various forms of nonparametric FFT-based techniques, Papoulis-Gerchberg interpolation (Gerchberg, 1974; Papoulis, 1975), and a reduced implementation of Sacchi interpolation (Sacchi et al., 1998), but all were found to be less successful than the employed AR interpolation scheme.

4. Performance Evaluation

The goal of any ionogram RFI mitigation strategy, operating on either time or frequency domain representations of the data, is to recover the spectral properties of the signal of interest and background noise in the ionogram (frequency) domain. As such, performance must be evaluated on the final ionogram image, integrating multiple good and bad channels, rather than across individual narrow-band channels. Perfect reconstruction of the signal over each corrupted channel is not necessary or indeed possible under most circumstances.

In order to assess the performance of the RFI rejection algorithm in the ionogram domain, three general requirements were defined:

1. Preserve the guard band noise floor while suppressing the RFI level.
2. Minimize ionospheric mode spreading in group delay (i.e., spectral leakage of strong signals) and, more broadly, avoid raising the power of individual ionogram pixels.
3. Restore weak ionospheric signals otherwise hidden by the presence of RFI.

For each requirement, the aim was to derive statistics that would permit quantitative statements to be made about the algorithm performance on real data. The statistics that are presented in this section, as probability density plots and in tabular form, are based on combinations of the following metrics:
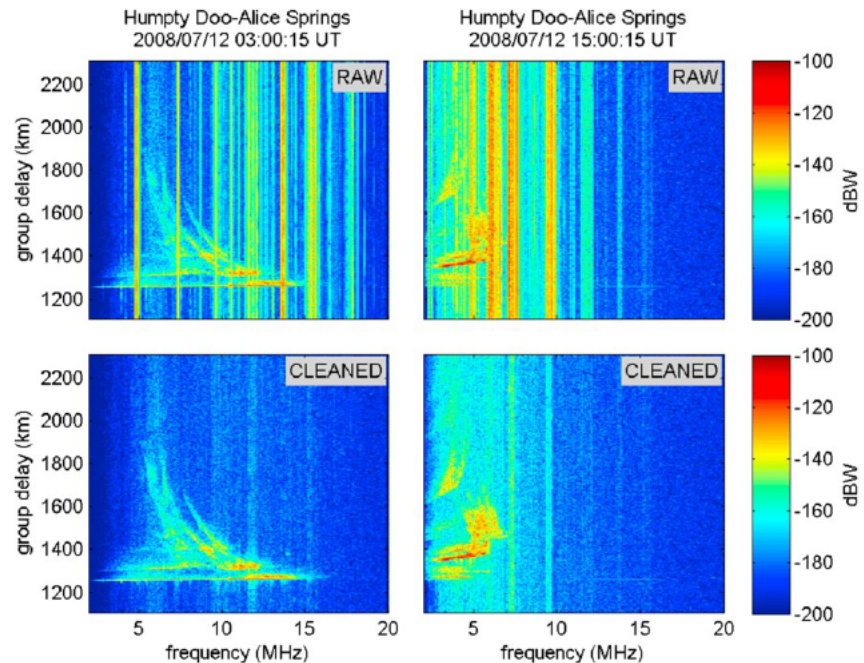
Figure 8. Two moderately disturbed ionograms from the sample data set, representing daytime (left column) and nighttime (right column) observations from 12 July 2008, when F_2 ionization was depressed due to a high-speed solar wind stream. Despite the high-frequency band being more congested and trace features more spread, the radio frequency interference rejection algorithm still performs remarkably well in the raw (top row) versus clean (bottom row) comparisons.

1. *RFI level.* The interference-to-background-noise ratio for a given frequency slice in the raw ionogram (prior to RFI rejection). It is calculated by subtracting an estimate of the background noise floor (being the smoothed lower decile of the median guard band noise, filtered over frequency) from the median guard band noise itself. Where RFI is not present, this will simply relate to the width of the background noise distribution. Note that congested bands of interferers may result in an overestimate of the background noise floor, and hence, an underestimate of the RFI level.
2. *RFI suppression.* The RFI difference ratio for a given frequency slice, calculated by subtracting the median guard band noise in the cleaned ionogram (after RFI rejection) from the corresponding value in the raw ionogram. It may be negative if the median noise level has been inadvertently raised by the RFI rejection process.
3. *Pixel cluster difference.* The power difference ratio for a given pixel cluster, consisting of five adjacent delay bins at a single frequency. It is calculated by subtracting the median power level in the cleaned ionogram cluster from the corresponding value in the raw ionogram cluster, for all nonoverlapping five-pixel clusters in the ionogram grid. Some of these clusters will contain signal power, while others will represent the background noise only, and any differences which are negative indicate raised power. Note that five samples are chosen to provide a trade-off between localization in delay and reduced sensitivity to outliers in the noise distribution.
4. *Weak signal restoration.* The weak signal difference ratio for a given frequency slice, calculated by subtracting the maximum power level in all group delays of the cleaned ionogram from the median guard band noise in the raw ionogram.

The above metrics were accumulated over a month-long set (July 2008) of ~6,000 DORS ionograms from the Humpty Doo to Alice Springs path. These ionograms cover a variety of propagation conditions, including both day and night. Two examples from one of the more disturbed periods are shown in Figure 8. The effect of variations in solar activity, although not presented here, appears to be less than the routine day/night variation, and while the exact percentages do change in proportion to the number of frequencies supporting
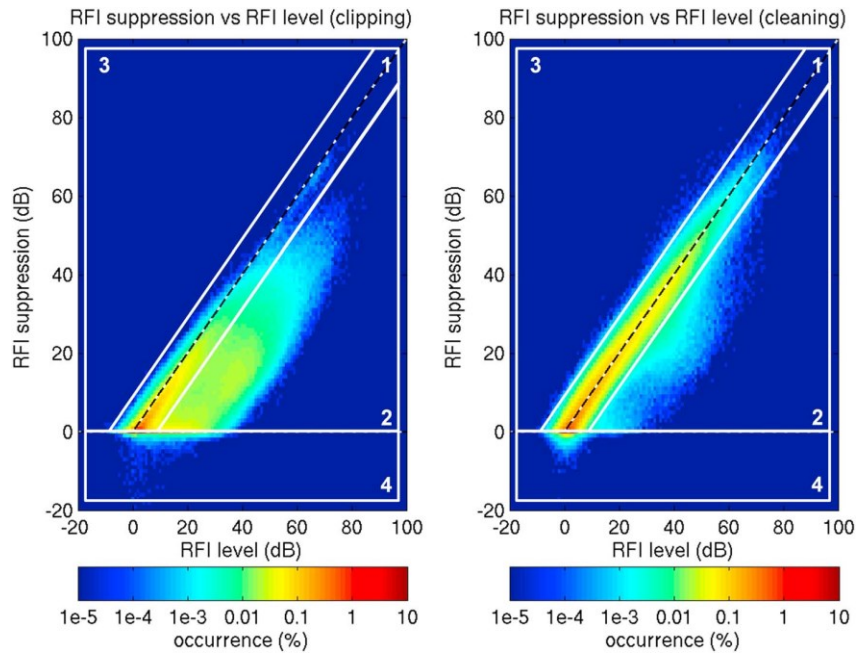
Figure 9. Probability density plots of RFI suppression versus RFI level for Humpty Doo to Alice Springs in July 2008. Four classes of results are identified: (1) complete suppression, (2) undersuppression, (3) oversuppression, and (4) raised noise. A single occurrence corresponds to a single frequency in one ionogram. RFI = radio frequency interference.

propagation, the relative differences between clipping and cleaning remain consistent. As many of the metrics are derived from the ~80 delay bins in the guard band, the results are not strongly dependent on the OI path, although they will be specific to the radio frequency environment at the receiver site. To eliminate the influence of the DORS receiver passband (5–35 MHz), including edge effects in the noise floor estimate, only frequencies between 7 and 33 MHz were considered.

In the following subsections, comparisons are made between cleaned ionograms, using the RFI rejection algorithm described in section 3, and clipped ionograms, using an amplitude-clipped version of the acquired data (with the amplitude limit set at 12 dB above the median-filtered data). The latter technique crudely reduces the RFI level in a manner similar to older-generation chirp sounders and therefore serves as a suitable baseline. An example of the relative improvement in ionogram feature clarity after cleaning is discussed in the final subsection.

4.1. Requirement A: Effect on Guard Band Noise

The performance against requirement A is visualized in the probability density plots of Figure 9, for both the clipping (left) and cleaning (right) algorithms. By conceptually dividing the 2-D space of *RFI suppression* versus *RFI level* into four regions, as shown, the relative occurrences for different classes of results can be compared. These classes are defined as follows:

Class 1. *Complete suppression.* Points lying on the 1-to-1 (dashed) line, plus or minus a tolerance factor of 9 dB, indicate successful RFI rejection as measured in the guard band.

Class 2. *Undersuppression.* Points lying to the right of the 1-to-1 line (beyond the 9 dB tolerance) indicate only partial RFI rejection.

Class 3. *Oversuppression.* Points lying to the left of the 1-to-1 line (beyond the 9 dB tolerance) indicate overly aggressive RFI rejection (i.e., reducing the apparent background noise to an artificial level).

Class 4. *Raised noise.* Points lying below the horizontal line where RFI suppression is zero indicate raised background noise levels (i.e., negative RFI suppression).

Table 1
Relative Occurrences for the Different Classes of RFI Rejection Results

Algorithm	Class 1 complete suppression	Class 2 undersuppression	Class 3 oversuppression	Class 4 raised noise
"Clipping" (baseline)	85.7% *(78.3% ≤ 10 dB, 7.5% > 10 dB)*	14.1%	<0.01%	0.18%
"Cleaning" (AR interpolation)	98.8% *(78.5% ≤ 10 dB, 20.3% > 10 dB)*	0.93%	0.05%	0.22%

Note. These occurrences relate to requirement A, using the Humpty Doo to Alice Springs (July 2008) data shown in Figure 9. Background colors indicate best (green), worst (red), and no significant difference (yellow). Class 1 occurrence statistics have also been further broken down into cases of weak (≤10 dB) and strong (>10 dB) RFI, listed in italics. RFI = radio frequency interference; AR = autoregressive.

Table 1 lists the class occurrence statistics corresponding to Figure 9. Note that the percentages take into account all frequencies (7–33 MHz) regardless of whether or not they actually contain RFI. Therefore, the results are diluted to some extent by the large density of points around the 0-dB origin that contribute to class 1. By separating class 1 into those cases with little or no RFI (i.e., RFI levels ≤10 dB) and those with stronger RFI (i.e., >10 dB), as specified in Table 1 (italics), the relative occurrences of the three undesirable classes (2, 3, and 4) are brought into more stark relief.

While the severity of each individual case is not captured here, the results do highlight some important results:

1. The cleaning algorithm consistently achieves much more complete suppression than the clipping algorithm, with undersuppression being characteristic of the clipping approach but comparatively rare for cleaning. Indeed, if only the stronger instances of RFI (>10 dB) are included in class 1, complete suppression becomes substantially less likely than undersuppression for clipping (7.5% vs. 14.1%), whereas the opposite is true for cleaning (20.3% vs. 0.93%). Suppression ratios of up to 40–60 dB are common in the cleaned results, whereas an upper performance of 20–40 dB is more typical of the clipped results (i.e., around 20 dB poorer).
2. Oversuppression is highly infrequent for both algorithms, with no significant difference in the occurrence statistics. It usually occurs in cases where there are congested bands of RFI corrupting the signal, leaving a section entirely unrecoverable. The 50% bad data cutoff, before the cleaning algorithm also reverts to clipping, is designed to minimize these relatively rare cases. The slight corresponding increase in undersuppression is generally an acceptable compromise.
3. Raised noise is also infrequent for both algorithms, with no significant difference in the occurrence statistics. For the cleaning algorithm, the issue tends to arise where RFI levels are low (<10 dB) and the data are probably better left untouched; however, the overall degradation is very small.

Thus, the "best" algorithm in terms of guard band noise behavior becomes a trade-off between the likelihoods of each of the four classes. Requirement A essentially translates to minimizing the occurrence of classes 3 and 4 (for which the guard band noise distribution is not preserved), while maximizing the occurrence of class 1 and, to a lesser extent, class 2 (for which RFI is suppressed to some extent). In classes 3 and 4, there are insignificant differences between the algorithms, and though the cleaning algorithm does perform marginally poorer than clipping, these rare cases are typically not severe. On the other hand, there are significant performance advantages achieved by the cleaning algorithm in classes 1 and 2. For class 1 in particular, the outcome is almost 99% of frequencies deemed to be more or less RFI-free after cleaning, compared with only 86% after clipping.

4.2. Requirement B: Effect on Strong Ionospheric Signals

Outside the guard band of the ionogram, the ability of the RFI rejection algorithm to leave strong desired ionospheric signals undistorted is determined by the quality of the interpolation. If the spectral properties of the signal are not preserved (including maintaining coherency across the interpolated interval), the ionospheric mode power may become spread in group delay; this produces the spurs (spectral leakage) previously noted in the middle panel of Figure 1. Ordinarily the spurs will be much weaker than the mode peak itself and the original RFI level. However, in the worst cases, the background noise in the vicinity of the mode may actually be raised, leading to a loss in local signal-to-noise ratio (SNR). Note that this kind of artifact would not necessarily have been detected as a class 4 ("raised noise") problem as defined above, because the spread is localized in delay and predominantly lies outside the guard band.
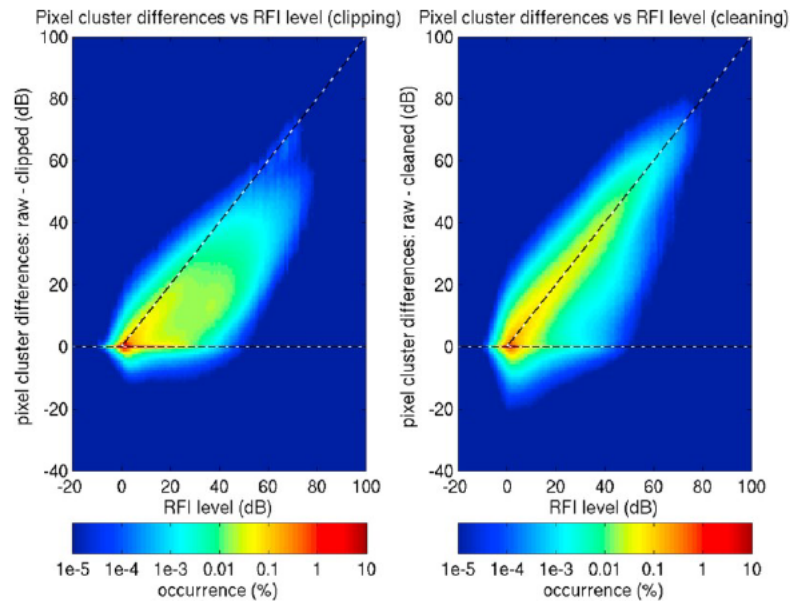
Figure 10. Probability density plots of pixel cluster differences versus RFI level for Humpty Doo to Alice Springs in July 2008. Unlike Figure 9 (above) and Figure 11 (below), these plots use an estimate from every 5-pixel cluster of each ionogram, rather than a single estimate per frequency. RFI = radio frequency interference.

A quantitative measure of success against requirement B is derived from the localized power differences between the raw ionograms and the clipped or cleaned ionograms. The distribution of these *pixel cluster differences* is plotted against *RFI level* in Figure 10, for both clipping (left) and cleaning (right). Negative differences (below the horizontal "zero difference" line) are generally a bad outcome, whereas positive differences tend to indicate successful RFI rejection, particularly near the 1-to-1 (diagonal) line. In many ways, the representation in Figure 10 is quite similar to that in Figure 9, although by considering all pixels, rather than just those within the guard band, it is also sensitive to effects that occur in the image region surrounding the desired ionogram trace.

The superior performance of the cleaning algorithm, particularly for high RFI levels, is again readily apparent in Figure 10. However, in comparison with the clipping algorithm, the cleaning algorithm does contribute to slightly more cases of raised power (i.e., negative differences below −0.5 dB). The total percentages of pixel clusters with raised power are listed in Table 2 for the two alternative algorithms; the results show marginal differences, being about 3% for both clipping and cleaning, but they are nevertheless more than the raised noise alone (addressed as part of Requirement A). If only raised powers above 10 dB in magnitude are included, these statistics drop to less than 0.002% (clipping) and 0.02% (cleaning), so the overall impact on ionogram SNR is not considered to be significant.

4.3. Requirement C: Ability to Restore Weak Ionospheric Signals

The final issue to address is requirement C: that is, the ability to restore low-level signals after clipping or cleaning is applied. To visualize this aspect of the algorithm performance, the distribution of *weak signal restoration* is plotted against *RFI level* in Figure 11, for both clipping (left) and cleaning (right). Only results for which the maximum power level after clipping or cleaning is at least 9 dB above the median guard band noise in the clipped or cleaned ionogram are included in Figure 11, as this is seen to be a good indicator of signals which are physical in origin (not just part of the background noise distribution). As such, the occurrence statistics inherently depend on propagation conditions; for instance, results are excluded when there is no underlying propagation at a given

Table 2
Relative Occurrences of Raised Power and Weak Signal Restoration

Algorithm	Raised power	Weak signal restoration
"Clipping" (baseline)	2.6%	6.2%
"Cleaning" (AR interpolation)	3.2%	6.9%

Note. These occurrences relate to requirement B (raised power) and requirement C (weak signal restoration), using the Humpty Doo to Alice Springs (July 2008) data shown in Figures 10 and 11, respectively. Background colors indicate best (green) and worst (red). AR = autoregressive.
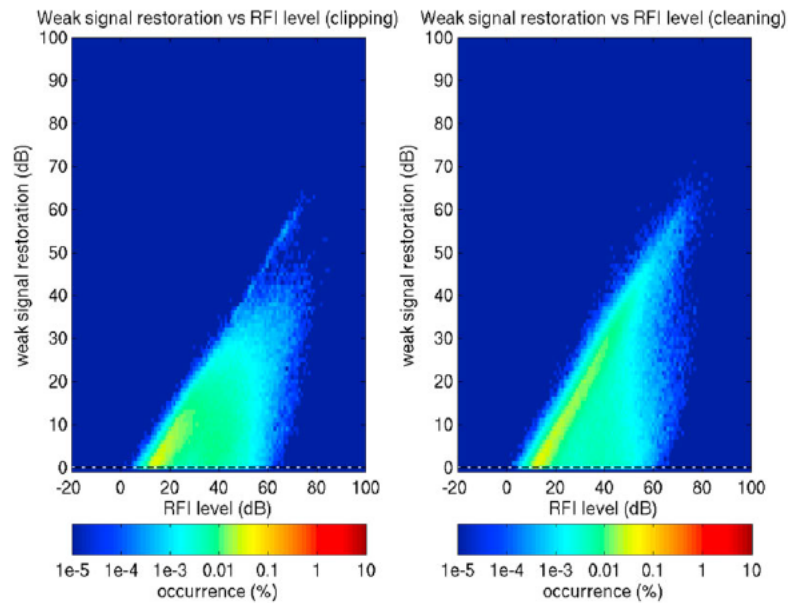
Figure 11. Probability density plots of weak signal restoration versus RFI level for Humpty Doo to Alice Springs in July 2008. A single occurrence corresponds to a single frequency in one ionogram. RFI = radio frequency interference.
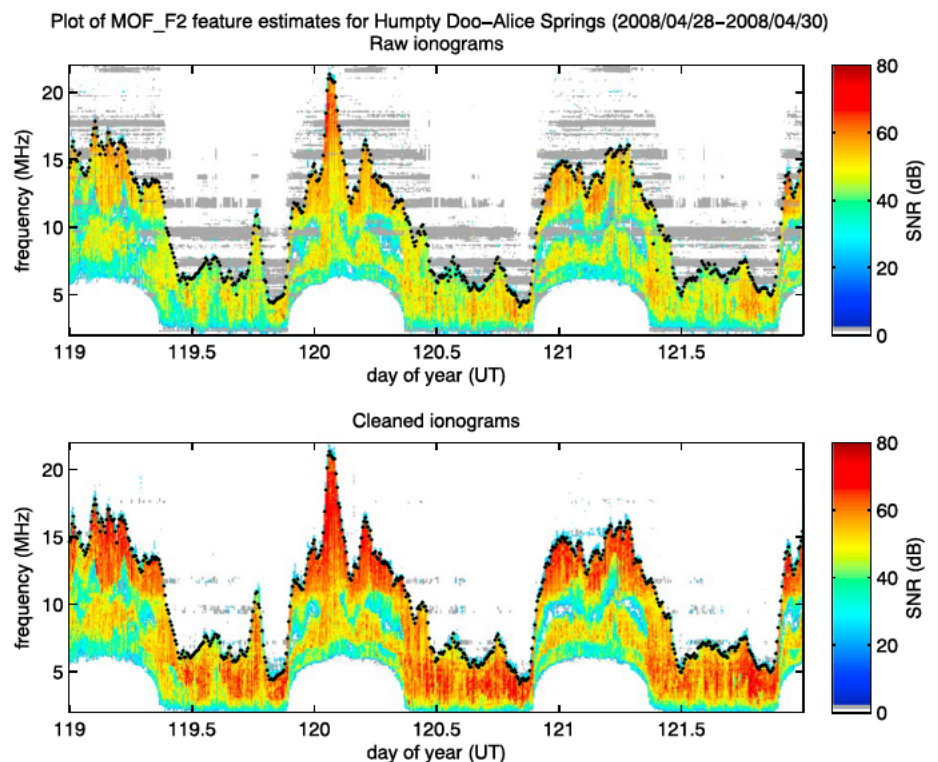
Figure 12. Estimates of the maximum observed frequency (MOF) of F_2 propagation for 3 days (28–30 April 2008) on the Humpty Doo to Alice Springs path. The edge feature estimates (black markers) are overlaid on an image of the time-varying band availability, defined as the maximum intensity (SNR) of ionospheric propagation at each frequency in the one-hop F_2 region of the ionogram. SNR = signal-to-noise ratio.
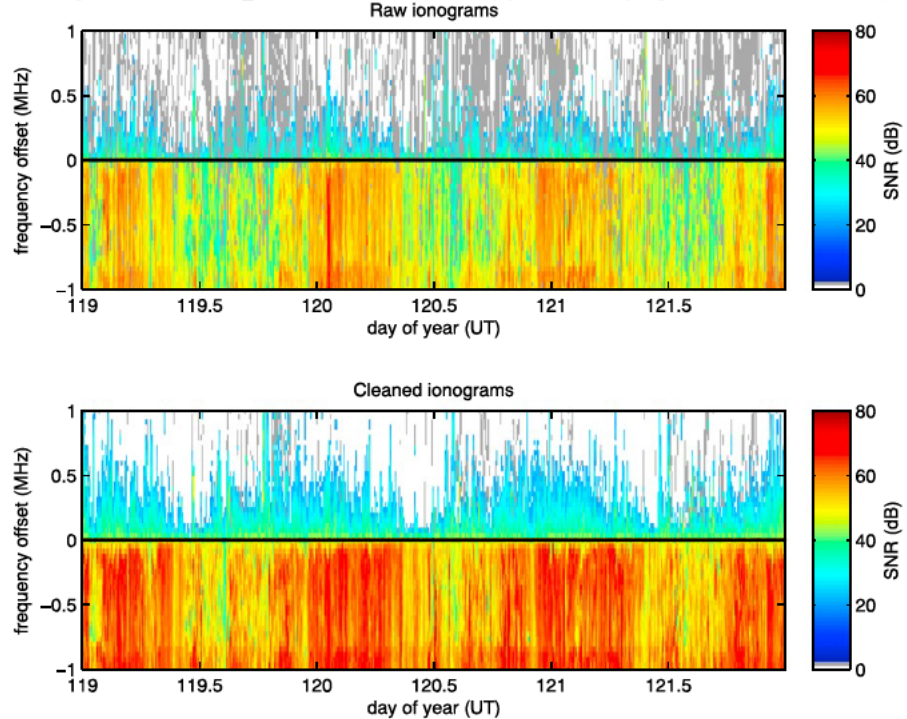
Figure 13. An alternative representation of the data in Figure 12, differenced with respect to the F_2 MOF feature estimate and presented on a reduced frequency scale. It shows the weaker off-angle and micromode structure around the MOF which is recovered by the cleaning process, as well as the robustness of the feature extraction algorithm in dealing with these weaker returns and unremoved RFI. MOF = maximum observed frequency; RFI = radio frequency interference; SNR = signal-to-noise ratio.

frequency (although the occurrences are still expressed as a percentage of all observations). Nevertheless, for the same set of soundings, a fair comparison can be made between the alternative algorithms.

The main conclusion to draw from Figure 11 is that cleaning tends to restore more weak signals overall, and in particular, is able to restore significantly more of the very weak signals (often sporadic E) that are more than 30 dB below the original RFI level (i.e., a greater density of points in the upper-right part of the distribution). Examples of such weak signals in Figure 1 include the patchy sporadic E trace between 16 and 20 MHz at 1,270-km delay, and the surface back-scatter return extending out in frequency and delay from the main two-hop trace at 11 MHz and 1,700 km.

The total percentages of frequencies with weak signals restored (>0.5 dB) are listed in Table 2, with both algorithms being around 6–7%. If only restoration factors above 10 dB are included, these statistics drop to 2.4% (clipping) and 3.8% (cleaning), thus widening the improvement offered by cleaning. Although far less common, the difference is even more pronounced above 30 dB: 0.12% (clipping) versus 0.69% (cleaning). In other words, cleaning is over 5 times more effective than clipping at recovering these very weak signals.

4.4. Improvement in Ionogram Feature Clarity

While the above requirements address the ability of the AR interpolation (cleaning) process to enhance the ionogram image integrity, it is worth emphasizing that the main beneficiary of these improvements is the DORS autoscaling algorithm (Heitmann & Gardiner-Garden, 2019), which relies on the clarity of a few key trace features. A practical test of the cleaning performance is therefore its impact on the feature extraction results. This is illustrated by way of example in Figure 12 and Figure 13, which are derived from a three-day sequence of DORS ionograms on the Humpty Doo to Alice Springs path (1234 km).

Figure 12 shows the F_2 maximum observed frequency (MOF) estimates (black markers) overlaid on an image of the maximum SNR at each ionogram frequency across all F_2 delay bins (≥ 160 km virtual height). The F_2 MOF is extracted as an edge feature, independently for each ionogram, and typically represents the extraordinary polarization mode. The top panel uses the raw ionograms (i.e., without cleaning), while the bottom panel uses the corresponding cleaned ionograms.

Although image thresholding is carried out in both cases to remove residual RFI in the amplitude-only ionogram, this has a tendency to leave gaps in the F_2 traces of the raw ionograms. Those pixels in gray highlight the data that have been removed by thresholding; the impact is clearly much more severe in the raw ionograms, with some frequency bands (e.g., 9–10 MHz) largely annihilated. Nevertheless, there are actually very few instances where missing data have noticeably degraded the MOF estimates, due to the ability of the autoscaling algorithm to effectively smooth over gaps (Heitmann & Gardiner-Garden, 2019). It is the improvement in overall image SNR associated with cleaning that has a more pronounced effect on the MOF estimates. For example, from around 12 UT (post-sunset) on day 120, the raw results contain more "jumps" where the estimate has overshot a weaker edge.

The same results are shown on an enlarged frequency scale in Figure 13, centered on the MOF estimates themselves. Ideally there should be a sharp edge in the image power at zero frequency offset. In most cases this is indeed observed, despite the weaker off-angle and micromode structure at higher frequencies, and the cleaning process greatly accentuates this edge in general. Other ionogram peak, edge and line features, although not shown here for sake of brevity, also exhibit comparable improvements in clarity with application of the cleaning process.

5. Conclusions

An algorithm for rejecting RFI in ionospheric sounder data has been described and demonstrated, using a collection of DORS OI soundings. It is based on applying impulsive noise detection and AR interpolation and operates directly on the acquired (deramped) data in the time domain. With conventional ionogram processing, the resultant cleaned ionograms exhibit RFI suppression ratios of 40–60 dB (or more) and are much better suited to autoscaling (ionogram feature extraction) than prior to RFI rejection. While some residual artifacts remain, most of the strong signals are preserved without distortion, and weaker signals are successfully restored, even with power typically distributed across multiple ionospheric propagation mode structures.

The AR interpolation (cleaning) algorithm consistently outperforms the baseline clipping algorithm for RFI rejection in a variety of real HF environments, as measured against the three general requirements and their corresponding metrics. The clarity of ionogram trace features after cleaning is almost always improved, as illustrated in the small sample of ionogram data and F_2 MOF estimates derived from the DORS autoscaling algorithm. A future extension of the performance analysis to operate on synthetically generated RFI may be useful for further validation.

Although soundings from only one DORS OI path have been evaluated here, the performance statistics are broadly representative of what has been seen on numerous other midlatitude paths with similarly quiet receiver sites. Indeed, the algorithm has now been running routinely as part of the DORS and PRIME onboard processing since 2012 and 2014, respectively, covering a wide variety of paths (both OI and VI) and ionospheric conditions. On this basis, the technique is considered to be an effective and robust RFI mitigation strategy in general. While the performance is yet to be evaluated for low- and high-latitude soundings, which are prone to more routine spreading in the ionospheric propagation modes, it is expected that the AR model order would remain sufficient, based on the disturbed midlatitude examples shown.

References

Abramovich, Y. I., Bohme, J. F., Gorokhov, A. Y., & Maiwald, D. (1993). GEM-algorithm for sea echo Doppler analysis with data corrupted by sparkling interference. *Proceedings of SPIE, 2027,* 194–203. https://doi.org/10.1117/12.160434
Andersen, N. (1974). On the calculation of filter coefficients for maximum entropy spectral analysis. *Geophysics, 39*(1), 69–72. https://doi.org/10.1190/1.1440413
Ayliffe, J. K., Durbridge, L. J., Frazer, G. J., Gardiner-Garden, R. S., Heitmann, A. J., Praschifka, J., et al. (2019). The DST Group high-fidelity, multi-channel oblique incidence ionosonde. *Radio Science, 54.* https://doi.org/10.1029/2018RS006681

Burg, J. P. (1967). Maximum entropy spectral analysis. Proceedings of the 37th Annual International Meeting of the Society of Exploration Geophysicists. Reprinted in D. G. Childers (1978). In *Modern spectrum analysis* (pp. 34–41). Washington, DC: IEEE Computer Society Press.

Burg, J. P. (1975). Maximum entropy spectral analysis (PhD thesis). Stanford, CA: Stanford University.

Fabrizio, G. A. (2013). *High frequency over-the-horizon radar*. New York, NY: McGraw-Hill.

Gerchberg, R. W. (1974). Super-resolution through error energy reduction. *Optica Acta: International Journal of Optics*, *21*(9), 709–720. https://doi.org/10.1080/713818946

Harris, F. (1978). On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, *66*(1), 51–83. https://doi.org/10.1109/PROC.1978.10837

Harris, T. J., Quinn, A. D., & Pederick, L. H. (2016). The DST group ionospheric sounder replacement for JORN. *Radio Science*, *51*, 563–572. https://doi.org/10.1002/2015RS005881

Heitmann, A. J., & Gardiner-Garden, R. S. (2019). A robust feature extraction and parameterized fitting algorithm for bottom-side oblique and vertical incidence ionograms. *Radio Science*, *54*. https://doi.org/10.1029/2018RS006682

Makhoul, J. (1975). Linear prediction: A tutorial review. *Proceedings of the IEEE*, *63*(4), 561–580. https://doi.org/10.1109/PROC.1975.9792

Nuttall, A. H. (1976). *Spectral analysis of a univariate process with bad data points, via maximum entropy and linear predictive techniques, NUSC Technical Report 5303*, New London, CT: Naval Underwater Systems Center.

Papoulis, A. (1975). A new algorithm in spectral analysis and band-limited extrapolation. *IEEE Transactions on Circuits and Systems*, *22*(9), 735–742. https://doi.org/10.1109/TCS.1975.1084118

Porat, B., & Friedlander, B. (1987). The exact Cramer-Rao bound for Gaussian autoregressive processes. *IEEE Transactions on Aerospace and Electronic Systems*, *AES-23*(4), 537–542. https://doi.org/10.1109/TAES.1987.310887

Sacchi, M. D., Ulrych, T. J., & Walker, C. J. (1998). Interpolation and extrapolation using a high-resolution discrete Fourier transform. *IEEE Transactions on Signal Processing*, *46*(1), 31–38. https://doi.org/10.1109/78.651165

Swingler, D. N., & Walker, R. S. (1989). Line-array beamforming using linear prediction for aperture interpolation and extrapolation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, *37*(1), 16–30. https://doi.org/10.1109/29.17497

Turley, M. (2003). Impulsive noise rejection in HF radar using a linear prediction technique. Proceedings of the IEEE International Radar Conference 2003 (pp. 358–362). https://doi.org/10.1109/RADAR.2003.1278767

Turley, M. D. E., & Netherway, D. J. (1990). OTHR signal reconstruction for data corrupted by impulsive noise. Proceedings of First Australian Radar Conference, Radarcon 90 (pp. 203–209).

Vaseghi, S. V. (2008). *Advanced digital signal processing and noise reduction* (4th ed.). Hoboken, NJ: John Wiley. https://doi.org/10.1002/9780470740156

F.5   Gardiner-Garden et al., *Advances in Space Research*, 2019

The following is a reprint of Gardiner-Garden, R., Cervera, M., Debnam, R., Harris, T., Heitmann, A., Holdsworth, D., et al. (2019), 'A description of the Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE)', *Advances in Space Research*, 64(10), pp. 1,887–1,914.

This paper summarises the ELOISE 2015 experiment as a whole, of which the AoA ionosonde was but one important part. It thus provides contextual information on the purpose of the wider experiment and some of its preliminary outcomes, which are related to but mostly outside the scope of this PhD project.

The manuscript was written by Dr Robert Gardiner-Garden, with myself as a co-author, mainly contributing to the description and figures for the AoA and CSF capabilities. The paper was published during my PhD candidature.

A description of the Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE)

R. Gardiner-Garden [a], M. Cervera [b,c], R. Debnam [b], T. Harris [b,c], A. Heitmann [b,c],
D. Holdsworth [b,c], D. Netherway [b], B. Northey [b], L. Pederick [b,c], J. Praschifka [b], A. Quinn [b],
M. Turley [b], A. Unewisse [b], B. Ward [b,c], G. Warne [b]

[a] National Security Intelligence, Surveillance and Reconnaissance Division, DST Group, Locked Bag 7005, Liverpool, NSW 1871, Australia
[b] National Security Intelligence, Surveillance and Reconnaissance Division, DST Group, PO Box 1500, Edinburgh, SA 5111, Australia
[c] School of Physical Sciences, The University of Adelaide, Adelaide, SA 5005, Australia

Abstract

The Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE) was an extensive experiment undertaken by the Defence Science and Technology (DST) Group that focused on collecting ionospheric sensor data from multiple overlapping ionospheric paths in the Australian region in order to improve understanding of the characteristics of ionospheric variability and its effect on HF radio propagation. The experiment ran from July to October 2015 and included a period of three weeks of increased sample density and three days of dedicated over-the-horizon (OTH) radar operations. It was anticipated that ELOISE would sample a wide range of environmental conditions and present an opportunity to characterise periods of "normal variability" and periods of "exceptional variability" in the ionosphere.

This report is a general description of the aims of this experiment and the types of data collected. Particular interest focused on observing and measuring variability in ionospheric electron density gradients and their effect on oblique HF propagation. To this end, ELOISE established a pair of two dimensional HF receiver arrays to directly measure the oblique angle of arrival (AoA) on many overlapping oblique paths. ELOISE also established a dense sub-network of spatially separated quasi-vertical incidence soundings in the vicinity of Alice Springs in central Australia. This enabled a comparison of gradients observed in a dense network of vertical sounders with gradient effects observed in oblique propagation passing overhead. Several additional ionospheric observing systems were also used to give complementary pictures of the fine scale characteristics of ionospheric variability in the region.

Plain Language Summary: This paper is an overview of the 2015 Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE), an experiment designed to observe and characterise mid-latitude ionospheric disturbances in the Australian region and understand their impact on high frequency (HF) signal propagation.

ELOISE involved the simultaneous operation of a large collection of ionospheric sounders enabling ionospheric variability to be characterised on a finely sampled large scale.

Particular efforts were made to provide direct high fidelity measurements of the angle of arrival (AoA) on many oblique HF propagation paths. These direct AoA measurements imply horizontal electron density gradients that can be compared to ionospheric gradients estimated from conventional models of the ionosphere derived from the spatial network of sounder sites.

The size and scope of the experiment are detailed in this paper and some preliminary results are presented.

1. Introduction

Long distance HF radio signal propagation is affected by the natural variability of the ionosphere (Davies, 1990; McNamara, 1991; Fabrizio, 2013). HF users mitigate the effects of these environmental variations on their systems by understanding the variability and modelling its impact on HF propagation. One of those HF users is the Australian Defence Force who operates the Jindalee Operational Radar Network (JORN), a network of over-the-horizon (OTH) surveillance radars. This system is supported by a network of HF vertical, oblique and backscatter sounders all observing the ionosphere in real time. While many aspects of ionospheric variability are well understood, it is desirable to increase the reliability and precision of our understanding and our models in order to improve the performance of HF systems.

Many studies of ionospheric variability concentrate on the large scale space and/or time variability with limited spatial and/or temporal resolution (Ayub et al., 2009; Danilov, 2009; McNamara et al., 2008; McNamara and Wilkinson, 2009; Rishbeth et al., 2009; Rishbeth et al., 2000; Sojka et al., 2007; Sojka et al., 2001). These studies focus on climatologically averaged changes, hourly time scales, and patterns of variation resolved on the scales of 500 km or 5° of latitude/longitude. Conversely, other ionospheric variability studies focus on fine temporal and/or small spatial scale variability without characterizing the large scale spatially and temporally varying environment (Harris et al., 2012; Hernandez-Pajares et al., 2012; Paznukhov et al., 2012; Valladares and Hei, 2012). These studies focus on a few days of data, time scales of minutes to hours, and spatial patterns of variation resolved on the scales of tens of kilometres but spanning at most a few degrees of latitude/longitude.

To further our understanding of the ionosphere and its variability, the simultaneous observation of variations in ionospheric parameters across a greater range of scales is necessary. Of particular interest is this question: "how well is the ionospheric variability on different scales being represented by the ionospheric sounder observations and our models?" Ionospheric gradients also affect the HF propagation path though the ionosphere (Davies, 1990; McNamara, 1991). Small spatial scale variations potentially have a disproportionate effect on the steepness of gradients and hence the sounder observations. Of particular interest for this study are the questions: "how well are the observed ionospheric angle of arrival (AoA) variations being represented by our models?" and: "is there significant ionospheric variability that is under-sampled by our sounder networks?"

The Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE) was an experiment conducted by the Defence Science and Technology (DST) Group which was designed to address these questions by observing ionospheric variability in the Australian region at both large and small scales simultaneously and with multiple different instruments. This dataset has subsequently been used to test models of ionospheric variability and understand the different instrument responses to variability on both small and large scales. Particular interest was focused on the measurement of gradients and the spatial and the temporal correlations in the patterns of variability.

ELOISE was active for approximately 84 days from July to October 2015. The ionospheric variability observed during this period included several ionospheric storms; days with substantial mid-day collapses of the F_2 critical frequency, foF2 (Lynn et al., 2014, 2016); many days with periods of general geomagnetic disturbance and large scale and medium scale travelling ionospheric disturbance (TID) activity; and several nights with distinct mid-latitude spread-F.

ELOISE was an experiment with many related components:

- A core component of the experiment was an existing network of high quality ionospheric sounders, which included vertical incidence sounders (VIS) and oblique incidence sounders (OIS), routinely making measurements over north and west Australia. This network was modified and expanded during ELOISE to increase its spatial and temporal resolution.
- A pair of two dimensional antenna arrays was constructed to estimate the oblique propagation angle of arrival (AoA) from 11 OIS paths simultaneously. This permitted the observation of the effect of variability in ionospheric gradients on HF propagation path delay, azimuth and elevation. This builds on the work of Black et al. (1995) and extends it with modern equipment to many paths.
- Twelve OIS paths (some in common with the AoA paths) were modified so that they collected both conventional OIS ionograms and simultaneously made channel scattering function (CSF) measurements of the Doppler shift resulting from ionospheric motion (iDOP). These measurements allow an understanding of the pattern of iDOP variations induced by ionospheric variability on different scales. This builds on the work of Lynn (2008, 2009).
- A small, dense and rapidly sampled network of vertical and quasi-vertical incidence sounders (VIS/QVIS) was established in the Alice Springs region, under the apogee of two separate overlapping OIS propagation paths. This builds on the work of Harris et al. (2012) by simultaneously understanding the spatial correlations of disturbances on small and large scales.
- A JORN radar was used to collect tracking data derived from an expanded network of eight transponders in the region of interest and many commercial air targets of opportunity. Co-located JORN backscatter sounders (BSS) (see Earl and Ward, 1987) were also used to collect HF propagated and backscattered surface returns in the region of interest. In addition, an experimental radar task was also designed to collect data showing

the leading edge of the backscattered surface clutter, over a limited set of frequencies and all azimuths, but with much greater azimuth resolution than the conventional BSS.

Auxiliary experiments utilizing airglow cameras and collaborative experiments with the Australian Bureau of Meteorology Space Weather Services (BOM/SWS) using dual frequency GPS Total Electron Content (TEC) receivers were also performed as part of ELOISE but are not reported here. Section 2 of the paper describes the design goals of this experiment, including how the experiment was designed to generate numerous opportunities to compare different and complementary ionospheric sounders observing the same instance of ionospheric variability. In Section 3 the different data types and processing steps are described. Section 4 provides preliminary analysis and discussion, with some initial conclusions on ELOISE made in Section 5.

2. The experiment design

ELOISE was an experiment designed to improve our understanding of ionospheric variability in a large region but on a fine scale. It ran for approximately 84 days from July to October 2015. The equinox period of the trial was selected because sporadic E (Es) is at a seasonal minimum (in our experience of observations). The intent of the trial was to focus on the non-Es related causes and patterns of ionospheric variability.

2.1. Understanding the large scale pattern of ionospheric variability

The existing Jindalee Operational Radar Network currently monitors ionospheric variability via a large combined network of twelve VIS, six OIS transmitters, three OIS receivers and four BSS systems. The VIS network has a typical separation of 7° of latitude and longitude and is sampled every 3.75 min, while the OIS network prior to ELOISE was sampled every 7.5 min. To improve the spatial and temporal density of samples collected during ELOISE, the OIS network was upgraded and expanded with three additional OIS receiving systems and three additional OIS transmitting systems so that the typical spatial separation between observations over the area of interest was reduced to 3° of latitude and longitude and the default sample revisit time was decreased to 3.75 min across the combined VIS and OIS networks. One of these OIS receiving systems was new hardware co-located with existing equipment for cross validation.

The purpose of this expansion was to improve the spatial and temporal sample density of observations of ionospheric variability and produce a better measure of ionospheric density gradients. Each ionogram in this integrated set of sounder measurements can be automatically fitted in order to contribute to a model of the ionospheric

electron density $N_e(x,y,z,t)$ and its spatial and temporal gradients. A map of the expanded VIS and OIS networks is presented in Fig. 1a. A map of the nominal BSS spatial coverage is presented in Fig. 1b. The exact equipment locations are catalogued in Appendix A.

2.2. Measuring the effects of ionospheric gradients on HF propagation angle-of-arrival

ELOISE was designed to measure the oblique propagation phase front angle of arrival (AoA) from a wide range of OIS paths in order to estimate the effect of ionospheric gradients on HF propagation in the region. This was achieved by the construction of two separate phased antenna arrays, each being two dimensional HF "receiver per element" systems. These arrays permitted the routine observation of variability in the oblique propagation AoA from 11 OIS paths simultaneously. This propagation path AoA variability is expected to be related to variability in the underlying ionospheric electron density gradients because gradients in the ionosphere's electron density are controlling the propagation path refractive index and are likely to produce path refraction asymmetries around the path midpoint. The aperture of the arrays (90 × 180 m) was large enough to enable good precision in the estimation of the AoA over most paths. The design, array calibration and signal processing is discussed in a separate paper (Heitmann et al., 2018) but in summary, each array consisted of a crossed pair of uniform linear arms with a total of 19 antennas (10 per arm) to measure bearing (or azimuth) and elevation of arrival for each OIS transmission.

Fig. 2 shows the basic layout of antennas within the Laverton AoA array. The 6.5 m high monopoles have a broad HF response from 5 to 30 MHz.

The AoA estimation process is a least squares planar fit to the phase front across both arms. This is calculated for each pixel in the OIS images (all frequencies and delays). An assumption of simple mirror geometry (i.e., a one-hop propagation path for each delay) provides an approximate reference azimuth and elevation for each ionogram cell which are used to resolve grating-lobe ambiguities. The slope of the fitted phase front hence gives a reliable and unambiguous estimate of each component of angle of arrival (azimuth and elevation). Much of the possible propagation path multimode is well resolved in both frequency and delay within this high resolution one-way path data. The residual mean square error of the fit is used to estimate the measurement uncertainty and define the limits of this planar fit.

A map of the OIS paths from Fig. 1a that observe AoA is presented in Fig. 3a. One installation at Laverton, Western Australia used mains power and operated 24 × 7 observing eight OIS paths from late July 2015 to mid-January 2016. The second installation at Coondambo, South Australia used generator power and only operated when staffed; i.e., three paths and 8–16 h per day for the trial period. These locations were selected because they presented a broad range of OIS path lengths and because they
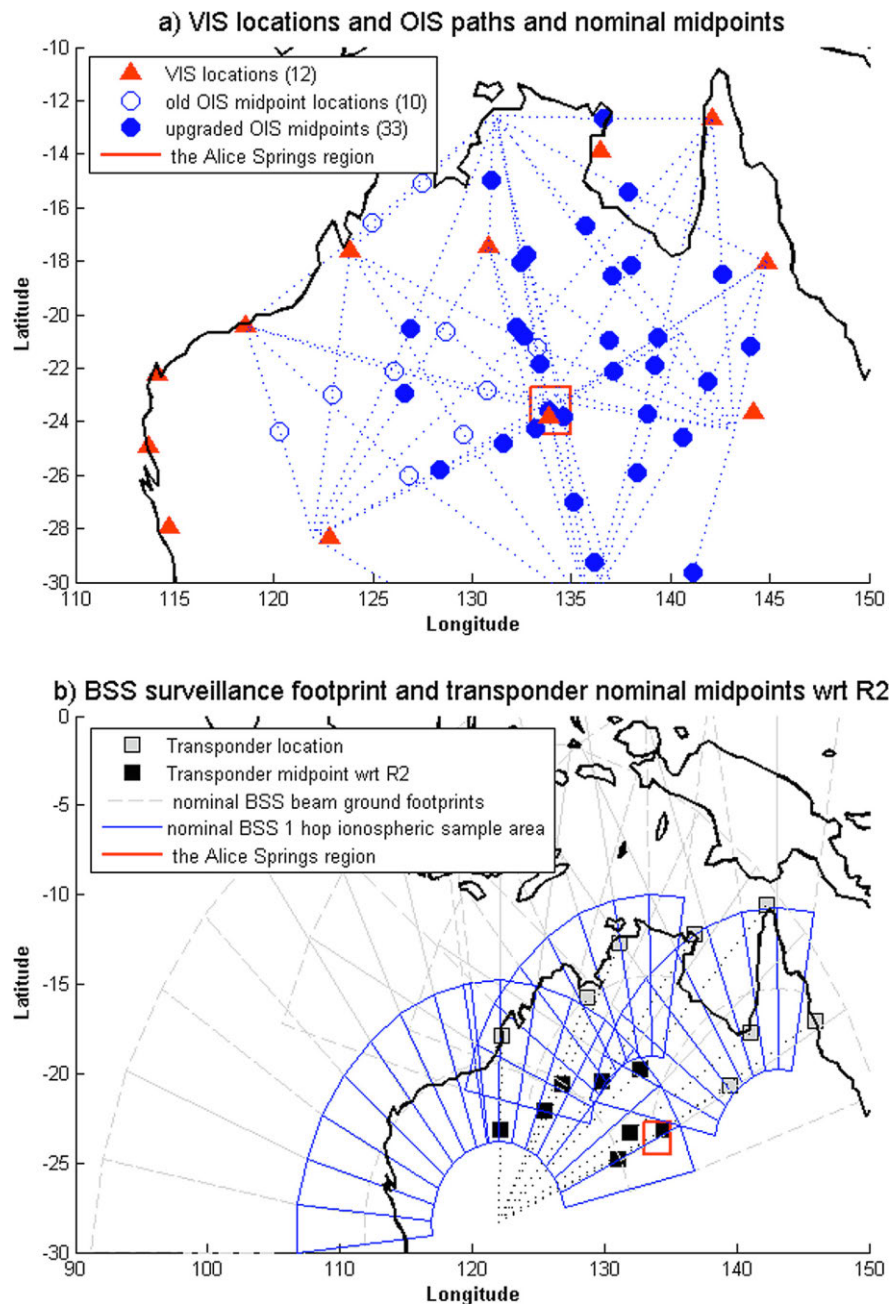
Fig. 1. (a) A map of the nominal sounder ionospheric sample points; i.e., the VIS locations and OIS path midpoints. (b) A map of the expanded transponder network and the JORN BSS surveillance footprints and the nominal ionospheric area influencing the measurements; i.e., the transponder path midpoints with respect to JORN radar No. 2 near Laverton, Western Australia.

crossed over a single common region in the area of Alice Springs in the Northern Territory (red square) where there was already a VIS site, allowing the impact of ionospheric variability and disturbances on very different propagation geometries to be observed.

2.3. Observing fine scale disturbances and gradients

A key goal of ELOISE was to understand how fine scale disturbances and gradients will affect observations made by different classes of sounding instruments. Ionospheric variability measured by probing steeply up into the ionosphere with a VIS may potentially have different characteristics to variability observed with long, low-elevation OIS paths, propagating through nearly the same region of the ionosphere. To understand the effect of small scale ionospheric variations on different sounding instruments, ELOISE constructed a small quasi-vertical incidence sounder (QVIS) network (two QVIS receivers and one QVIS transmitter in addition to the existing VIS transmitter/receiver) in the
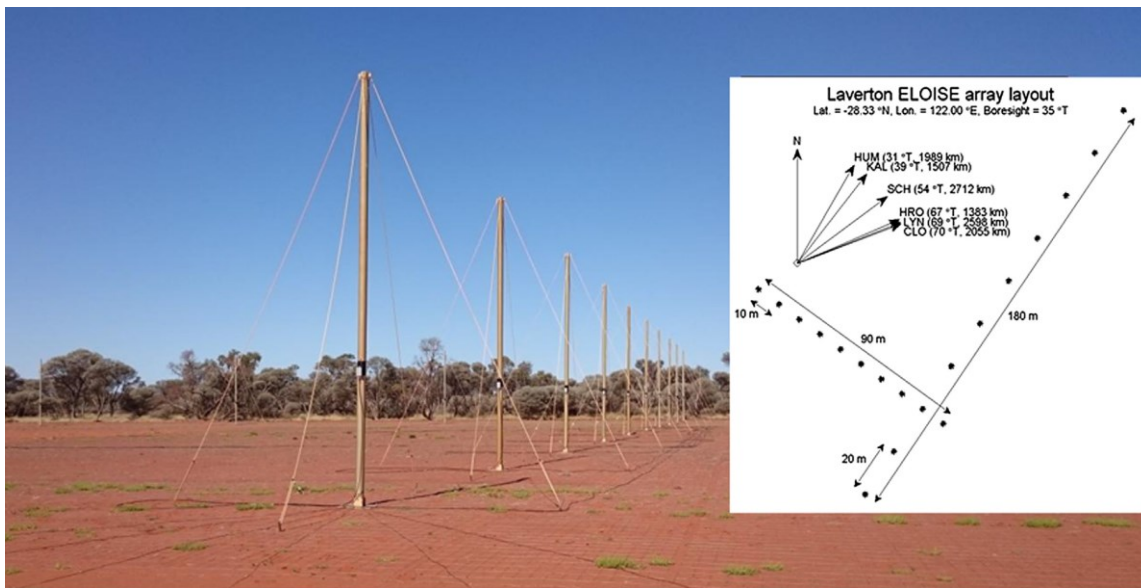
Fig. 2. A schematic of the layout of the Laverton 2D receive array configuration and a picture of the antennas.
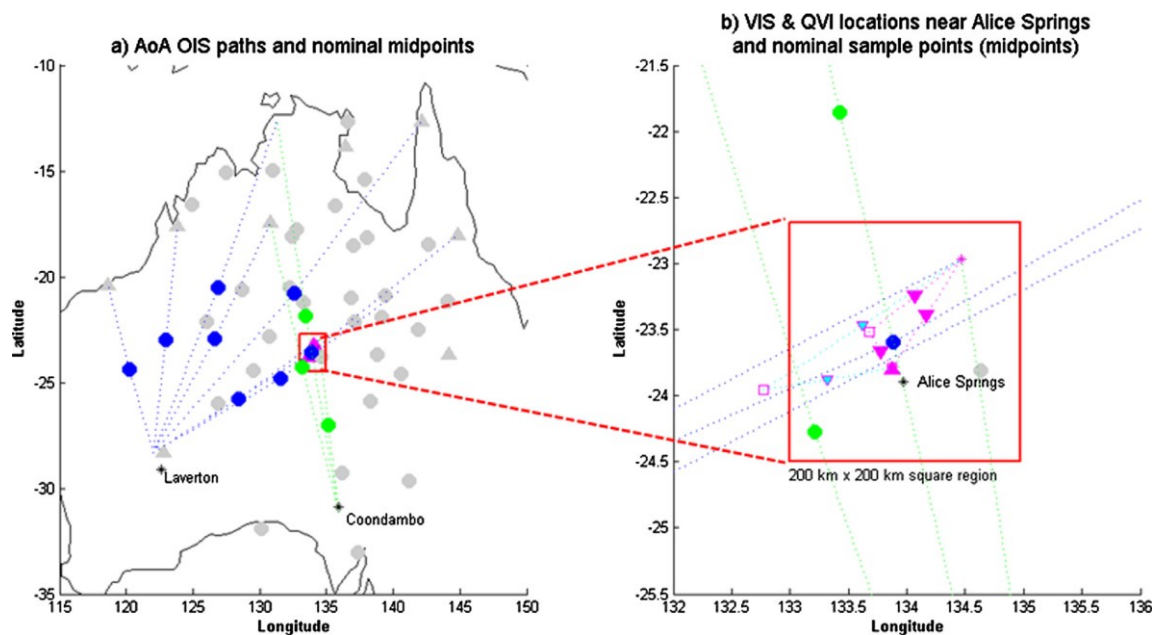
Fig. 3. (a) A map of the ELOISE OIS AoA paths (dashed lines) and midpoint positions (circles) from Fig. 1a with respect to Laverton (blue) and Coondambo (green) and the focus region near Alice Springs (red square). (b) A detailed map of the start(plus), end(square) and midpoint positions (magenta triangles) of the small scale QVIS network near Alice Springs. The long OIS path midpoints (from part a) in the region are also included. The unfilled QVIS midpoints (blue triangles) suffered from low data quality. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Alice Springs region. This equipment and the existing VIS were operated with a rapid 75 s revisit rate. A spatial map of the VIS site and five QVIS midpoints (marked as magenta triangles) is presented in Fig. 3b. The locations of QVIS transmitters (plus symbol) and QVIS receivers (magenta squares) are also marked. Regrettably, one of the QVIS receiver sites (Hermannsburg) had very high local noise and produced very little high quality data which

is represented by the unfilled midpoint symbols in Fig. 3b. This still left three good QVIS sample points and one good VIS in the Alice Springs network (marked as solid magenta triangles in Fig. 3b).

Additionally, a dual filter all-sky high fidelity air glow camera was co-located with the dense VIS/QVIS network. This dataset presents an opportunity to observe disturbances on a very small scale and relate the pattern of vari-

ability observed in the airglow to that observed in the dense sounder network. Results and descriptions of this part of ELOISE have been presented separately (Pederick et al., 2017; Unewisse et al., 2017; Unewisse and Cool, 2017a).

2.4. Comparing sounder and OTH radar observations of HF propagation

Another key goal of ELOISE was to relate the sounder observations of ionospheric variability to the JORN OTH radar observations of Earth surface backscatter power and associated transponder tracks (i.e. sequences of associated radar peaks in azimuth/range/Doppler space). A map of the BSS network and the expanded transponder network is presented in Fig. 1b. Two of the JORN radars (No's 1 & 3) and their BSS subsystems have an arc of coverage of approximately 90°. The third JORN radar (No. 2), in the West at Laverton, is different and has two separate segments each with an arc of coverage of 90°. HF radio transponders require illumination (and reception) from one of the JORN OTH radars (Fabrizio, 2013), in order to obtain HF propagation path information (delay, azimuth and ionospheric Doppler) for a particular two way pair of nearby point to point propagation paths at a particular HF frequency (one path from the transmitter array to the transponder and one path back from the transponder to the receiver array location). Variations in these transponder observations are an additional source of information associated with the behaviour of variations in the ionosphere near the path midpoint.

To maximise the available data, the existing transponder network was extended to contain eight transponders in the area of interest. In addition, nine eight-hour data collection periods were made with the JORN radar at Laverton, producing OTH radar propagation observations to complement sounder observations over the course of the ELOISE trial. In order to understand the variability in HF propagation over a large region with fine azimuthal resolution, an experimental radar processing task was developed to enable simultaneous measurement of the Earth surface backscattered clutter power over the full 90° field of regard. This new processing task was interleaved with traditionally processed JORN transponder observations to generate JORN radar observations of ionospheric variability that were readily comparable to the ionospheric sounder measurements routinely being collected in the same region.

2.5. Evaluating ability to estimate ionospheric Doppler

To understand how the characteristics of inherent ionospheric motion (ionospheric Doppler or iDOP) compared to the rate of change of phase path estimated from ray tracing through the sounder electron density profile fits, ELOISE temporarily configured six of the OIS transmitters to transmit a small set of low power channel scattering function (CSF) waveforms along with the routine OIS

sweep (see Section 3.6 for more detail). To be permitted to transmit such waveforms required that the frequency be established to lie in an unoccupied frequency channel. For the frequency to be useful it had to be below the maximum observed frequency (MOF) for a given OIS path. On-board processing of the noise and MOF data associated with each path was completed in near real time to estimate if any given frequency channel was unoccupied so suitable waveform parameters could be selected. Any changes were distributed to each transmitter before the start of the next sounder run to keep up with channel occupancy changes and maintain OIS synchronisation. Typically this method enabled three single frequency CSF samples of iDOP versus delay to be observed for each 3.75 min ionogram observation.

2.6. Managing a large remote environmental data collection experiment

Managing such a large trial with equipment at 28 separate, remotely located sites, producing greater than 50,000 VIS, OIS and CSF observations per day, was a complex issue. To monitor these expanded networks, the communication links between various sounder sites were upgraded and a "sounder network health check" software tool was used to centrally monitor the accumulation of data from every sounder. To manage the large volumes of data, it was desirable to have sounder processing and data reduction occurring automatically and in real time at each remote sounder site. The processed data were distributed in real time to a trial management centre in Adelaide, South Australia, while the raw data were archived locally for later retrieval. This meant the automatic sounder processing had to be reliable, and the inevitable examples of poor data had to be automatically flagged.

3. A description of the experimental data and data qualities

3.1. Overview and typical data quantities

A key output of ELOISE was the observation of ionospheric variability with a variety of different and complementary ionospheric sensors. Because the number of sounders and the period of the experiment were large, the only practical way of quantifying the variability was to have reliable automatic processing of the data in real time on the various sounders. This section describes those data types and processing steps. Each ionospheric sounder produces data in the form of "raw ionograms", which provide two-dimensional frequency (MHz) versus delay or virtual height (km) representations of the propagation power and other observables. Each ionogram is typically a stack of these images. In the case of a VIS, the stack comprises two images representing the ordinary (or "O-ray") and extra-ordinary (or "X-ray") returned power (dBW). These raw images are modified or "cleaned" by algorithms that separate the wanted signal from the environmental noise

and spurious interference and generate a picture of the propagation in terms of signal-to-noise ratio (SNR) (Harris et al., 2016; Turley et al., 2019).

A presentation of the typical volumes of different types of valid sounder image data as a function of the day of year for the trial is presented in Fig. 4. In this and all subsequent results, times are measured in Universal Time (UT), with local times ranging from 7 to 10 h ahead of UT.

3.2. Ionogram feature data

Each of the different ionospheric sounders also executes additional processing to extract image "features" in the pattern of propagation observed within these cleaned ionograms. These features are influenced by the state of the ionosphere and the HF propagation path in use by the sounder.

In the case of O-X separated VIS ionograms, the feature extraction process is either ARTIST (Reinisch et al., 1997) or PRIME (Harris et al., 2016). The main feature derived from a VIS ionogram is a one-dimensional scalar "trace", a curve representing the observed propagation in terms of

one-hop delay versus frequency. With an assumption that the observed propagation path is from directly overhead, an O-trace can be used to define the overhead electron density profile, $N_e(z)$.

In the case of the OIS ionograms, collected using DST Group's Digital Oblique Receiving System (DORS) (Ayliffe et al., 2019), the O and X components of the image are unseparated and the "features" extracted from each image are typically a set of edges of propagation observed within the image. These VIS and OIS features characterise the ionospheric variation observed within the one-hop F2 and E regions. In the F2 region, three of these features are fxF2, MOFx_F2 and delay_F2 (Heitmann and Gardiner-Garden, 2019). A representative sample of values is presented in Fig. 5.

In the case of the QVIS (and VIS) systems, the image processing is forked and executed with two processing threads per image. One fork treats the data like the PRIME VIS processing after the images are O/X separated (and range converted to an effective zero range). The second fork processes the data like an OIS image (after the image is range converted to an effective 500 km range). This pre-
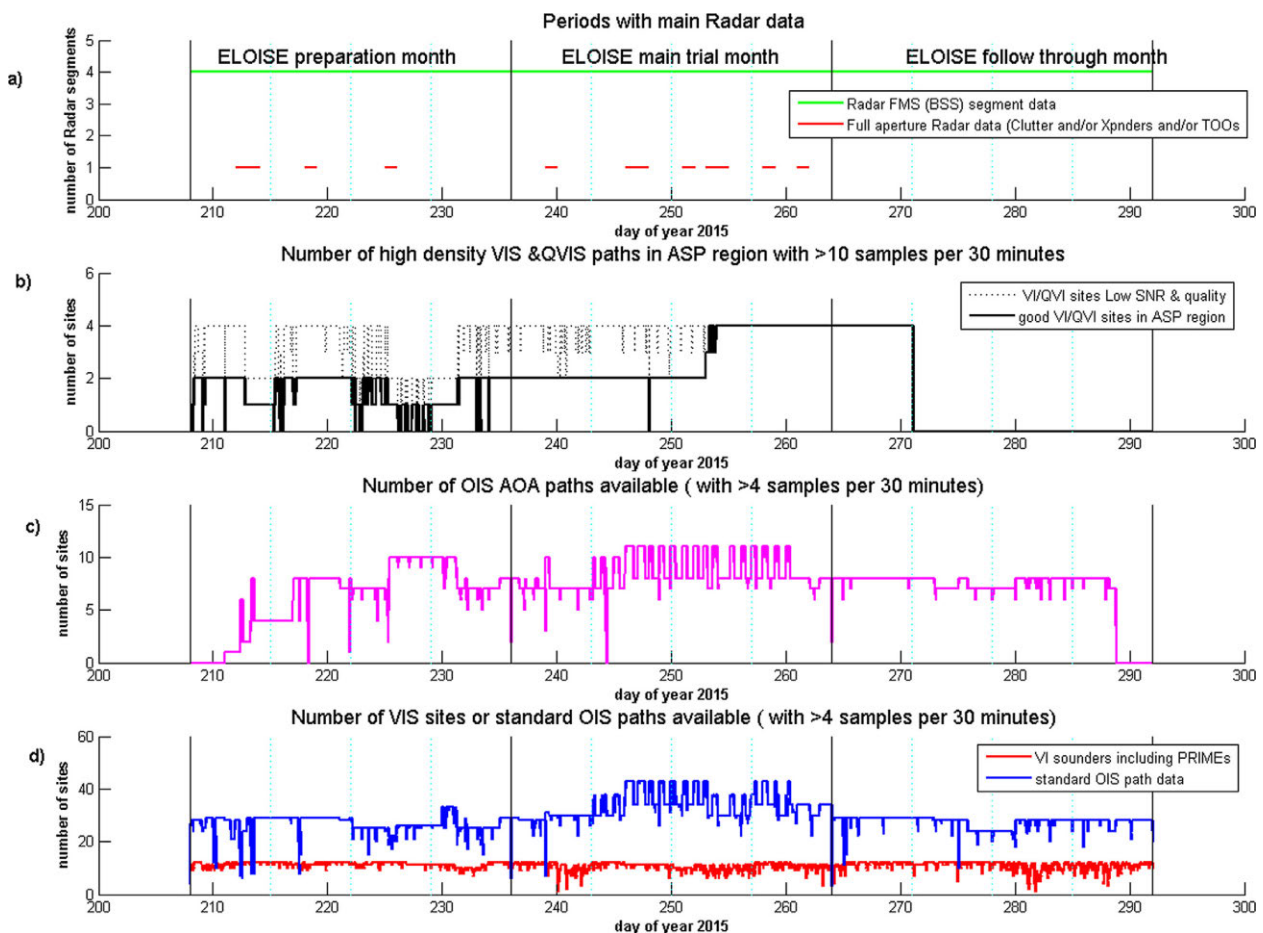
Fig. 4. A catalogue of the data availability from different instruments (a) the BSS and OTH radar back scattered surface clutter and transponder data; (b) VIS/QVIS network data near Alice Springs; (c) AoA OIS data and (d) the volumes of routine VIS and OIS data.
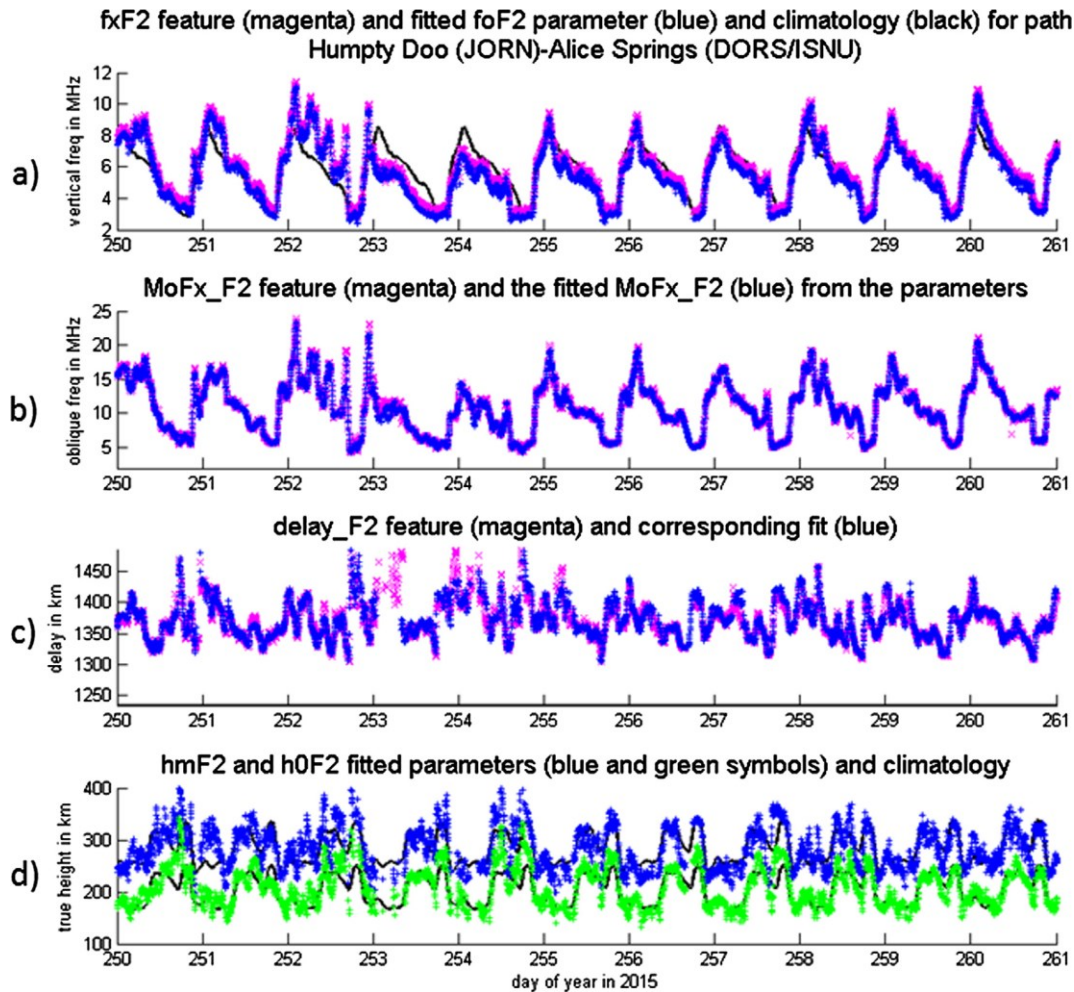
Fig. 5. F2 features (magenta) and fitted F2 parameter data (blue and green), plus 30 day median estimates of the climatological diurnal cycle (black) from a 1234 km long OIS path from Humpty Doo to Alice Springs. Panel (a) foF2 data (b) MOFx_F2 data, (c) delay_F2 data, and (d) hmF2 and h0F2 data. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

sents an opportunity to understand the effect of different processing methodologies on the estimated ionospheric parameter variation.

In the case of the BSS images (Fabrizio, 2013), the cleaned data are a stack of eight images from each radar segment, where each image is associated with the signal accumulated from a given azimuthal beam direction (with a significant extent in the range and azimuth footprint producing the accumulated backscattered power and significant complexity in interpretation being necessary because of the two way and bi-static nature of the propagation paths). The feature extracted by the image processing is an estimate of the "leading edge" (LE); i.e., the edge between a region of no propagation (or noise) and the area of backscattered surface clutter. Typically there can be multiple LEs per image; for example, one associated with the one-hop F returns and one associated with the edge of visible Es and potentially combinations of the two.

In the VIS, QVIS, OIS and BSS data, the sounder feature extraction and "features" do not assume anything specific about the HF propagation path other than that it is well defined, and that the algorithms can separate one-hop returns from other propagation features. The reliability and the accuracy of the characterisation of these features is a key to the value of any data. The accuracy of a feature estimate is dictated by the resolution of the sounder, the quality of the propagation conditions, and the quality of the equipment. Typical sounder image resolutions are presented in Table 1. The feature uncertainty scales in proportion to the image resolution. Distributions of the reliability of the data processing are presented in Appendix B.

In the case of the CSF sounder data, the image dimensions are Doppler frequency and delay and the typical resolutions depend on the waveform parameters selected. Typically the range resolution of the CSF is poorer than

Table 1
Typical VIS, OIS and BSS image resolutions and sample rates.

Sounder class	Type	Frequency limits (MHz)	Image frequency cell size (kHz)	Image one-way delay cell size (km)	Typical peak SNR (dB)	Sample rate (s)
VIS	Lowell DPS	1–22	100	2.5	10–30	225
VIS & QVIS	PRIME	2–22	40	1.25	10–30	75
OIS	DORS	2–36 or 2–45	40	1.5	30–60	225 or 450
BSS	JORN	2–45	200	50.0	20–60	300–450

the corresponding OIS because the bandwidth of the CSF waveform (typically around 10 kHz) is much smaller than the OIS processing range window (typically resulting from a bandwidth of 100 kHz). The Doppler resolution of the CSF is affected by the wave forms repetition frequency and number of repetitions. The automatically extracted features are, however, conventional; i.e., each CSF image can be characterised by a noise floor and a set of two-dimensional peaks where the signal is distinctly above the noise floor. An example is presented later in Fig. 11.

The feature extraction algorithms have been developed over years and tested over millions of ionograms (Galkin and Reinisch, 2008; Harris et al., 2016; Earl and Ward, 1987; Heitmann and Gardiner-Garden, 2019). Any consistent inability of the algorithms to estimate a set of ionogram features is usually due to, a major equipment fault, a cause flagged by the automatic processing, and/or an example of extreme ionospheric conditions. This is discussed further in Section 3.3.

3.3. Electron density profile parameter data

To compare the magnitude of ionospheric variations from the many different types of sounder features, it is desirable to interpret and quantify the observed feature variations on a common basis for every sounder; i.e., with a common parameterised model of the electron density and the propagation affecting the measurement. To this end, in the case of the VIS, QVIS and OIS data, each set of features is "fitted" with a simple parameterised model of the electron density profile $N_e(z)$ (Gardiner-Garden et al., 2017), assuming that the propagation is overhead in the case of the VIS path or that the propagation is symmetric about the great circle midpoint in the case of the QVIS or OIS data. The set of 10 parameters used to describe the ionosphere's $N_e(z)$ profile and fitted to each set of features is:

$$P = [f_{oE}, h_{mE}, y_{mE}, f_{oF1}, h_{mF1}, y_{mF1}, f_{oF2}, h_{mF2}, y_{mF2}, q_{F1}], \quad (1)$$

where f_{oE}, f_{oF1} and f_{oF2} are the critical frequencies of the three layers (E, F1 and F2); h_{mE}, h_{mF1} and h_{mF2} are the layer peak heights; y_{mE}, y_{mF1} and y_{mF2} are the layer half-thicknesses; and q_{F1} is the F1 layer join parameter (influencing the steepness of the join between the F1 and the F2 layers). Propagation path properties and complete synthetic OIS traces are estimated using an analytic ray-trace

(ART) solution to Haselgrove's equations (Chen et al, 1992) and a ray homing algorithm once this P (parameter state vector) is known. The notional base height of a particular ionospheric layer is the peak height minus semi-thickness e.g. $h_{0F2} = h_{mF2} \text{-} y_{mF2}$.

Variations in the patterns of the scalar "features" observed are hence transformed into variations within the "fitted parameters" deduced from these features. In the case of the VIS data, the fit to the trace features is overdetermined. In the case of the OIS data, the parameterised fit to the measured features specifies a unique solution. These fitted parameters come with their own associated quality factors (deduced from uncertainties in the features). At the core of ELOISE is a set of fitted parameter time series over 84 days at 55 different VIS and OIS geographic sample points characterising the observed ionospheric variations.

As with all experimental data, occasional missing data, data outliers and sampling noise all have a potential to complicate any interpretation of the patterns of observed variations. To this end, the raw parameters are tested for outliers (with respect to neighbours). Throughout the paper, parameter p is used to refer to any single component of vector P in Eq. (1). Observations of p from sample location j at time t are represented as $p_j(t)$. Sampling noise and data gaps for each time series are reduced by fitting a straight line through the data:

$$\widehat{p}_j(t) = \widetilde{p}_j(t_0) + (t - t_0)\frac{dp_j}{dt}\Big|_{t_0}, \quad (2)$$

where $\widetilde{p}_j$ is the locally smoothed estimated mean value of p and $\frac{dp_j}{dt}$ is the estimated rate of change over time, assumed to be constant for a small sample of data under test (typically, a fit of up to 5 samples or a 15 min window is used). This smoothed line approximation to the fitted parameter is more continuous and noise free and convenient for follow on analysis of the geophysical nature of the variation. Typical fit errors are less than 100 kHz for the OIS foF2 residual fit root mean square error (rms) and 5–10 km for the F2 height parameters residual fit rms. In the case of the QVIS data, a 5 point smoothing window represents a 5 min smoothed line approximation to the temporal variation in the fine scale fitted data because of the QVIS faster update rate.

An example of F2 feature and parameter variations from a single OIS path is presented in Fig. 5. The top three

panels show the features extracted from the ionograms. The bottom panel shows the fitted height parameters derived from these features. The features derived from the synthetic OIS traces are constructed from the fitted and smoothed parameters (using an Analytic Ray Trace (ART) model of the HF propagation). The values of the synthesised features are over-plotted upon the original extracted features within the top three panels. This figure shows the consistency of the features and fits and the dominance of the repeated climatological pattern of diurnal variation of the F2 region of the ionosphere, as well as a pattern of additional day-to-day and hour-to-hour variations.

To understand and visualise the pattern of behaviour of the variation on a finer scale and to visualise the size of the estimated uncertainty, the same data as Fig. 5 is presented in Fig. 6 but with a shorter time window and with the

30 day median of the diurnal cycle (the climatology) removed. The smooth parameter data with the spatially and temporally varying climatology removed are referred to as the 'parameter anomaly data'. A more detailed description of this anomaly field is presented in Section 4.2. These anomaly data enable a focus on the day-to-day and hour-to-hour variations within the F2 layer.

In this figure, the differences between the raw parameter data and the smoothed parameter data (i.e., smoothing errors) are so small compared to other disturbances that they disappear into the width of the line and only are visible within the enlarged bottom panels. The differences between the fxF2 feature anomalies and foF2 fitted parameter anomalies in Fig. 6(c) are because the same climatology value of foF2 is removed from both original time series leaving essentially the O-X geomagnetic split due to the gyro frequency in the fxF2 anomalies. The dominant
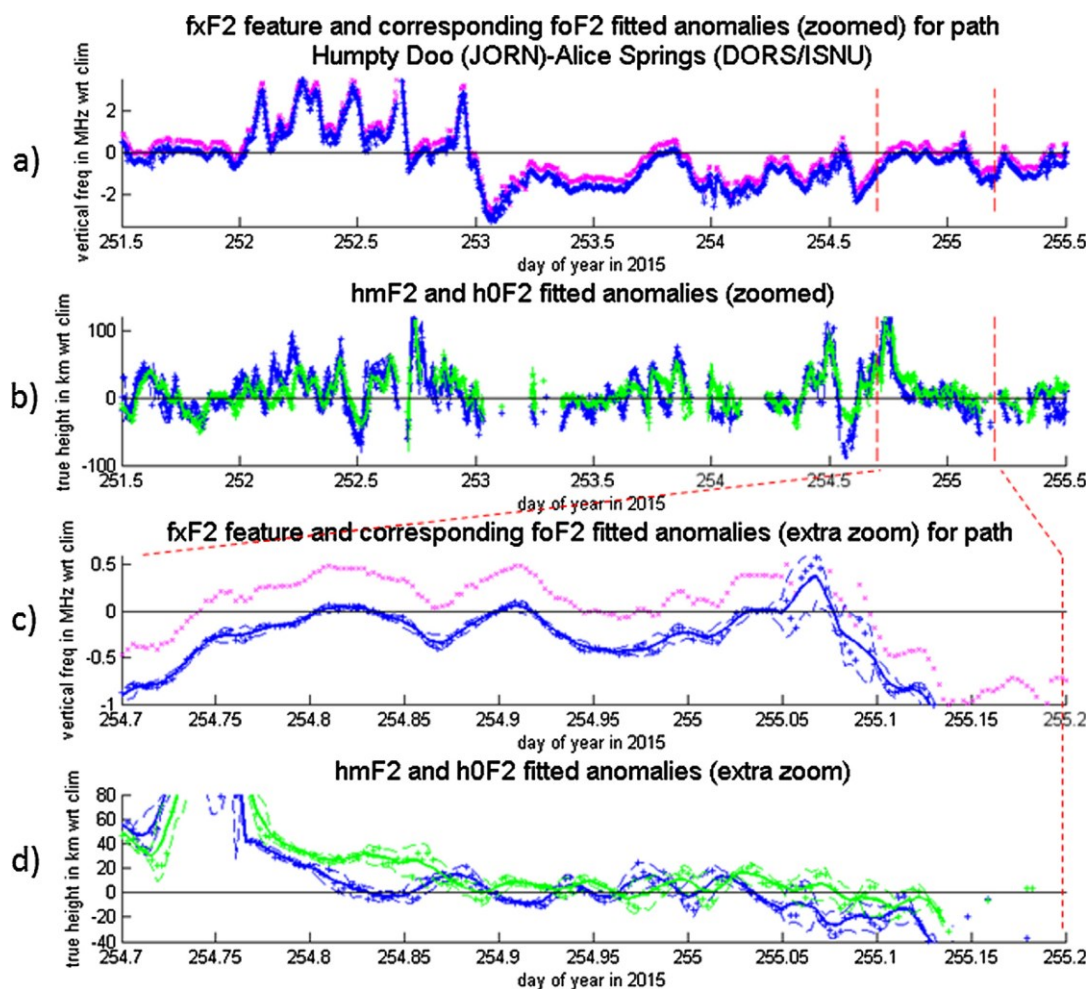
Fig. 6. Panel (a) the observed fxF2 feature anomaly data (magenta) and the fitted foF2 parameter anomaly data (blue) as per Fig. 5a, but for a sample of 4 days with the climatology removed. Panel (b) the fitted F2 height parameters hmF2 and h0F2 with the respective climatology removed over the same sample of 4 days. Panel (c) contains the observed fxF2 and fitted foF2 anomalies as per panel a, but with a shorter time window. The raw data is presented as isolated symbol and the smoothed fitted estimate is presented as a solid line (with dash lines at ± one standard deviation). Panel (d) contains the raw fitted data and the smoothed fitted lines of hmF2 (blue) and h0F2 (green) anomalies as per panel b, but with a shorter time window. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

pattern of disturbance for this particular period is the solar storm on day 252 and the negative storm phase that follows on day 253. The large spatial network of sample locations with ELOISE fitted data enables a description of the spatial pattern of these variations and the horizontal gradients in the regional ionosphere. The assumption that the oblique propagation is symmetric (aligned with a great circle path) has been observed (in separate more complete synthesis of the effects of observed large scale gradients) to have little effect on the synthetically reconstructed features. Regular bursts of small amplitude ripples in the parameters and features, potentially from small and medium scale travelling ionospheric disturbances (TIDs), see Hernandez-Pajares et al. (2012), are also observed.

3.4. Data quality

Before exploring the pattern of variation in the parameterised ionospheric anomaly fields, it is worthwhile to note that not all ionogram images necessarily produce spatially and temporally consistent parameters or fits. Periods of strong Es can obscure the F2 region of the ionosphere and this can be readily detected and reported. Periods containing a collapse in F2 ionisation (e.g. due to negative storms or rapid recombination just before dawn) can leave very little F2 region to be observed and parameterised. This also can be detected and reported. Typical rates of

availability of "good" fit data are described later in Appendix B. One of the more troublesome conditions is when the ionosphere is not particularly smooth or well-structured and hence the sounder propagation shows many multiple paths or spread within the image. This can be clearly visible in the sounder image to the naked eye but is not consistently identified within the sounder features and hence produces an incomplete or inconsistent parameterised fit.

An example of a single such OIS image is presented in Fig. 7 and a time series of slices through such images is presented in Fig. 8. ELOISE AoA measurements are used here simply to identify perturbed and multi-path propagation conditions that may affect feature and fit data; the AoA data properties themselves will be described in further detail in Section 3.5.

Figs. 7 and 8 show a small example of some of the ionospheric variation observed (in a similar fashion to Lynn et al., 2013) and are interpreted to show a substantially off-angle multipath return approaching the symmetric propagation path midpoint location and then moving away on the other side, on three separate occasions. These figures also show the propagation variation synthesised from the smoothed fit to each ionogram. While the disturbance is significant, because the poor fits are automatically detected, the state of the ionosphere reconstructed from the smoothed fit is still able to represent some of the vari-
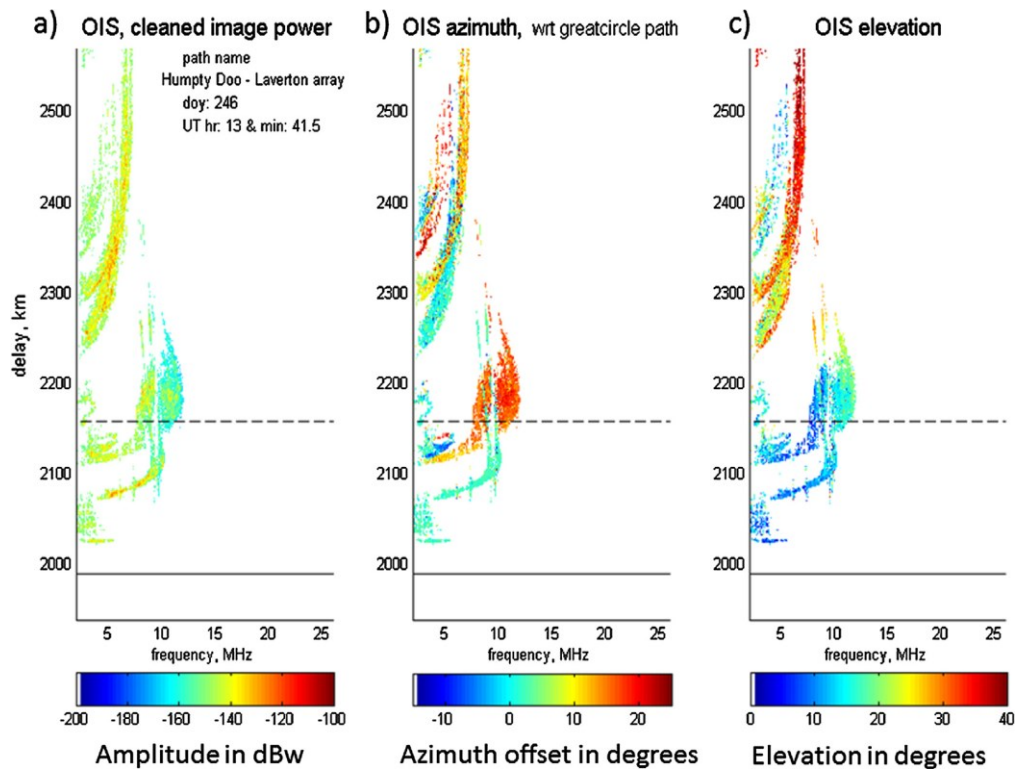
Fig. 7. An example of an AoA sounder observation from Humpty Doo to Laverton (1989 km) showing (a) the power of multipath returns, (b) the azimuthal direction (bearing) of arrival of these same returns, measured as an offset with respect to the great circle path, and (c) the elevation of arrival of the returns.
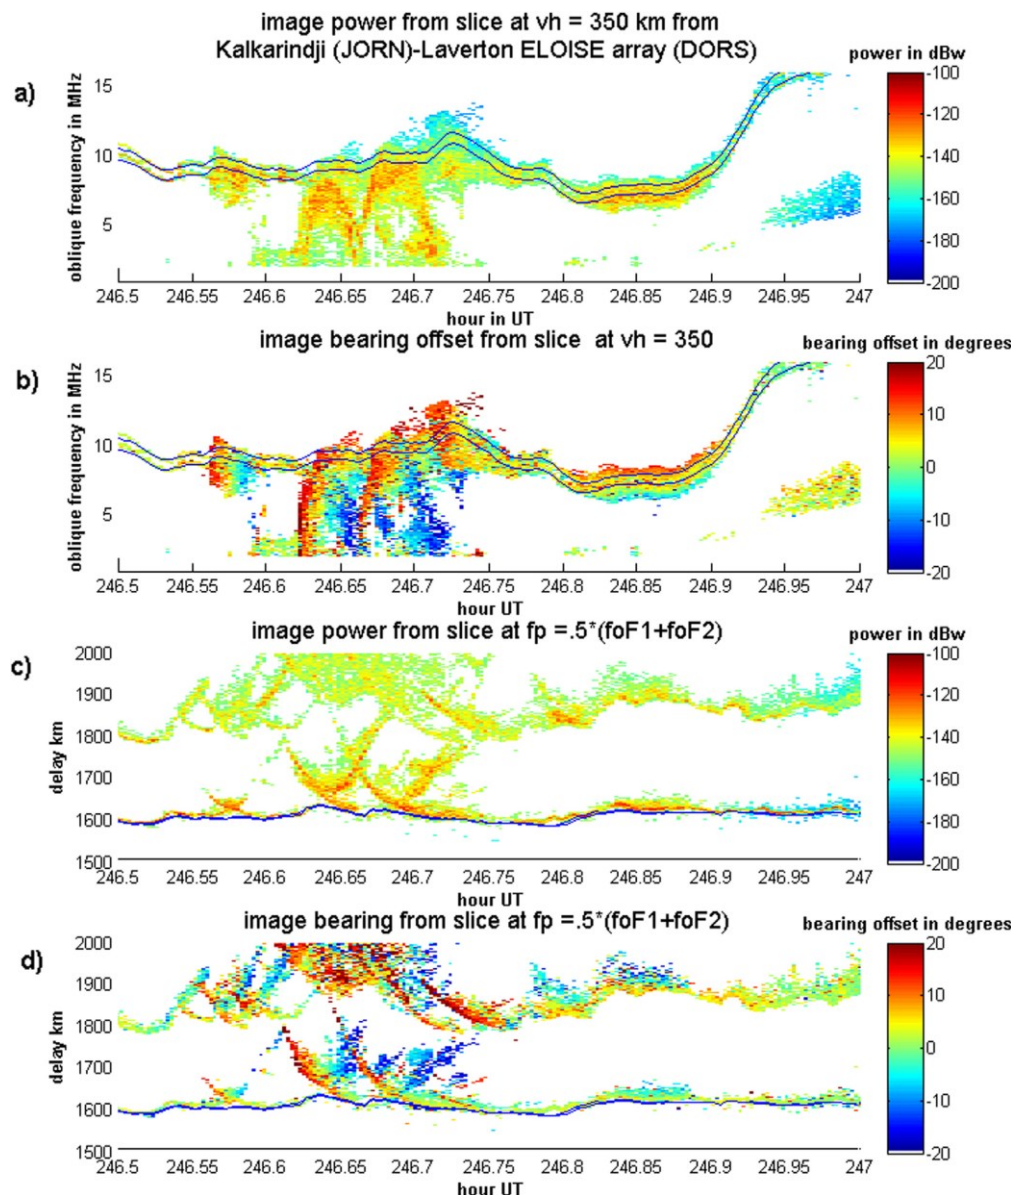
Fig. 8. Slices of AoA properties varying over time from a single OIS path, Kalkarindji into Laverton (1506 km); (a) a slice of OIS image power at an equivalent virtual height of 350 km for part of day 246 versus frequency and time; (b) a corresponding slice showing the bearing offset with respect to the great circle azimuth; (c) & (d) corresponding slices through the OIS image at a frequency half way between the observed large scale values of foF1 and foF2 showing image properties versus delay and time (c) the image power and (d) the bearing offset with respect to the great circle azimuth for the same path. Overlaid in black are the corresponding estimates of frequency and delay from the slices extracted from synthesised OIS traces based on the fitted parameters.

ation in the dominant features. Despite ELOISE being conducted in the equinox season and at a descending phase of the solar cycle, when mid-latitude spread-F is observed to be common (Hajkowicz, 2007), these night-time spread-F conditions did not create an entire breakdown of the parameterised description of the ionosphere except for a limited number of runs on a limited number of nights (less than 1% of the night time hours lost more than 20 min of data within an hour long interval when the path lengths is less than 2000 km).

3.5. AoA measurements

The OIS AoA data generated from the ELOISE arrays allow the assumption that oblique propagation is mainly symmetric about the great circle midpoint to be tested. The OIS AoA data record is a stack of five OIS-like ionogram images where one image is the conventional power or SNR and the other images (on the same domain) are the estimated azimuth, elevation and the uncertainties in these two measurements. This set of measurements enables vari-

ations in the oblique ionospheric propagation path with respect to expectations to be studied directly.

An example of a complete single OIS AoA image is presented in Fig. 9 for the OIS path from Harts Range (near Alice Springs) into Laverton (1383 km). A limited sample of data derived from these images at a single time-varying equivalent vertical frequency and accumulated over time (representing a slice of the AoA data over time) is presented in Fig. 10 demonstrating small deviations from great circle propagation and elevations that differ from the equivalent mirror expectations based on delay and large scale measured gradients.

In the first instance, the array calibration was performed manually using known reference sources in the far field to estimate array alignment and small inter-element timing correction factors. With these in place, geophysical values of the measured azimuth and elevation can be described with respect to the unperturbed symmetric path; i.e., the azimuth and elevation expected given the path range and delay. These offset values of azimuth (or bearing) and elevation are the differences between the measured angles and the predictions based on the measured great circle range assuming one-hop mirror geometry.

These complete AoA ionograms can be "reduced" in order to simplify the image data for subsequent analysis. Reducing an image involves producing a set of multi-dimensional peaks over all vertical frequencies in the iono-

gram and just retaining those peaks that are within a threshold (30 dB) of the strongest peak returned at that equivalent vertical frequency (or height). This enables minor alternative propagation paths that are present and observed to be eliminated from most follow-on analysis. In general, these figures show a small sample of the raw observed azimuth and elevation deviations away from the expected case of great circle and symmetric mirror like propagation paths.

A small (1°) azimuthal calibration error is not included in this plot of raw data (panel c). Fig. 10 presents a time series of these observed deviations at just one frequency (out of many possibilities) and for just one path (out of the 11 OIS AoA paths available).

3.6. CSF measurements

Direct measurement of ionospherically-induced Doppler (iDOP) is another aspect of the ELOISE observation and characterisation of ionospheric variations. Observations by HF sounders are influenced by both changes in the underlying electron density content and by changes in the propagation path (phase path) that are part of the measurement process. Successive OIS electron density profile fits enable the indirect estimation of phase path changes and hence iDOP between adjacent soundings. The interleaved channel scattering function (CSF) measurement
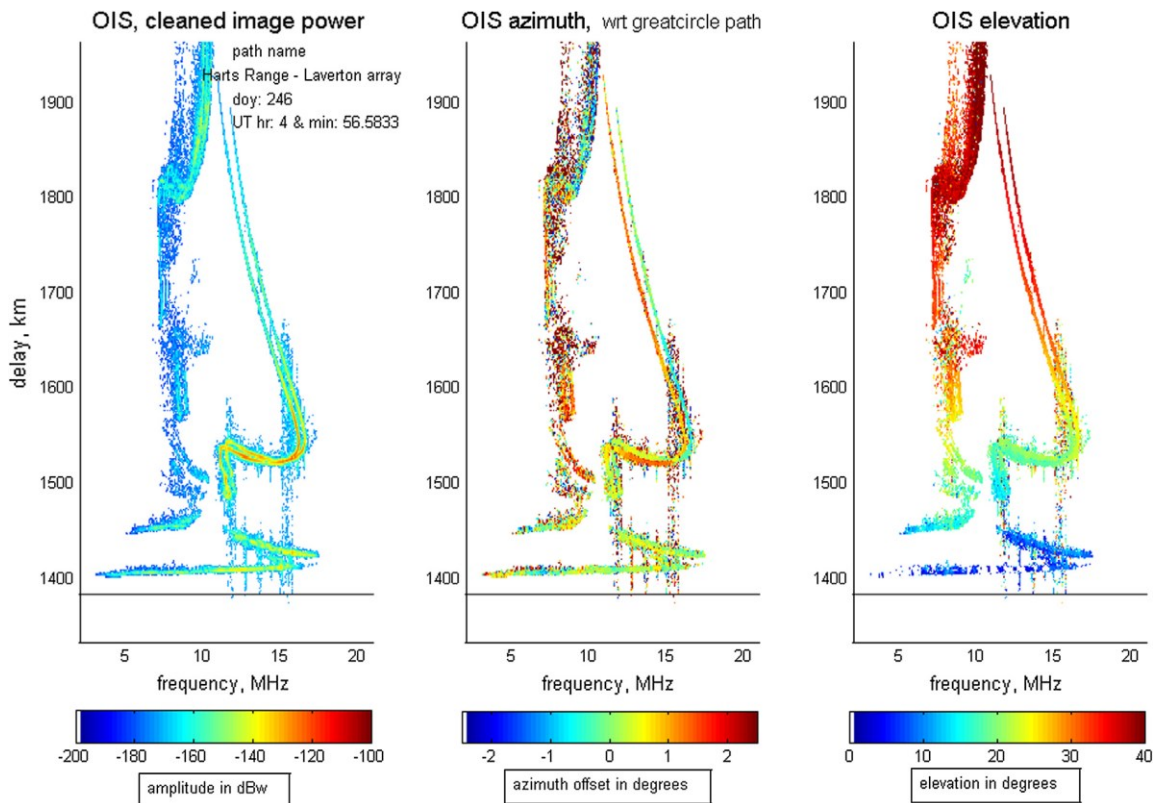
Fig. 9. An example of AoA sounder observations as per Fig. 7 but for a less disturbed ionosphere showing (a) the power of returns, (b) the azimuthal offset (with respect to the great circle bearing of the path), and (c) the elevation of arrival of the returns (with respect to the expected mirror one-hop delay).
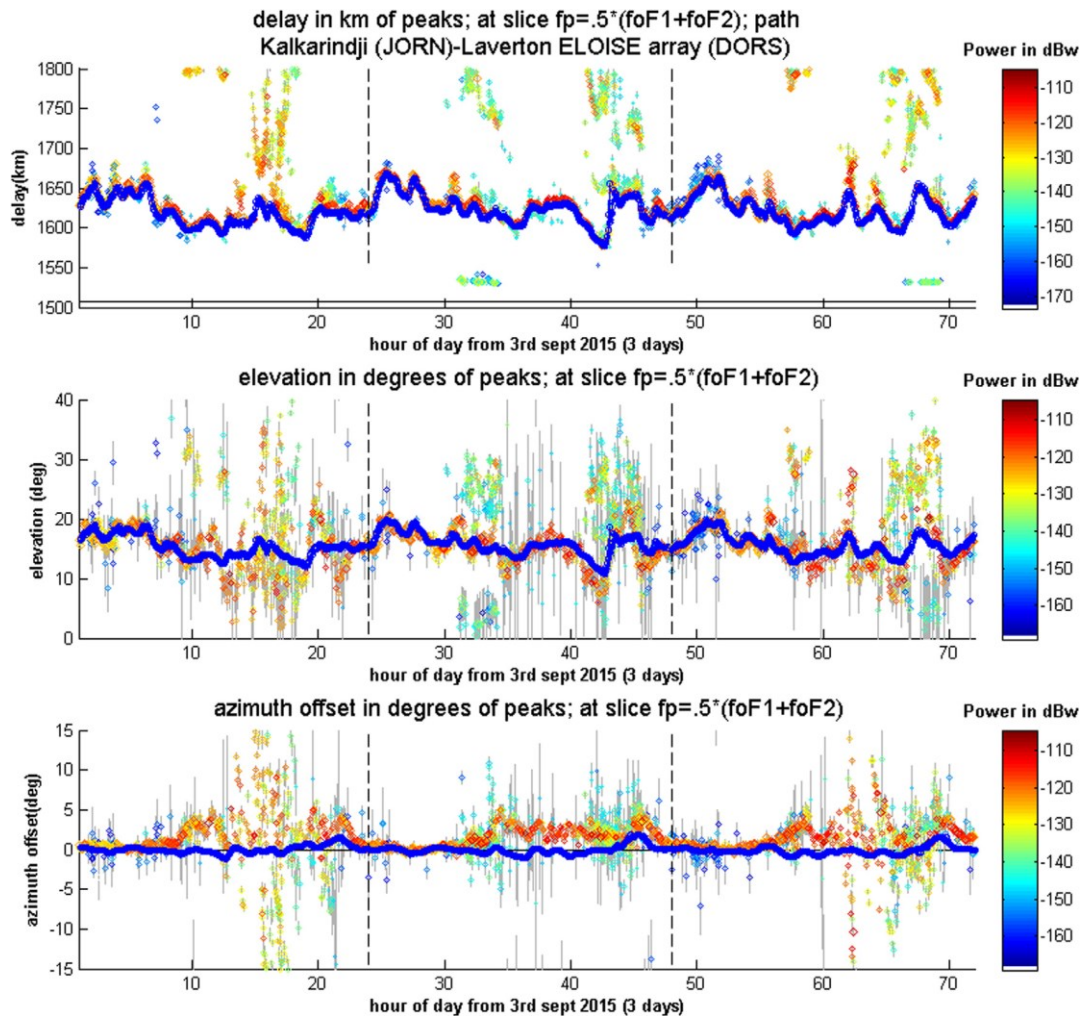
Fig. 10. Slices over time of the OIS peaks extracted from many ionograms recorded from a single OIS path and at a particular time-varying equivalent vertical frequency, fv = 0.5*(foF1 + foF2). Panel (a) describes the power in the extracted peaks as a function of signal delay and time of day, and the synthesised (parameterised) OIS fits to the raw OIS images versus delay and time. Panel (b) describes the power in the extracted peaks versus peak elevation and time and also the elevation from the corresponding synthesised ionograms (estimating and assuming gradients in the large scale fit described in more detail in Section 4.2). Panel (c) describes the power in the extracted peaks versus peak azimuth and time and also the azimuth from the corresponding synthesised ionograms (estimating and assuming gradients derived from the large scale fit). The grey lines in all three panels represent the width of all the estimated peaks (where the power had fallen 6 dB with respect to the maximum).

makes a comparable direct estimate of this iDOP. This method was demonstrated previously (Harris and Frazer, 2005; Lynn, 2009) and its application to ELOISE is described in more detail by Heitmann et al. (2018).

ELOISE utilized six OIS transmitters to also transmit CSF waveforms. Two sites (Laverton and Coondambo) had OIS receiver channels scheduled to synchronise with these CSF transmissions. This gave a total of 12 one-way simultaneously measured CSF paths (and typically 3 frequencies per path per 3.75 min epoch). A time-staggered lane scheme, with a fixed 10 Hz waveform repetition frequency (WRF), was implemented to avoid co-channel interference between the six CSF transmitters. To accommodate the additional processing and CSF dwell time,

the frequency limits of the wide-band sweep were reduced from 2–45 MHz to 2–36 MHz, without any change to the 3.75 min revisit rate. Three time-varying frequencies were selected at fractions of the observed MOF for the path in order to allow the reflection height of the iDOP data to span the F2 region of the ionosphere.

An example of a single CSF run for a 1383 km path is shown in Fig. 11, alongside measurements from an ionogram that probed the same frequency 73 s later. The same group delay structure is clearly seen in both images, albeit at poorer resolution in the CSF owing to its lower bandwidth waveform.

A comparison of the direct ionospheric Doppler measurements against those synthesised from fitted ionospheric
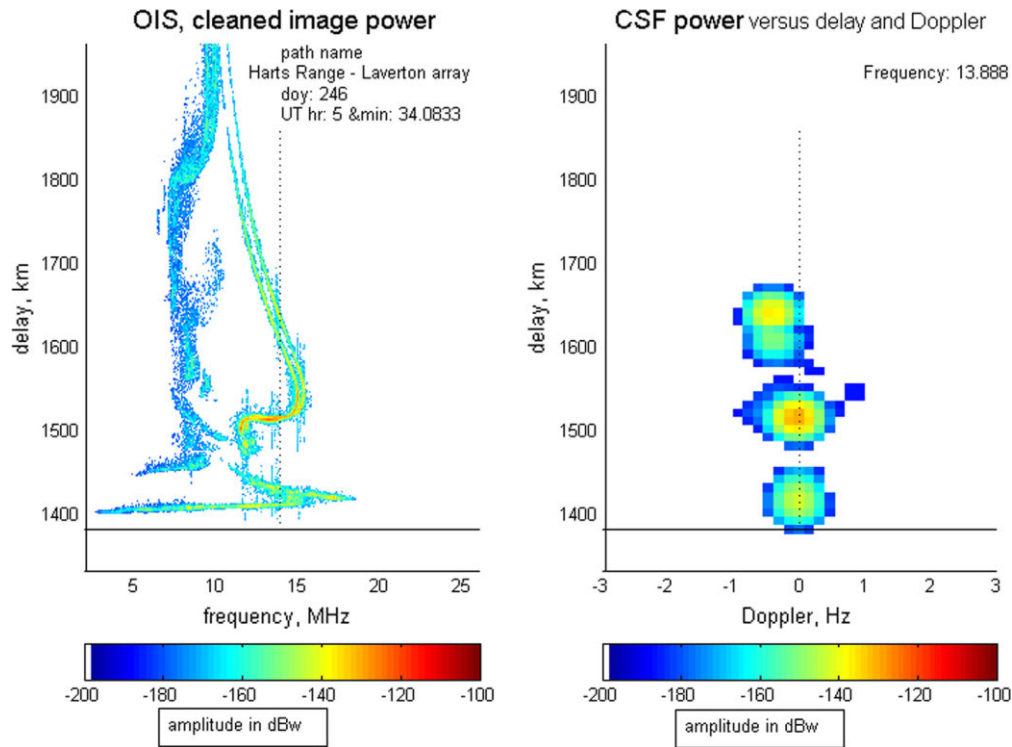
Fig. 11. A sample ionogram (left) and the corresponding CSF range-Doppler map at 13.888 MHz (right), for a 1383 km OIS path.

profiles, will be the subject of a separate paper. The CSF data can be subjected to a similar analysis to the other sounder data where the characteristic features are now the image peaks (fitted Gaussians functions in delay and Doppler). Nine out of the twelve CSF paths were also collected and analysed on the AoA arrays and hence include AoA properties in addition to just propagation path SNR.

3.7. Backscattered surface clutter measurements

Direct measurement of two-way backscattered HF propagated surface clutter is another way ELOISE observes ionospheric variability in the region. The JORN has a network of four backscatter sounders (BSSs) that

operated 24 × 7 throughout the ELOISE period. Early versions of these sounders are described at length by Earl and Ward (1987). Like other HF sounders, variations in the principal observed features (leading edges) results from a combination of ionospheric content changes and propagation path changes, and in the case of backscatter, the two-way mixing of different modes of propagation. An example of this clutter power and leading edge variation is presented in Fig. 12.

These observations of ionospheric variability are continuous in range but are averaged in azimuth (due to the relatively wide beam widths, varying between 18° at 5 MHz to 3° at 30 MHz). This provides a complementary way of measuring the continuous spatial pattern of particular
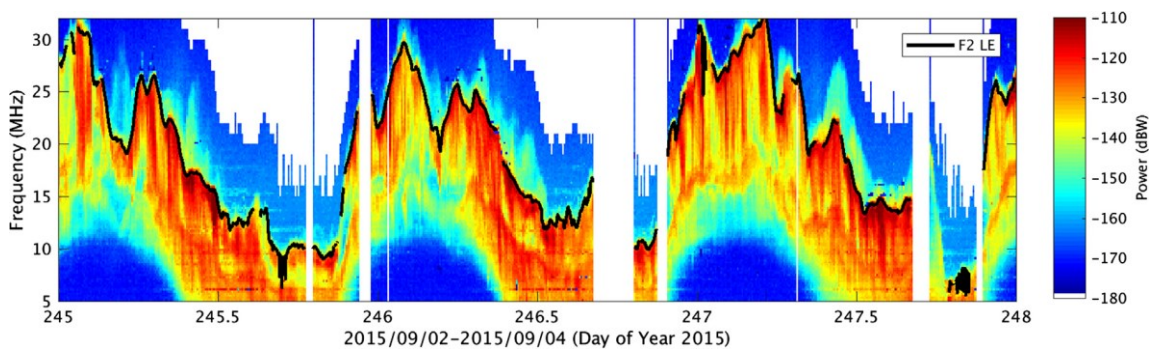
Fig. 12. BSS observations of the backscattered surface clutter (and corresponding leading edge frequencies, F2_LE) as they vary over time and frequency from a single BSS beam and at a fixed group delay, in this case with a 2833 km one-way group delay.

ionospheric variations compared to the more pointwise measurements reflected within the OIS fits. Any OIS (fitted or observed) needs to be moderately modified (to represent the bi-static path differences, the integrated azimuthal beam width of the BSS and the differences between a back scattered leading edge and an OIS MoF) if a precise comparison between OIS variability and BSS leading edge variability is to be measured and understood. This is the subject of separate work not discussed further in this paper.

JORN radars observe two way backscattered surface clutter with an even better sensitivity and azimuth resolution compared with BSS sub-systems but only at a limited number of frequencies. Any JORN radar can be configured to operate in an adaptively selected frequency channel and flood all radar azimuths simultaneously. This generates backscattered clutter over a wide area but at a fixed frequency. ELOISE obtained approximately nine 8-hour blocks of this type of radar observation, including maps of Earth surface backscattered power and observations of transponders. An example of the OTH radar Earth surface backscattered power is presented in Fig. 13 (see Holdsworth, 2017). In this case the OTH radar clutter observations have a much narrower beamwidth (varying between 1° at 5 MHz to near 0.2° at 30 MHz), a range cell size of 15 km and a typical revisit rate of 60 s.

4. Preliminary analysis and discussion

ELOISE produced a very large and high fidelity dataset of ionospheric observations that enabled many different geophysical processes to be characterized in great depth. The variability observed during ELOISE included: several ionospheric storms (e.g. on day of year 240, 252–253 and 281); some days with substantial mid-day and pre-dawn collapses of f_{oF2}; many days with many periods of general geomagnetic disturbances and large scale and medium scale

TID activity; and several nights with distinct (and well sampled) mid-latitude spread-F. While there were some periods of sporadic E, this was seldom dominant and will not be covered further in this paper.

4.1. The climatology

One of the key concerns of ELOISE was to observe ionospheric variability on small and large (space and time) scales simultaneously so that small scale perturbations can be properly understood in the context of associated large scale variability. External climatological models of the large scale variability in the diurnal cycle are based on spatially averaged historical observations indexed by estimates of the smoothed sunspot number R12 (v1), an index capturing the broad effects of the external solar forcing via ionising EUV and X-rays. Examples are IRI (Bilitza et al., 2017) or the climatological model in use within JORN (Barnes et al., 2000). The extended period of ELOISE data collection (three months) enabled the observed 30 day median climatological diurnal cycle to be derived directly from the data itself with great spatial fidelity.

The ELOISE derived climatological variation provided a much finer scale description of the monthly median diurnal cycle in both space and time compared to the historically smoothed climatological maps of ionospheric parameters used within JORN. These are not qualitatively different from the external IRI or JORN estimates within northern Australia, but do establish the amplitude and pattern of the migrating tidal component of variability more accurately than the smoothed historical data. In addition, this 30 day analysis enabled an estimation of the "normal" amplitude of day to day and hour to hour variations. These maps were found to be smoothly varying from month to month and will be reported on in a separate paper.
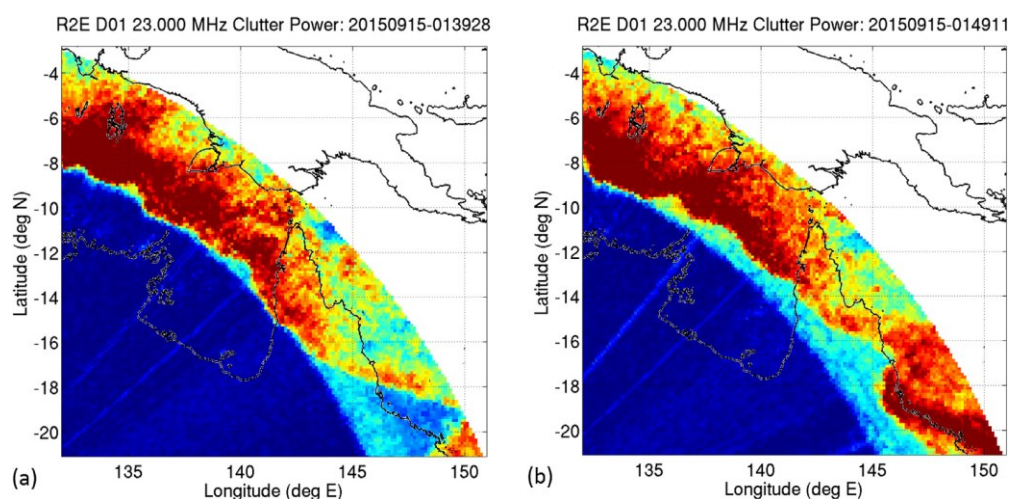
Fig. 13. Radar observations of the back scattered clutter with respect to the radar at Laverton in Western Australia at a single frequency but at two different times (ten minutes apart). This shows an example of clutter variation observed near the North-East Australian (Queensland) coast that is suspected of being an ionospheric electron density perturbation passing over Alice Springs in central Australia (the approximate path midpoint).

4.2. The patterns of large scale variability

While the climatology explains a great part of the observed ionospheric variability our interest here is also on the pattern of the day to day and hour to hour anomalies in the ionosphere that are deviations from the 30 day median climatological pattern; i.e., the anomaly field where:

$$p = \bar{p} + \Delta p + \epsilon p \qquad (3)$$

where p is the observed value of a given ionospheric parameter, $\bar{p}$ is the expected climatological value of this parameter, Δp is the total anomaly, i.e. the deviation from the climatological pattern, and ϵp represents "noise", consisting of both geophysical noise and sampling errors. The ionospheric parameter "noise" is identified and removed by temporal smoothing as part of the parameter post processing, described by Eq. (2). Its amplitude is typically small, being less than 100 kHz for the foF2 frequency parameter and less than 5–10 km for the F2 height parameters.

The dominant spatial correlations within the F2 parameter anomaly variations can be visualised by examining the pattern of semi-variance between all the 55 separate data locations over the core ELOISE trial period and presenting the derived semi-variance as a function of the inter-site separation between sample location j and sample location k defined by:

$$\gamma_{jk}(\Delta p) = \frac{1}{2N} \sum_{t=1}^{N} \left(\Delta p_j(t) - \Delta p_k(t) \right)^2 \qquad (4)$$

where γ_{jk} is the semi-variance of anomaly parameter Δp between sample locations j and k, and $\Delta p_j(t)$ is the total anomaly of parameter p observed at sample location j at time t.

Fig. 14 shows the semi-variance of the f_{oF2} anomaly data and the h_{mF2} anomaly data, both in terms of the longitude and latitude separations between samples (Δx and Δy) and in terms of the combined and scaled one dimensional separation distance metric $\Delta s = \sqrt{\left[(\Delta x/2)^2 + \Delta y^2 \right]}$.

The latitude versus longitude aspect ratio of the variations can be estimated from the data itself but a 2:1 aspect ratio is generally a good approximation. This simple and common functional form has been adopted to enumerate
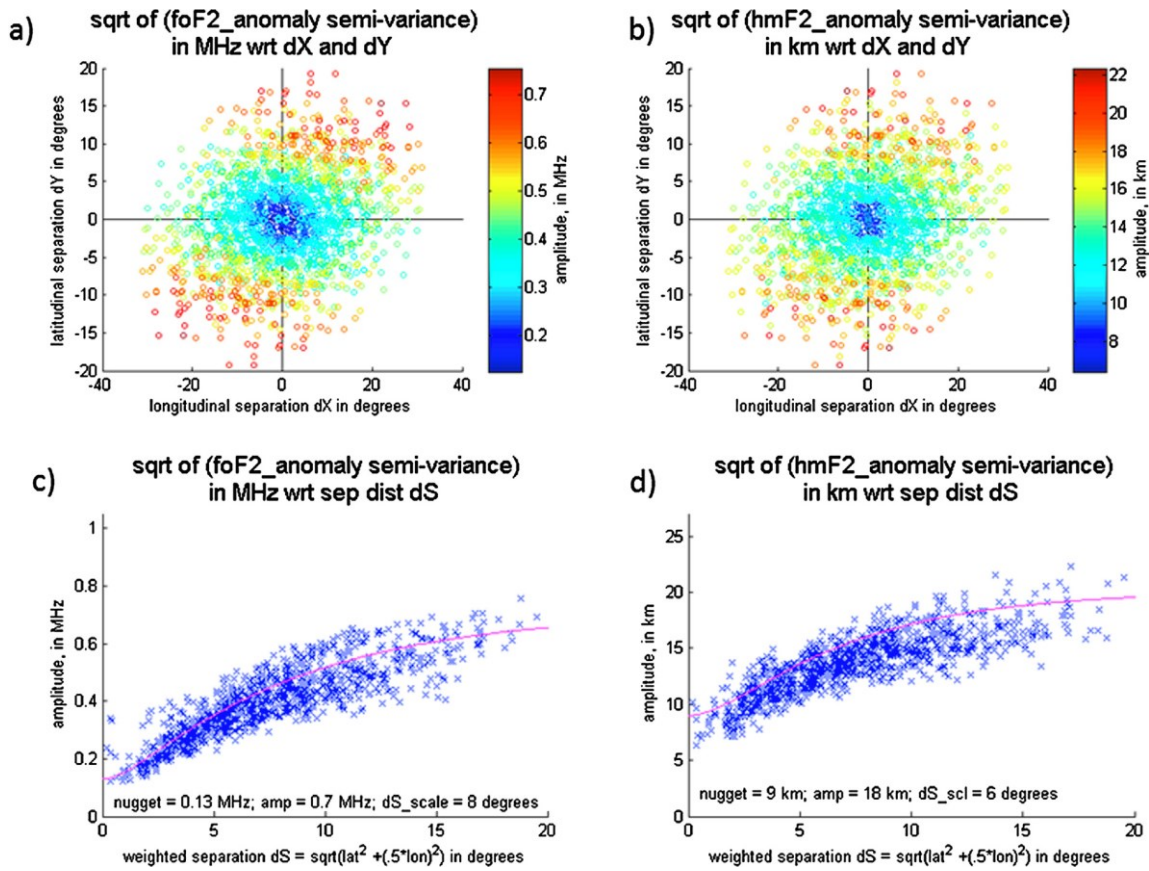
Fig. 14. Semi-variance of the f_{oF2} anomaly data (a) as a function of the longitude and latitude separations between sample locations Δx and Δy and (c) as a function of the joint and scaled separation distance Δs. Over plotted in magenta is the fitted functional model used to describe the semi-variance. The semi-variance of the h_{mF2} anomaly data is presented in panels (b) & (d). It has a similar functional shape but different amplitudes. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the dominant spatial scale of correlated variations. This shows that most of the f_{oF2} anomaly variation is highly correlated over spatial scales of 5 degrees of latitude and 10 degrees of longitude. A related calculation shows that the same is true for h_{mF2} and h_{oF2} anomaly parameters. An additional calculation of the lagged autocorrelation of the Δf_{oF2}, Δh_{mF2} and Δh_{oF2} samples shows that the F2 anomaly parameter variation is commonly well correlated for periods longer than 30 min.

These large decorrelation scales show that this ELOISE network of sounders is adequately positioned to sample a great fraction of the spatial variations observed in the F2 parameter anomaly data. This suggests that a useful model representation of F2 parameter variations can be constructed by fitting a simple linear model to the surrounding sounder data. The (total) anomaly term in Eq. (3) can hence be further broken down into a "large-scale anomaly" δp and a "residual anomaly" μp as follows:

$$\Delta p = \delta p + \mu p. \tag{5}$$

The large scale anomaly field can be characterised by adopting a functional form that has locally linear gradients in space and time. Estimates of the large scale anomaly field are derived from a weighted least squares fit to the data, such that, for the grid point at longitude x_0, latitude y_0 and time t_0 (which may or may not coincide with a sample location), the local planar model estimate $\widehat{\delta p}$ has following form:

$$\widehat{\delta p}(x,y,t) = \delta \widetilde{p}(x_0,y_0,t_0) + (x - x_0)\frac{\partial(\delta p)}{\partial x}\bigg|_{x_0,y_0,t_0}$$
$$+ (y - y_0)\frac{\partial(\delta p)}{\partial y}\bigg|_{x_0,y_0,t_0} + (t - t_0)\frac{\partial(\delta p)}{\partial t}\bigg|_{x_0,y_0,t_0}. \tag{6}$$

For the large scale F2 parameter anomaly estimates, for example in Fig. 15, the weights within the calculation of the fit weighted root mean square (rms) error are constructed to decrease with increasing spatial separation between the
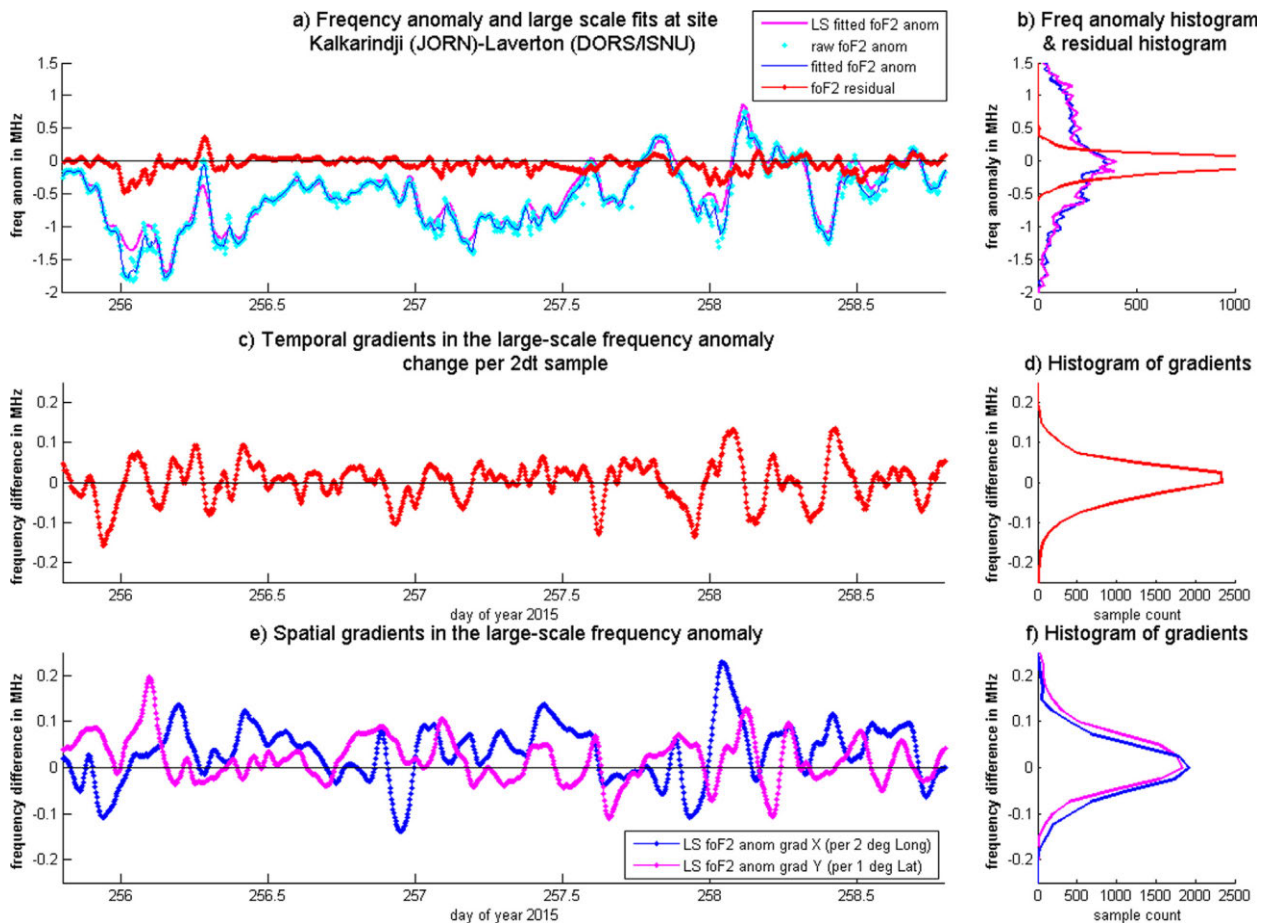
Fig. 15. (a) A time series of three days of the foF2 anomaly data (including the total raw anomaly, in cyan and the fitted anomaly (in blue) and both its large scale part, in magenta, and its residual part, in red) at a single sample location, and (b) the accumulated histograms of these foF2 anomaly variation values over 30 days. The corresponding time series of large scale gradients in the anomaly data with respect to time is reproduced in panel c in red. The histogram in panel (d) shows the distribution of this rate of change of the large scale foF2 anomaly. The spatial rate of change is reproduced in the next two panels. The blue is the change with respect to two degrees of longitude and magenta is the change with respect to one degree of latitude. These are presented in panel (e) as a time series and panel (f) as a histogram from 30 days of data. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

point where δp is being estimated and the original sounder data location. Results are derived from data accumulated only over the surrounding 30 min (i.e. $\pm$ 15 min) enabling the large scale parameter estimates to rapidly vary. This linear model quantitatively describes the large scale variation of any parameter, at any grid location or any sounder sample location, including its local gradients.

The ELOISE analysis adopts a scale of 3.5 degrees with respect to latitude and 7 degrees with respect to longitude within the spatial weighting function so that a common set of spatial weights are applied to all F2 parameters. This scale is large enough to give a smooth well sampled (and well populated) estimate of the large-scale anomaly field with respect to the network's sounder separations.

Fig. 15 shows a small sample of these anomaly variations once the climatological variation has been removed. A time series of the residual data (the original smoothed data with both the climatology and the large scale value removed) is also presented in red in the top panel. All the preceding analysis can be executed on each and every iono-

spheric parameter (including E, F1 &F2 layers) to produce a complete electron density profile at any time and location. The spatial scales determining the weights are found to be larger for the E and F1 parameters (5° latitude × 10° longitude) compared to the F2 parameters (3.5° latitude × 7° longitude). The magnitude of the residual is typically small. The large scale fit to the data represents a large fraction of the observed ionospheric variations and the large scale gradient effects, but leaves unrepresented the finer scale and smaller amplitude (but potentially steep) residual variations. This residual part of the F2 variation is typically less than 200 kHz for the f_{oF2} parameter or 10–15 km for the true height parameters h_{mF2} and h_{oF2}. The temporal gradient in the fit to the large scale variations in foF2 anomaly are presented as gradients in foF2 over two data sample times (2*dt) i.e., 7.5 min. The spatial gradients in the fitted foF2 anomaly variation are presented over one and two degrees of latitude and longitude respectively. This is basically the same ratio of variation as the underlying scale lengths adopted by the algorithms, that
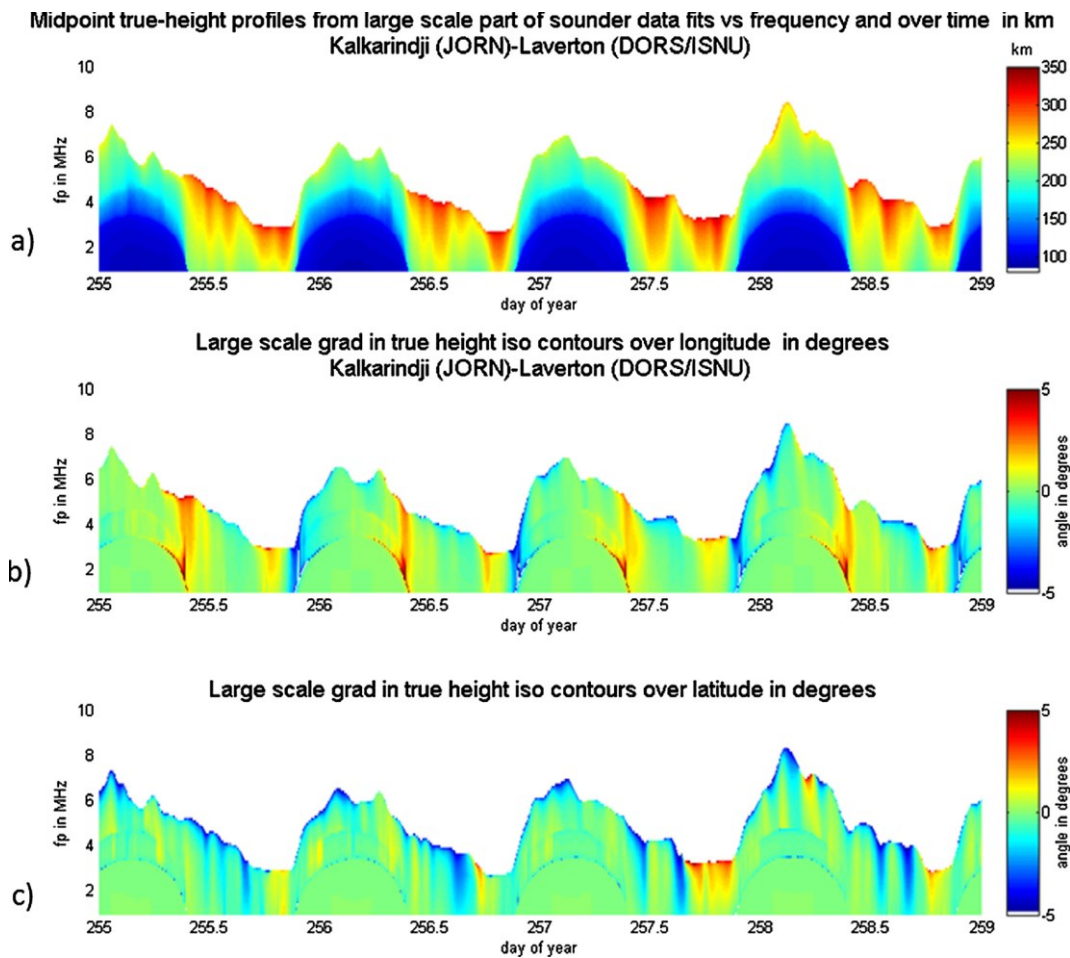
Fig. 16. A time series of (a) the true height of each plasma frequency value, (b) the corresponding large scale $N_e(z,x,y)$ surface gradient expressed as a tilt angle in degrees with respect to longitude, and (c) the large scale surface gradient in degrees with respect to latitude, for a few days at the midpoint of the Kalkarindji to Laverton OIS path.

is, 3.5 degrees of latitude, 7 degrees of longitude and 30 min of time.

4.3. The patterns of gradients and propagation path angle of arrival

A key goal of ELOISE was to better understand the gradients estimated and observed within a sounder based model of the ionosphere. The relationship between gradients in each of the separate parameters and the final electron density profile at any single height or frequency is quite complex. Fig. 16 combines the various parameters and presents (in panel a) the estimated true height of the overhead electron density as a function of plasma frequency f_p, which in turn is directly related to the electron density profile $N_e(z)$. The corresponding large scale horizontal gradients in electron density (and hence the oblique propagation refractive index) as a function of f_p (and hence true height) are also presented in panels (b) and (c). These gradients are the slopes of constant plasma frequency (f_p) iso-surfaces where each fitted parameter has a contribution to the final slope.

This figure shows that apart from the regular periods of gradients at dawn and dusk there are other significant periods of gradients when large scale disturbances ripple through the domain. This is a useful presentation of the gradients that are likely to affect the oblique propagation. The large gradients near the peak of the layer have little impact on the propagation that does not reach those heights. In addition, the in-plane and cross-plane components of electron density gradients will have a very different impact on the oblique propagation properties. To convert these estimated gradients into asymmetries in the oblique propagation properties measured by an AoA system requires a functional model of the HF propagation that depends on the path range and aspect. For simplicity the model adopted is a modified version of the ART used before. The large scale gradients described in Figs. 15 and 16 are used to estimate the synthesised OIS path elevation and azimuth data presented and over-plotted in Fig. 10.

The example (see Fig. 10) of the AoA observations and inferred AoA estimates (derived from the fitted and parameterised ionograms) is sampled at a frequency f_p that varies
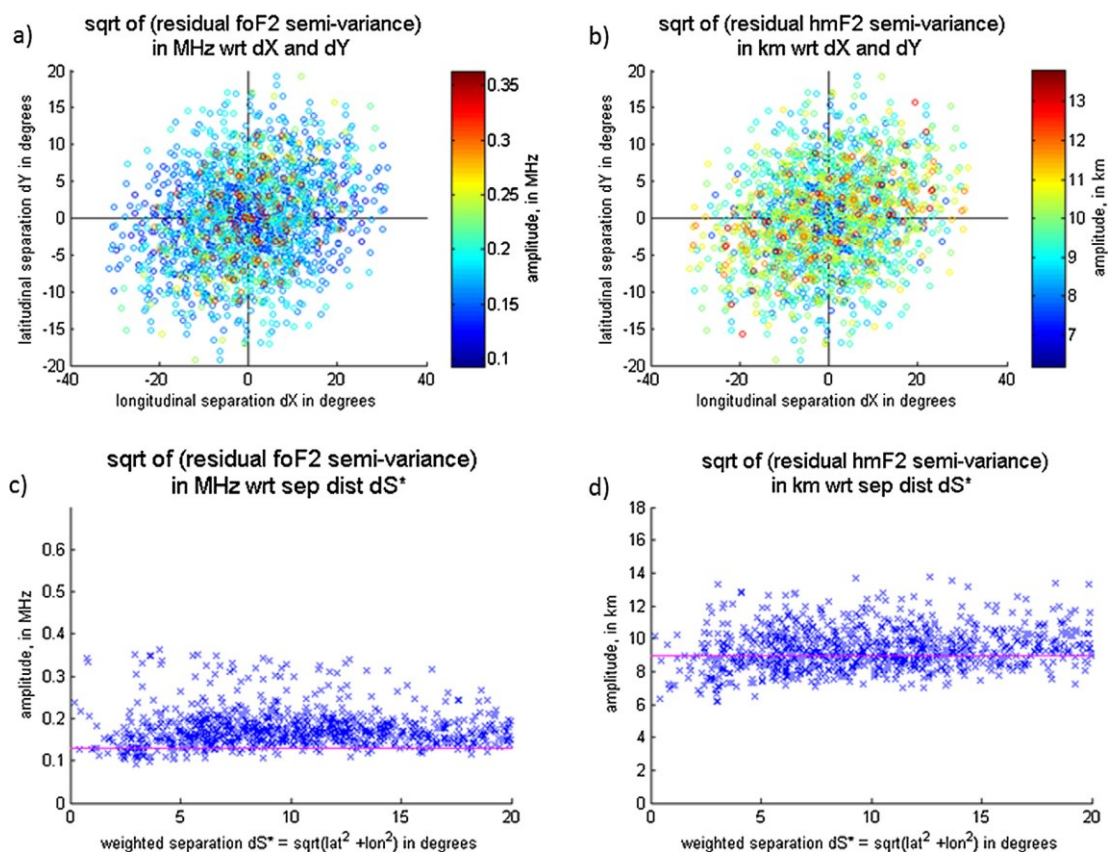
Fig. 17. (a) The square root of the semi-variance (semi-deviation) of the foF2 residual anomaly $\gamma_{jk}(\mu f_{oF2})$ as a function of latitude and longitude and (b) a similar spatial semi-deviation for the residual hmF2 parameter. Panels (c) & (d) are the corresponding semi-deviation data presented with respect to the normalised and symmetric inter-site separation distance: (c) for the foF2 residual anomaly and (d) for the hmF2 residual anomaly. Over plotted are the estimated nugget (uncorrelated noise) parameters (magenta lines) from the fit to the total anomaly data presented in Fig. 14. Note that the 2:1 aspect ratio factor between longitude and latitude does not manifest in the residual semi-variance calculation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

over time such that it is half way between the large scale estimates of f_{oF1} and f_{oF2}. Figs. 16 and 10 shows that a large part of the observed path delay, elevation and azimuth variation is explained by a combination of changes in the midpoint $N_e(z)$ profile and the horizontal gradients in the large scale N_e. We note that in the presentation of azimuth within Fig. 10, apart from a known one degree calibration bias (not included in the plot of raw data), there appears to be an additional unexplained consistent nighttime difference on the scale of 1–2°. This is much smaller than the very large excursions observed on many nights but is never-the-less important and the subject of ongoing analysis. Observations of propagation with additional alternative (multiple) paths between the sounder's two end points are observed but are most commonly weaker in amplitude (lighter in colour) and are a small overall fraction of the observed variability.

Because the large scale gradients in iso-surfaces of N_e are typically small (tilt angles of less than 3 degrees), it is observed that for OIS paths with ranges <2000 km, the large scale disturbances displace the effective apogee away from the OIS path midpoint by only a moderate distance (<100 km) and hence the effective N_e profile affecting the propagation is not greatly different from the OIS path midpoint profile. This in part justifies the common assumption used when integrating large scale VIS and OIS data into a single common model, that the OIS data represents the N_e profile at the OIS path midpoint. This may not be true later when trying to estimate the much smaller residual parameter variations.

4.4. The patterns of residual small scale variability

The final part of this report describes something quantitative about the pattern of variation visible in the residual and small scale part of the observed ionospheric variations. The semi-variance diagrams such as Fig. 14, when derived from the parameter residual values (μp), show no consistent spatial cross correlations (at zero time lag) using a complete set of ELOSIE sample sites. An example of this
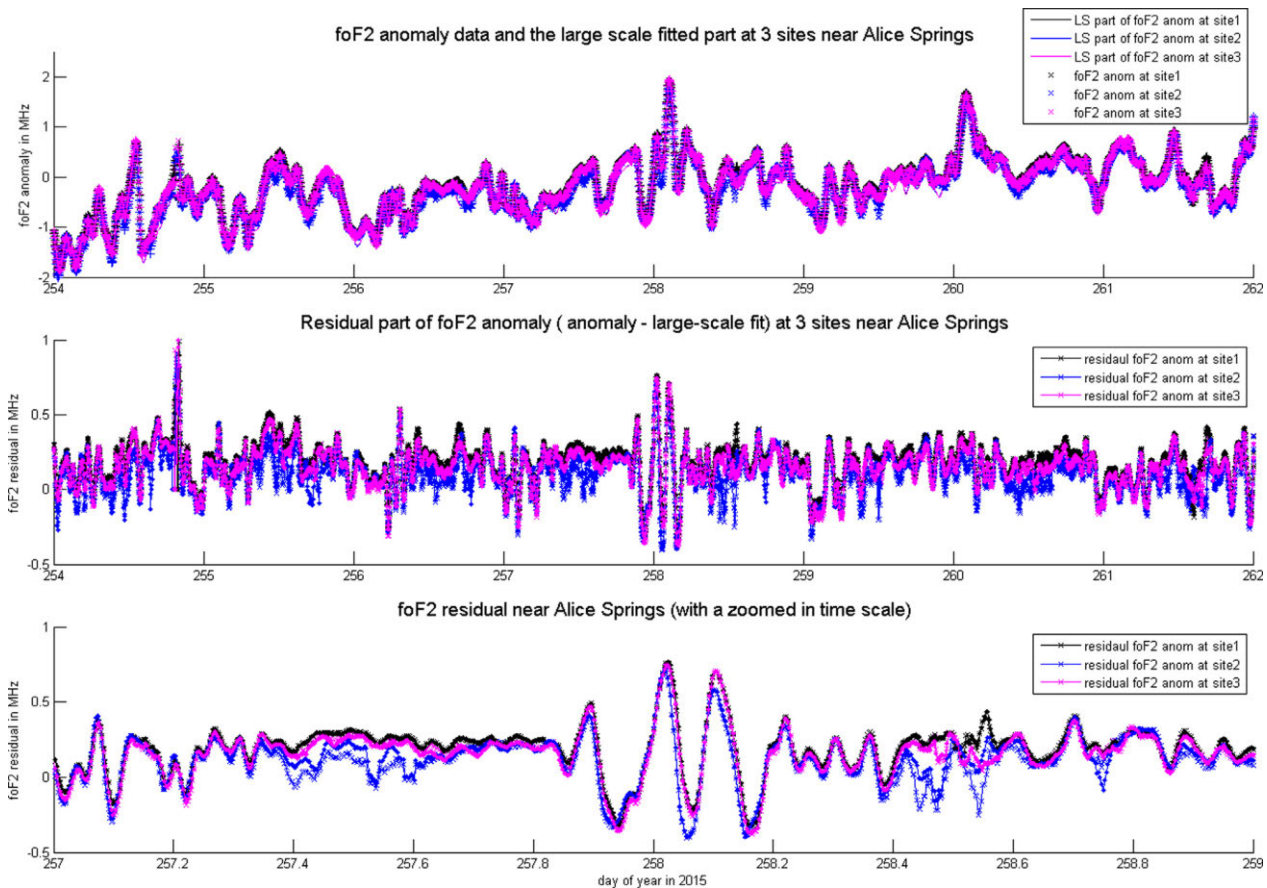
Fig. 18. A time series of the foF2 parameter anomaly variations and differences between sample locations with 20–50 km separation (within the fine scale QVIS network) around Alice Springs. Panel (a) is the total foF2 anomaly Δf_{oF2} at each of the 3 different QVIS sample locations in 3 different colours. Separate lines for the large scale part are almost indistinguishable from the total anomaly data when plotted on this scale. Panel (b) shows the corresponding foF2 residual μf_{oF2} estimates (when the large scale part of the foF2 variation has been removed). Panel (c) shows the same data as panel (b) but for a zoomed (4×) time scale of 2 days to show up the small spatial lag between different sample locations. On this scale, the data from the two separate antenna arms at a given site is sometimes distinctly different. In general two separate lines from each separate path or site are indistinguishable and labelled as one curve.

pattern of residual semi-variance is presented in Fig. 17. We note that the 2:1 aspect ratio factor between longitude and latitude used in the normalised separation distance Δs does not impact the basically flat pattern of residual semi-variance.

The apparently random pattern could be used to suggest there is little consistent structure in the sounder residual data but a more careful examination of the pattern of residual variations tells a different story. It is valuable to first look at the spatial differences in foF2 residuals over the small spatial scale of the QVIS network (20–50 km) and compare the fine temporal differences that are visible (after representing and removing the effects of large scale gradients between locations). This is presented in Fig. 18 for a small sample (8 days) of a single parameter (f_{oF2}) from the ELOISE analysis period using the best 3 QVI sites available (those with the longest continuous records).

This presentation shows that some part of the local gradient is from the large scale variations but an important part is also from the smaller scale disturbances (suspected in part to be medium scale TIDs (Hocke and Schlegel, 1996). A different set of spatial and temporal weights within the large scale fitting process would have left a slightly different pattern of variation within the residual anomaly variation because the broad spectrum of parameter variations is not completely separated between large scales and small scales. Never-the-less the change would be small and the separation is useful when interpreting variation because the large scale variation is very smooth on the scale between samples and the residual amplitude is typically small. The differences in the residual foF2 values between the 3 sample locations (displayed in Fig. 19) are when averaged typically less than 100 kHz and represent variations on the same scale as the measurement accuracy. This means that reliably and quantitatively measuring the lags and variations in disturbances across the spatial domain is a subtle process. Never-the-less, the amplitude of the residual foF2 variations (the non-large scale part of the data) can be significant (200–500 kHz) for some periods of time. On this fine scale, the data derived from the two separate antenna arms (traditionally used to separate o &x within a sounding, one receiver per arm) at a given
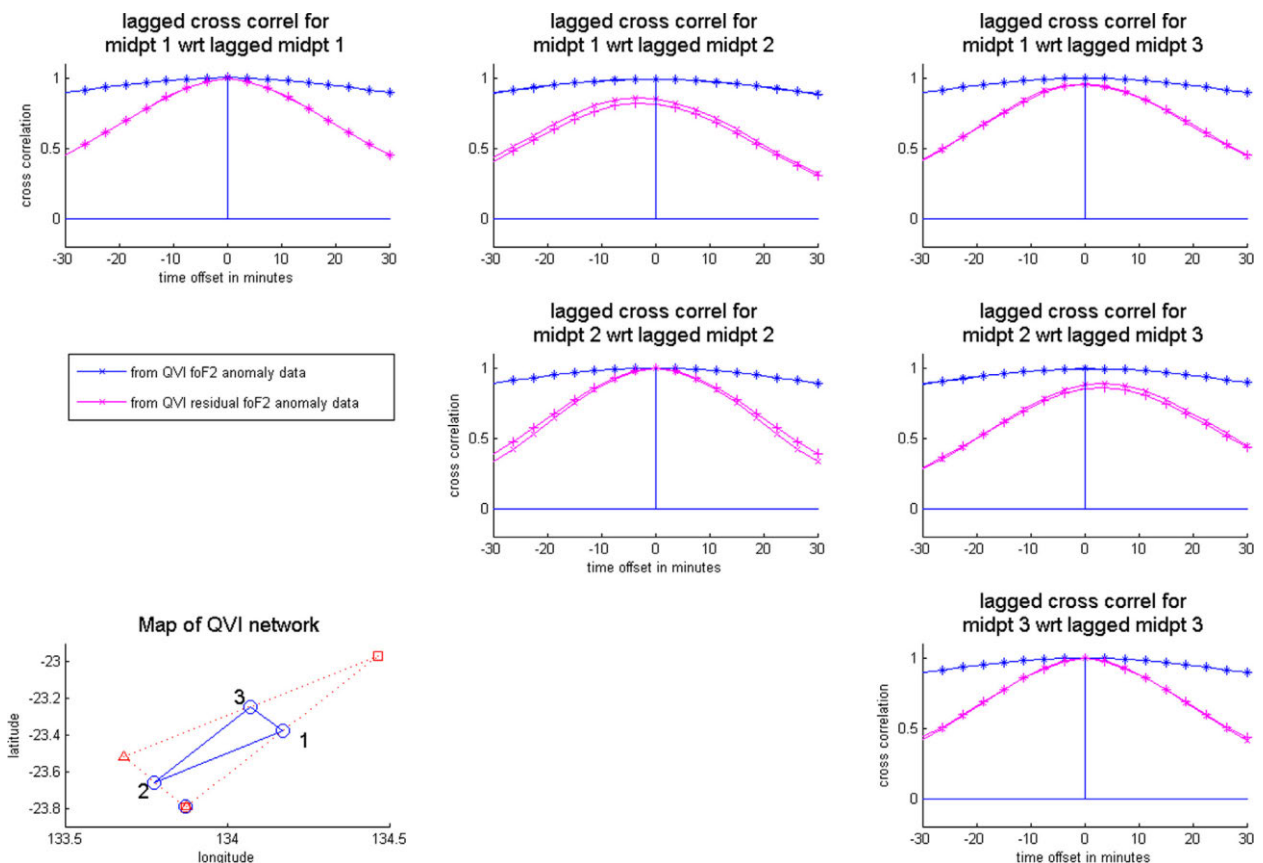
Fig. 19. Plots of the lagged cross correlated dependencies of the large scale anomaly δf_{oF2} (blue) and the smaller scale residual μf_{oF2} (magenta) parameter data within the QVIS network near Alice Springs. Diagonal elements show the local lagged autocorrelation functions (sites 1 versus 1, 2 versus 2 & 3 versus 3) and hence show the time-scale of the different components of variation. Off diagonal panels show the cross-correlation versus lag (for sites 1 versus 2, 1 versus 3 & 2 versus 3). The panel in the bottom left shows the spatial layout of the network of QVI sounder midpoints (in blue) in detail. On this scale, the data from the two separate antenna arms at a given site is sometimes slightly different. In general two separate independent lines based on data from each separate path or site are indistinguishable and labelled as one curve. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

site are sometimes distinctly different. In general two separate lines from each separate path or site are indistinguishable and for convenience are labelled as one curve. In fact, each line is made up of data from too highly correlated signals. During disturbed periods with poorer measurement quality it becomes apparent when two independent measurements of the residual foF2 from the same path (data collected with different receivers attached to different antenna at the same location) occasionally produce significantly different residual values. These less coherent spread conditions are only a small part of the overall residual variation observed.

To dig deeper into the patterns of variation within the residual (smaller scale) part of the F2 variation, it is valuable to examine the time lagged spatial cross-correlation within both the original anomaly parameter data Δf_{oF2} and the residual parameter data μf_{oF2} (when the sample locations have a small spatial separation, for example 20–50 km). This is presented in Fig. 19 which shows the strong consistency over the total time series (11 days from day 254–264 in this case) between the variations of the ionosphere on both a small and large spatial scale, but also that there exist small (but detectable) temporal lags in the small spatial scale variation. While the anomaly data (blue curve

in Fig. 19) hints at the differences and lags between the sample locations, removing the large scale data and examining the lagged cross correlation in the residual data (the magenta curves in Fig. 19) shows these differences in clearer relief because the lagged autocorrelation function is narrower. Never-the-less, Fig. 19 shows the consistent part of the lag over the total time series is small and on the order of less than 3–5 min. This does not mean that periods of stronger correlated patterns are not present in the data, just that they vary over time of day and day of year and Fig. 19 is an averaged result.

In addition to looking at the inter-site correlation in the residual of a single parameter within the QVIS network, this paper also examines the cross correlated pattern of variation in between the different residual parameters $\mu f_{oF2}, \mu h_{mF2}, \mu h_{oF2}$ and $d(\mu f_{oF2})/dt$ at each sample location within the full VIS and OIS network. This is presented in Fig. 20 and shows a consistent pattern of lag at all ELOISE sample locations between the different parameter residuals describing different aspects of the F2 layer variation. The diagonal panels show the lagged autocorrelation for each residual parameter and are indicative of a broad spectrum of wave like disturbances not particularly well correlated after one wave period. The cross correlation in hmF2 and
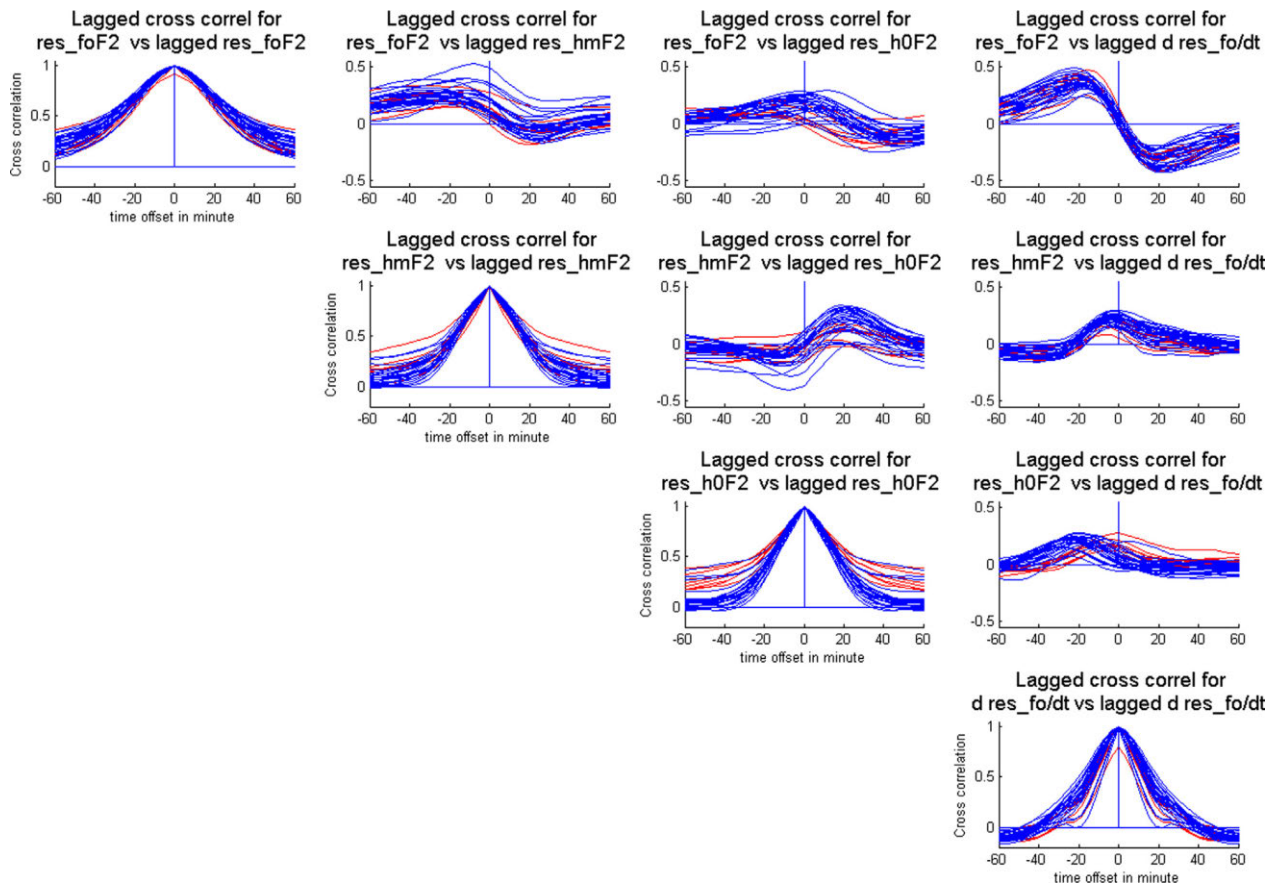
Fig. 20. An accumulation of the lagged cross correlated patterns of μfoF2, μhmF2, μh0F2 and d(μfoF2)/dt at all VIS and OIS midpoint locations. Red lines are data from VIS sites and blue lines are data from OIS path midpoint fits. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

h0F2 residual variations (variations in the peak height and base of the F2 layer, i.e., row 2 and column 3) shows a significant part of the h0F2 residual variations (30% of the variance) consistently lags the hmF2 variations by 15–20 min (row 2 and column 3). Interestingly, variations in the rate of change of residual foF2 ($d(\mu f_{oF2})/dt$) (column 4) show a single dominant band of periods in the disturbances (rows 1 and column 4) and it consistently leads the h0F2 variations at most sites in the region by 15–20 min (row 3 column 4).

5. Summary and conclusions

In summary these preliminary ELOISE results show that:

(1) The greatest part of the F2 layer variability is well sampled and well modelled by the large scale F2 parameter variation; i.e., explaining most of the observed F2 layer variations and gradients (which start de-correlating over scales of 3.5 degrees of latitude and 7 degrees of longitude, see Fig. 14), and leave unexplained variation amplitudes of less than 200 kHz in foF2, and 10 km in hmF2 and h0F2. There are multiple geophysical causes of these large scale variations, from variations in the climatological equilibrium and migrating tides, to solar storms and the large scale TIDs.

(2) The parameter sampling error in each individual parameter estimate is typically smaller than the small geophysical variations observed (see the residual parameter cross correlations in Fig. 20) and can be reliably removed by fitting a smooth curve (15 min piecewise linear fit) to the raw data. This means that the small scale variations (with durations over 10–15 min) can frequently be reliably measured at specific sample locations even if it is difficult to model them.

(3) The horizontal gradients within the large scale variations observed during the ELOISE period had a relatively small impact on the OIS AoA properties observed, compared to the direct impact of virtual height changes (see Fig. 10). This resulted in most of the observed OIS path elevation variation being directly and strongly associated with OIS midpoint profile variations. This did not exclude the possibility of short periods of distinct changes of gradients within the large scale ionospheric variation, changes that created multipath effects. Never-the-less, the occasions when multipath returns were produced by large scale N_e surface curvature were highly localised. Some issues concerning the calibration of (or potential biases in) the measured azimuth deviations remain and are the subject of ongoing analysis.

(4) The dominant time scale (quarter period) of small residual variations is 10–30 min but the period changes over time and space. This causes any long term statistical averages to wash out the dominant period of variations (see Fig. 20) but distinct periods of variation remain visible within the analysis of particular places and times.

(5) The one property not washed out within large accumulations of data, the small scale residual parameter variation, has a clear cross correlation between the F2 parameters at different altitudes. Typically 30–50% of the residual geophysical F2 height variation is found to be consistently lagged, i.e., disturbances are descending through the F2 region. Typically residual μh_{0F2} lags residual μh_{mF2} by 10–15 min. Another less dominant but consistent cross correlation is that the residual μf_{oF2} rate of change leads the residual μh_{mF2} variation by a few minutes (something starts to force a μf_{oF2} change and the F2 peak height follows). This all suggests that there is enough good data available to drive a model for the representation of the residual part of the F2 variation, but this subject is too big and subtle for further discussion in this preliminary analysis.

In conclusion, this paper has briefly described the breadth and depth of ionospheric measurements produced by ELOISE. ELOISE observed ionospheric variability in the Australian region on an unprecedented scale, using multiple networks of high quality HF sounders to produce ionospheric measurements over an extended region of space and period of time. The experiment was designed to provide high density sampling (in space and time) of ionospheric variability to test and validate improved models of the ionosphere and investigate different instrument responses to small and large scale ionospheric disturbances. It was expected that observations with a high revisit rate collected over an extended period of time would create an opportunity to observe a variety of temporal and spatial ionospheric perturbations caused by different geophysical processes and this was largely achieved.

Preliminary results reported here confirm that our fitting and modelling methods are able to reliably describe the spatial and temporal pattern of large scale variations in the ionosphere over the Australian region, both in terms of a more precise climatology and as a clear picture of the large scale hour to hour and day to day variations in response to solar and thermospheric forcing. ELOISE also provides enough data to allow us to describe and understand in detail both the broad spectrum of these large scale variations and their gradients, and the patterns and variation within isolated events. Fitting flexible and smooth surfaces to each of the extracted parameter data produces a complete description of the ionosphere over the Australian region, including vertical and horizontal gradients in the electron density that are constructed to be consistent with all the different sounder data.

By collecting a large number of overlapping sounder measurements with different classes of instruments, where each instrument was observing the same periods of ionospheric variability, ELOISE has created an opportunity

to understand each particular instrument's response to a wide variety of ionospheric disturbances and gradients on different scales. This understanding is critical to integrating ionospheric data from a diverse range of instruments into a single self-consistent model of the ionosphere and the HF propagation through it. ELOISE has demonstrated that the VIS and OIS data can be reconciled to explain the biggest part of the variation (the large scale variations) within an integrated model.

Analysis of how to best model the remaining small scale variability and separate N_e changes from propagation path changes is a subject of on-going work. ELOISE has nevertheless demonstrated that our equipment and processes can measure and observe the fine scale variability in the ionosphere.

Acknowledgments

The authors acknowledge the efforts of the whole team at High Frequency Radar Branch. A trial such as this requires a large team effort. Each component, each sensor, each remote facility, was built upon previous knowledge or other developments within the Branch. The authors also acknowledge the efforts of Lockheed Martin Australia (LMA) support staff that maintained and assisted in upgrading JORN sounders, and LMA and BAE Systems JORN site staff for assisting in site modifications and preparations.

One month of basic VIS trace & OIS ionogram feature and corresponding fitted parameter data from the ELOISE network is in the process of being made available on the Australian Bureau of Meteorology Space Weather Service web site http://www.sws.bom.gov.au/World_Data_Centre. One week of the underlying standard sounder images (and software readers) from a standard set of paths are also being made available at this same location.

Requests for access to additional HF data collected during ELOISE will be processed on a case-by-case basis pursuant to Defence Science and Technology Group policy for public release of information.

Airglow data collected during ELOISE is available from the DST Group Airglow database (Unewisse and Cool, 2017b) at https://www.dst.defence.gov.au/opportunity/dst-airglow-database.

Appendix A:. Equipment locations and OIS path lengths

(See Table 2 and 3).

Table 2
ELOISE sites and equipment.

Location	Lat, Lon in degrees	Sounder equipment	Ancillary equipment
Ajana	−27.94, 114.68	VIS_DPS	
Alice Springs Airport	−23.80, 133.90	VIS_PRIME, QVIS Tx, QVIS_Rx	GPS-TEC
Broome	−17.90, 122.26		Transponder
Carnarvon	−24.90, 113.71	VIS_DPS	
Cloncurry	−20.70, 140.51	OIS and CSF Tx	ADS-B Receiver
Coondambo	−31.06, 135.91	OIS_Rx, AoA_Rx	
Curtin	−17.62, 123.83	VIS_PRIME, OIS BR/TCI Tx	
Groote Eylandt	−13.85, 136.43	VIS_DPS	
Harts Range	−22.97, 134.45	QVIS_Tx, OIS and CSF Tx	ADS-B Receiver
Hermannsburg	−24.06, 132.74	QVIS_Rx	GPS-TEC
Horn Island	−10.59, 142.29		Transponder
Humpty Doo	−12.61, 131.27	OIS and CSF Tx	
Hughes	−12.68, 131.13		Transponder
Innisfail	−17.52, 146.03		Transponder
			ADS-B Receiver
Kalkarindji	−17.44, 130.83	VIS_DPS, OIS and CSF Tx	
Kowandi	−12.49, 131.05	OIS_Rx	
Kununurra	−15.78, 128.70		Transponder
Laverton	−28.33, 122.00	VIS_DPS, OIS_Rx, AoA_Rx, BSS	
Learmonth	−22.25, 114.09	VIS_DPS	
Longreach	−24.29, 143.19	VIS_DPS, OIS_Rx BSS	
Lynd River	−18.02, 144.87	VIS_DPS, OIS and CSF Tx	
Mount Everard	−23.52, 133.68	QVIS_Rx, OIS_Rx, BSS	GPS-TEC, TRACE
Mount Isa	−20.73, 139.48		Transponder
Nhulunbuy	−12.17, 136.78		Transponder
Normanton	−17.70, 141.07		Transponder
Scherger	−12.66, 142.09	VIS_DPS, OIS and CSF Tx	
South Hedland	−20.41, 118.58	VIS_DPS, OIS BR/TCI Tx	
Woodside	−34.97, 138.88	OIS_Tx	

Note: Rx refers to a Receiver system; Tx refers to a Transmitter system; VIS_DPS refers to the original Lowell Digital Portable Sounders in use within JORN; VIS_PRIME refers to the DST/HFR Branch PRIME VIS sounding system in use within JORN; TRACE refers to the HFR Branch optical all sky camera system.

Table 3

Standard OIS path names and ranges in km. Each row represents an OIS transmitter site and each column represents an OIS receiver site. Paths with ranges > 2000 km are unable to produce reliable information about one-hop E and Es.

	R1 Rx Longreach	R2 Rx Laverton	R3 Rx Mt Everard	Kowandi	Coondambo
South Hedland	2570	944	1597	1591	2095
Curtin	2145	1202	1218	959	1925
Kalkarindji	1493	1507	736	544	1593
Humpty Doo	1803	1989	1234	— 24 — (not used)	2097
Scherger	1293	2712	1495	1198	2136
Lyn River	715	2598	1314	1602	1706
Woodside	1256	1759	1366	2608	517
Cloncurry	488	2055	772	1350	1241
Harts Range	905	1383	—100—(not used)	1209	908
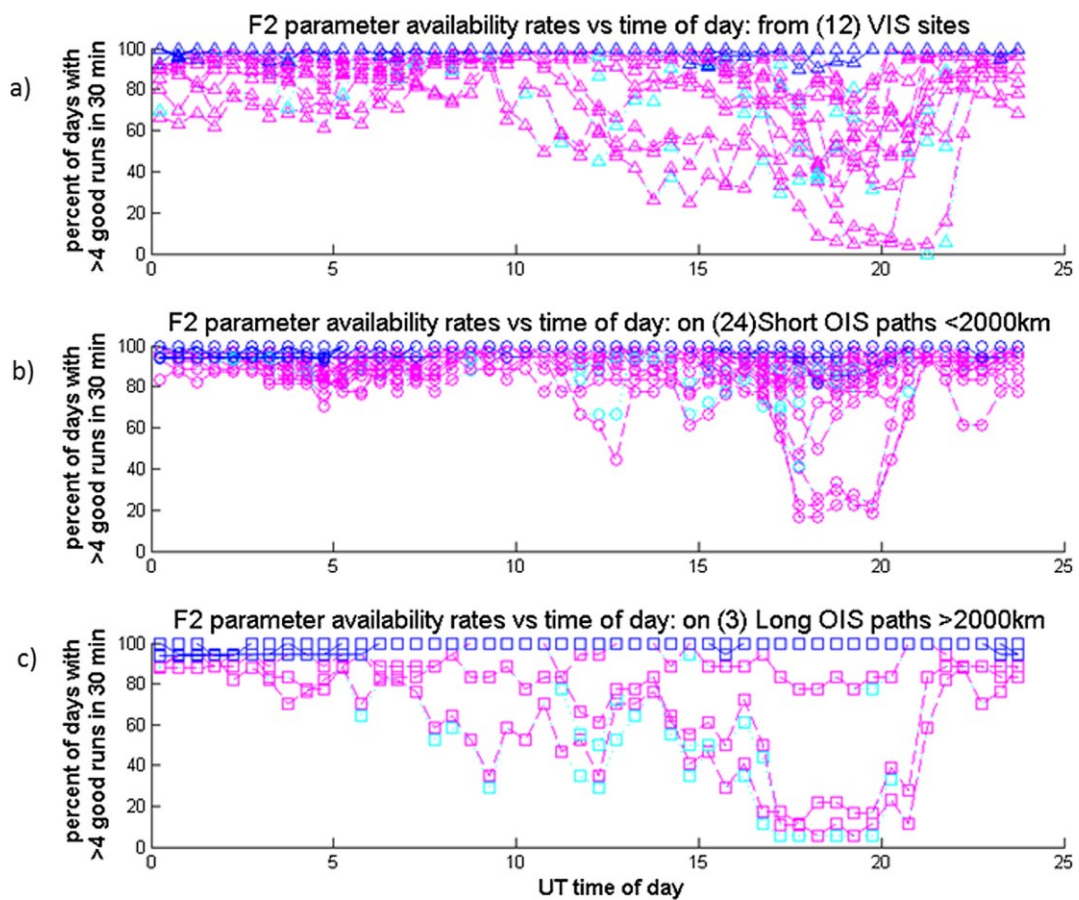
Fig. 21. Time series plots of the rates of the percentage availability of good sounder fit data versus time of day in UT (a) for all the VIS sounder sites, (b) for all the short OIS sounder paths, and (c) for the long OIS sounder paths. Symbols in blue represent the foF2 parameter. Symbols in magenta represent the availability rates for hmF2 parameters. Symbols in cyan are largely over-plotted by the symbols in magenta and represent good availability rates for h0F2 parameter data. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Appendix B:. Feature consistency and fitted parameter availability

Fig. 21 shows that all standard sites generate useful amounts of foF2 data 90–98% of the time. Most sites generate useful amounts of F2 height data most days.

Some sites have difficulty generating useful amounts of F2 height data over night. This is typically <20% of OIS sites but it mainly occurs in the long OIS paths. Just before dawn (18–21 UT) is a particularly challenging time to automatically generate F2 height data because for some paths the value of foF2 is very

low and the length of trace available for fitting is minimal.

References

Ayliffe, J.K., Durbridge, L.J., Frazer, G.J., Gardiner-Garden, R.S., Heitmann, A.J., Praschifka, J., Quinn, A.D., Scarman, G.C., Turley, M.D.E., 2019. The DST group high-fidelity, multichannel oblique incidence ionosonde. Radio Sci. 54 (1), 104–114. https://doi.org/10.1029/2018RS006681.

Ayub, M., Iqbal, S., Ameen, M.A., Reinisch, B.W., 2009. Study of maximum electron density NmF2 at Karachi and Islamabad during solar minimum (1996) and solar maximum (2000) and its comparison with IRI. Adv. Space Res. 43 (11), 1821–1824. https://doi.org/10.1016/j.asr.2008.09.034.

Barnes, R.I., Gardiner-Garden, R.S., Harris, T.J., 2000. Real time ionospheric models for the Australian Defence Force. In: Dyson, P., Norman, R. (Eds.), Workshop on Applications of Radio Science (WARS). La Trobe University, Melbourne, Australia, pp. 122–135.

Bilitza, D., Altadill, D., Truhlik, V., Shubin, V., Galkin, I., Reinisch, B., Huang, X., 2017. International Reference Ionospheric 2016: From ionospheric climate to real time weather predictions. Space Weather 15 (2), 418–429. https://doi.org/10.1002/2016SW001593.

Black, Q., Wood Jr, J.F., Sherrill, W.M., 1995. Mode angles of arrival in the 55 to 3500 km range. Radio Sci. 30 (3), 693–702. https://doi.org/10.1029/94RS03183.

Chen, J., Bennett, J.A., Dyson, P.L., 1992. Synthesis of oblique ionograms from vertical ionograms using a quasi-parabolic segment model of the ionosphere. J. Atmos. Terr. Phys. 54 (3), 323–331. https://doi.org/10.1016/0021-9169(92)90011-9.

Danilov, A.D., 2009. Time and Spatial variations of the foF2 (night)/foF2 (day) values. Adv. Space Res. 43 (11), 1786–1793. https://doi.org/10.1134/S0016793208020114.

Davies, K., 1990. Ionospheric Radio. Institution of Engineering and Technology, Stevenage, UK.

Earl, G.F., Ward, B.D., 1987. The frequency management system of the Jindalee over-the-horizon backscatter HF Radar. Radio Sci. 22 (2), 275–291. https://doi.org/10.1029/RS022i002p00275.

Fabrizio, G., 2013. High Frequency Over-the-Horizon Radar: Fundamental Principles, Signal Processing, and Practical Applications. McGraw Hill Professional, New York, NY, USA.

Galkin, I.A., Reinisch, B.W., 2008. The new ARTIST 5 for all Digisondes. Ionosonde Network Advisory Group Bull., Lowell, MA, USA, 69(8).

Gardiner-Garden, R.S., Heitmann, A., Northey, B.J., 2017. "A parameteric model of the Ionospheric electron denisty profile for JORN", DST-Group-Tech-Note-TN-1722. Department of Defence, Edinburgh, Australia, Retrieved from https://www.dst.defence.gov.au/publication/parametric-model-ionospheric-electron-density-profile-jorn.

Hajkowicz, L.A., 2007. Morphology of quantified ionospheric range spreadF over a wide range of mid-latitudes in the Australian longitude sector. Annales Geophys. 25 (5), 1125–1130. https://doi.org/10.5194/angeo-25-1125-2007.

Harris, T.J., Cervera, M.A., Meehan, D.H., 2012. SpICE: A program to study small–scale disturbances in the ionosphere. J. Geophys. Res. Space Phys. 117 (A6), A06321. https://doi.org/10.1029/2011JA017438.

Harris, T.J., Frazer, G.J., 2005. High Resolution Radar Measurements of Ionospheric Variations on an Oblique Path. In: Goodman, J.M. (Ed). 11th International Ionospheric Effects Symposium National Technical Information Services (NTIS), Alexandria, VA, USA, pp. 290–297.

Harris, T.J., Quinn, A.D., Pederick, L.H., 2016. The DST group ionospheric sounder replacement for JORN. Radio Sci. 51 (6), 563–572. https://doi.org/10.1002/2015RS005881.

Heitmann, A.J., Cervera, M.A., Gardiner-Garden, R.S., Holdsworth, D. A., MacKinnon, A.D., Reid, I.M., Ward, B.D., 2018. Observations and modelling of travelling ionospheric disturbance signature from an Australian Network of Oblique Angle-of-Arrival Sounders. Radio Sci. 53 (9), 1089–1107. https://doi.org/10.1029/2018RS006613.

Heitmann, A.J., Gardiner-Garden, R.S., 2019. A robust feature extraction and parameterized fitting algorithm for bottom-side oblique and vertical incidence ionograms. Radio Sci. 54 (1), 115–134. https://doi.org/10.1029/2018RS006682.

Hernandez-Pajares, M., Juan, J.M., Sanz, J., Aragón-Àngel, A., 2012. Propagation of medium scale traveling ionospheric disturbances at different latitudes and solar cycle conditions. Radio Sci. 47 (6), 1–22. https://doi.org/10.1029/2011RS004951.

Hocke, K., Schlegel, K., 1996. A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995. Annales Geophys. 14 (9), 917–940. https://doi.org/10.1007/s00585-996-0917-6.

Holdsworth, D.A., 2017. Skywave over-the-horizon radar track registration using earth surface and infrastructure backscatter. In: IEEE Radar Conference (RadarConf), edited by IEEE, pp.0986-0991, Seattle, WA, USA, 10.1109/RADAR.2017.7944347.

Lynn, K.J.W., 2008. Ionospheric observations made by a time-interleaved Doppler ionosonde. Adv. Space Res. 42 (7), 1218–1230. https://doi.org/10.1016/j.asr.2007.06.007.

Lynn, K.J.W., 2009. A technique for calculating ionospheric Doppler shifts from standard ionograms suitable of r scientific, HF Communications and OTH radar applications. Radio Sci. 44 (6), 1–11. https://doi.org/10.1029/2009RS004210.

Lynn, K.J.W., Otsuka, Y., Shiokawa, K., 2013. Ionogram-based range-time displays for observing relationships between ionosonde satellite traces, spread F and drifting optical plasma depletions. J. Atmos. Sol. Terr. Phys. 98, 105–112. https://doi.org/10.1016/j.jastp.2013.03.020.

Lynn, K.J.W., Gardiner-Garden, R.S., Heitmann, A., 2014. The spatial and temporal structure of twin peaks and a midday bite out in foF2 (with associated height changes) in the Australian and South Pacific low mid-latitude ionosphere. J. Geophys. Res. Space Phys. 119 (12), 10294–10304. https://doi.org/10.1002/2014JA020617.

Lynn, K.J.W., Gardiner-Garden, R.S., Heitmann, A., 2016. The observed compression and expansion of the F2 ionosphere as a major component of ionospheric variability. Radio Sci. 51 (5), 538–552. https://doi.org/10.1002/2016RS006036.

McNamara, L.F., 1991. The Ionosphere: Communications, Surveillance and Direction Finding. Krieger Publishing Company, Pennsylvania State University, Malabar, FL, USA.

McNamara, L.F., Baker, C.R., Decker, D.T., 2008. "Accuracy of USU-GAIM specification of foF2 and M(3000)F2 for a worldwide distribution of ionosonde locations",. Radio Sci. 43 (1), RS1011. https://doi.org/10.1029/2007RS003754.

McNamara, Leo F., Wilkinson, Phillip J., 2009. Spatial correlations of foF2 deviations and their implications for global ionospheric models: 1. Ionosondes in Australia and Papua New Guinea: SPATIAL CORRELATIONS: 1 n/a n/a. Radio Sci. 44 (2). https://doi.org/10.1029/2008RS003955.

Paznukhov, V.V., Galushko, V.G., Reinisch, B.W., 2012. Digisonde observations of TIDs with frequency and angular sounding technique. Adv. Space Res. 49 (4), 700–710. https://doi.org/10.1016/j.asr.2011.11.012.

Pederick, L.H., Cervera, M.A., Harris, T.J., 2017. Interpreting observations of large-scale traveling ionospheric disturbances by ionospheric sounders. J. Geophys. Res. Space Phys. 122 (12), 12556–12569. https://doi.org/10.1002/2017JA024337.

Reinisch, B.W., Haines, D.M., Bibl, K., Galkin, I., Huang, X., Kitrosser, D.F., Sales, G.S., Scali, J.L., 1997. Ionospheric sounding in support of over-the-horizon radar. Radio Sci. 32 (4), 1681–1694. https://doi.org/10.1029/97RS00841.

Rishbeth, H., Muller-Wodarg, I.C.F., Zou, I., Fuller-Rowell, T.J., Millward, G.H., Moffett, R.J., Idenden, D.W., Aylward, A.D., 2000. Annual and semi-annual variations in the ionospheric F2 layer: II: Physical Discussion. Annales Geophys. 18 (8), 945–956. https://doi.org/10.1007/s00585-000-0945-6.

Rishbeth, H., Mendillo, M., Wroten, J., Roble, R.G., 2009. Day-by-day modelling of the ionospheric F2 layer for the year 2002. J. Atmos. Sol. Terr. Phys. 71 (8), 848–856. https://doi.org/10.1016/j.jastp.2009.03.022.

Sojka, J.J., Thompson, D.C., Scherliess, L., Schunk, R.W., Harris, T.J., 2007. Assessing models for ionospheric weather specifications over Australia during the 2004 Climate and Weather of the Sun-Earth-System (CAWSES) campaign. J. Geophys. Res. Space Phys. 112 (A9), A09306. https://doi.org/10.1029/2006JA012048.

Sojka, J.J., Thomspon, D.C., Schunk, R.W., Bullett, T.W., Makela, J.J., 2001. Assimilation Ionosphere Model: development and testing with Combined Ionospheric Campaign Caribbean measurements. Radio Sci. 36 (2), 247–259. https://doi.org/10.1029/1999RS002411.

Turley, M.D.E., Heitmann, A.J., Gardiner-Garden, R.S., 2019. Ionogram RFI rejection using an auto-regressive interpolation process. Radio Sci. 54 (1), 135–150. https://doi.org/10.1029/2018RS006683.

Unewisse, A., Cervera, M., Pederick, L., Harris, T. J., Cool, A., 2017. Airglow Observations from ELOISE. In: Short, W., Caprarelli, G. (Eds.), Proceedings of the 16th Australian Space Science Conference, National Space Society of Australia Ltd, Melbourne, Australia, pp. 139–150.

Unewisse, A., Cool, A., 2017a. Imagers for Ionospheric Airglow Observations. In: Short, W., Caprarelli, G. (Eds.), Proceedings of the 16th Australian Space Science Conference, National Space Society of Australia Ltd, Melbourne, Australia, pp. 127–138.

Unewisse, A., Cool, A., 2017b. DST Airglow Database. Retrieved from, DST Group, https://www.dst.defence.gov.au/opportunity/dst-airglow-database.

Valladares, C.E., Hei, M.A., 2012. Measurements of the characteristics of TIDs using Small Regional Networks of GPS Receivers during the Campaign of 17–30 July of 2008. Int. J. Geophys., 1–14 https://doi.org/10.1155/2012/548784.

References

Abdu, M. A., Alam Kherani, E., Batista, I. S., de Paula, E. R., Fritts, D. C., & Sobral, J. H. A. (2009), 'Gravity wave initiation of equatorial spread F/plasma bubble irregularities based on observational data from the SpreadFEx campaign', *Annales Geophysicae*, 27(7), pp. 2,607–2,622. doi:10.5194/angeo-27-2607-2009

Abramovich, Y. I., Frazer, G. J., & Johnson, B. A. (2013), 'Principles of Mode-Selective MIMO OTHR', *IEEE Transactions on Aerospace and Electronic Systems*, 49(3), pp. 1,839–1,868. doi:10.1109/TAES.2013.6558024

Adams, G. W., Brosnahan, J. W., Walden, D. C., & Nerney, S. F. (1986), 'Mesospheric Observations Using a 2.66-MHz Radar as an Imaging Doppler Interferometer: Description and First Results', *Journal of Geophysical Research*, 91(A2), pp. 1,671–1,683. doi:10.1029/JA091iA02p01671

Afraimovich, E. L., Perevalova, N. P., & Voyeikov, S. V. (2003), 'Traveling wave packets of total electron content disturbances as deduced from global GPS network data', *Journal of Atmospheric and Solar-Terrestrial Physics*, 65(11–13), pp. 1,245–1,262. doi:10.1016/j.jastp.2003.08.007

Akchurin, A. D., Bochkarev, V. V., Ryabchenko, E. Y., & Sherstyukov, O. N. (2009), 'Improved precision of virtual height measurements with coherent radio pulse sounding based on the maximum likelihood method', *Advances in Space Research*, 43(11), pp. 1,595–1,602. doi:10.1016/j.asr.2008.03.033

Allison, J., Caddy, J., & Yip, P. (2019), 'Jindalee Operational Radar Network Phase 6: Over-the-horizon radar developments', *IEEE Potentials*, 38(4), pp. 28–33. doi:10.1109/MPOT.2019.2912436

Altshuler, E. E. (1998), 'Tropospheric Range-Error Corrections for the Global Positioning System', *IEEE Transactions on Antennas and Propagation*, 46(5), pp. 643–649. doi:10.1109/8.668906

Anderson, S. J., & Abramovich, Y. I. (1998), 'A unified approach to detection, classification, and correction of ionospheric distortion in HF sky wave radar systems', *Radio Science*, 33(4), pp. 1,055–1,067. doi:10.1029/98RS00877

Anderson, S. J., & Lees, M. L. (1988), 'High-resolution synoptic scale measurement of ionospheric motions with the Jindalee sky wave radar', *Radio Science*, 23(3), pp. 265–272. doi:10.1029/RS023i003p00265

Angling, M. J., & Cannon, P. S. (2004), 'Assimilation of radio occultation measurements into background ionospheric models', *Radio Science*, 39(1), RS1S08. doi:10.1029/2002RS002819

Angling, M. J., Cannon, P. S., Davies, N. C., Willink, T. J., Jodalen, V., & Lundborg, B. (1998), 'Measurements of Doppler and multipath spread on oblique high-latitude HF paths and their use in characterizing data modem performance', *Radio Science*, 33(1), pp. 97–107. doi:10.1029/97RS02206

Angling, M. J., & Jackson-Booth, N. K. (2011), 'A short note on the assimilation of collocated and concurrent GPS and ionosonde data into the Electron Density Assimilative Model', *Radio Science*, 46(6), RS0D13. doi:10.1029/2010RS004566

Angling, M. J., & Khattatov, B. (2006), 'Comparative study of two assimilative models of the ionosphere', *Radio Science*, 41(5), RS5S20. doi:10.1029/2005RS003372

Appleton, E. V., & Barnett, M. A. F. (1925), 'Local Reflection of Wireless Waves from the Upper Atmosphere', *Nature*, 115(2,888), pp. 333–334. doi:10.1038/115333a0

Araujo-Pradere, E. A., Fuller-Rowell, T. J., & Codrescu, M. V. (2005), 'Characteristics of the ionospheric variability as a function of season, latitude, local time, and geomagnetic activity', *Radio Science*, 40(5), RS5009. doi:10.1029/2004RS003179

Arthur, P. C., Lissimore, M., Cannon, P. S., & Davies, N. C. (1997), 'Application of a high quality ionosonde to ionospheric research', *Proceedings of the 7th International Conference on HF Radio Systems and Techniques*, IEE Conference Publication No. 441, pp. 135–139. doi:10.1049/cp:19970776

Ayliffe, J. K., Durbridge, L. J., Frazer, G. J., Gardiner-Garden, R. S., Heitmann, A. J., Praschifka, J., et al. (2019), 'The DST Group high-fidelity, multi-channel oblique incidence ionosonde', *Radio Science*, 54(1), pp. 104–114. doi:10.1029/2018RS006681

Azzarone, A., Bianchi, C., Pezzopane, M., Pietrella, M., Scotto, C., & Settimi, A. (2012), 'IONORT: A Windows software tool to calculate the HF ray tracing in the ionosphere', *Computers & Geosciences*, 42, pp. 57–63. doi:10.1016/j.cageo.2012.02.008

Balan, N., Liu, L. B., & Le, H. J. (2018), 'A brief review of equatorial ionization anomaly and ionospheric irregularities', *Earth and Planetary Physics*, 2(4), pp. 257–275. doi:10.26464/epp2018025

Balan, N., Otsuka, Y., Fukao, S., Abdu, M. A., & Bailey, G. J. (2000), 'Annual variations of the ionosphere: A review based on MU radar observations', 25(1), pp. 153–162. doi:10.1016/S0273-1177(99)00913-8

Ballard, K. A. (1981), *An experimental study of the propagation and dispersion of internal atmospheric gravity waves*, PhD thesis, Dartmouth College, Hanover, New Hampshire.

Barnes, R. I., Earl, G. F., Papazoglou, M., Burchett, L., & Terzuoli, A. J., Jr. (2010), 'The Instagram: A Novel Sounding Technique for Enhanced HF Propagation Advice', *Proceedings of the IEEE Radar Conference 2010*, Washington, DC, pp. 1,446–1,449. doi:10.1109/RADAR.2010.5494387

Barnes, R. I., Gardiner-Garden, R. S., & Harris, T. J. (2000), 'Real time ionospheric models for the Australian Defence Force', *Proceedings of the Workshop on Applications of Radio Science (WARS'00)*, Beechworth, Victoria, Australia, pp. 122–135. Retrieved from http://www.sws.bom.gov.au/IPSHosted/NCRS/wars/wars2000/commg/barnes.pdf

Barry, G. (1963), 'Ray tracings of acoustic waves in the upper atmosphere', *Journal of Atmospheric and Terrestrial Physics*, 25(11), pp. 621–629. doi:10.1016/0021-9169(63)90156-9

Basler, R. P., Price, G. H., Tsunoda, R. T., & Wong, T. L. (1988), 'Ionospheric distortion of HF signals', *Radio Science*, 23(4), pp. 569–579. doi:10.1029/RS023i004p00569

Baulch, R. N. E., Butcher, E. C., Devlin, J. C., & Hammer, P. R. (1984), 'A simple sounder to measure the properties of ionospherically reflected radio waves', *Journal of Atmospheric and Terrestrial Physics*, 46(10), pp. 895–903. doi:10.1016/0021-9169(84)90029-1

Bazin, V., Molinie, J. P., Munoz, J., Dorey, P., Saillant, S., Auffray, G., Rannou, V., & Lesturgie, M. (2006), 'A general presentation about the OTH-Radar NOSTRADAMUS', *Proceedings of the IEEE Radar Conference 2006*, Verona, New York, pp. 1,446–1,449. doi:10.1109/RADAR.2006.1631867

Bean, B. R., & Thayer, G. D. (1959a), 'Models of the Atmospheric Radio Refractive Index', *Proceedings of the IRE*, 47(5), pp. 740–753. doi:10.1109/JRPROC.1959.287242

Bean, B. R., & Thayer, G. D. (1959b), *CRPL Exponential Reference Atmosphere*, NBS Monograph 4, National Bureau of Standards, U.S. Department of Commerce, Washington, DC.

Beley, V. S., Galushko, V. G., & Yampolski, Y. M. (1995), 'Traveling ionospheric disturbance diagnostics using HF signal trajectory parameter variations', *Radio Science*, 30(6), pp. 1,739–1,752. doi:10.1029/95RS01992

Bell, K. L., Ephraim, Y., & Van Trees, H. L. (2000), 'A Bayesian Approach to Robust Adaptive Beamforming', *IEEE Transactions on Signal Processing*, 48(2), pp. 386–398. doi:10.1109/78.823966

Bennett, J. A. (1967), 'The relation between group path and phase path', *Journal of Atmospheric and Terrestrial Physics*, 29(7), pp. 893–896. doi:10.1016/0021-9169(67)90056-6

Bennett, J. A. (1968), 'The ray theory of Doppler frequency shifts', *Australian Journal of Physics*, 21(3), pp. 259–272. doi:10.1071/PH680259

Bennett, J. A., Chen, J., & Dyson, P. L. (1991), 'Analytic ray tracing for the study of HF magneto-ionic radio propagation in the ionosphere', *Applied Computational Electromagnetics Society Journal*, 6(1), pp. 192–210.

Bennett, J. A., & Dyson, P. L. (1993), 'On Pfister's theorem and its generalization', *Journal of Atmospheric and Terrestrial Physics*, 55(9), pp. 1,307–1,310. doi:10.1016/0021-9169(93)90055-4

Bennett, J. A., & Dyson, P. L. (2002), 'Effect of horizontal ionospheric structure on HF radio systems', *Proceedings of the 27th URSI General Assembly*, Maastricht, Netherlands. Retrieved from http://www.ursi.org/proceedings/procGA02/papers/p1337.pdf

Bennett, J. A., Dyson, P. L., & Norman, R. J. (2004), 'Progress in Radio Ray Tracing in the Ionosphere', *The Radio Science Bulletin*, 310, pp. 81–91. doi:10.23919/URSIRSB.2004.7909439

Bibl, K. (1998), 'Evolution of the ionosonde', *Annals of Geophysics*, 41(5–6), pp. 667–680. doi:10.4401/ag-3810

Bilitza, D. (2001), 'International Reference Ionosphere 2000', *Radio Science*, 36(2), pp. 261–275. doi:10.1029/2000RS002432

Bilitza, D., Altadill, D., Truhlik, V., Shubin, V., Galkin, I., Reinisch, B., & Huang, X. (2017), 'International Reference Ionosphere 2016: From ionospheric climate to real-time weather predictions', *Space Weather*, 15(2), pp. 418–429. doi:10.1002/2016SW001593

Bilitza, D., Altadill, D., Zhang, Y., Mertens, C., Truhlik, V., Richards, P., McKinnell, L.-A., & Reinisch, B. (2014), 'The International Reference Ionosphere 2012 – a model of international collaboration', *Journal of Space Weather and Space Climate*, 4, A07. doi:10.1051/swsc/2014004

Bilitza, D., & Reinisch, B. W. (2008), 'International Reference Ionosphere 2007: Improvements and new parameters', *Advances in Space Research*, 42(4), pp. 599–609. doi:10.1016/j.asr.2007.07.048

Black, Q. R., Wood, J. F., Jr., & Sherrill, W. M. (1995), 'Mode angles of arrival in the 55- to 3500-km range', *Radio Science*, 30(3), pp. 693–702. doi:10.1029/94RS03183

Black, Q. R., Wood, J. F., Jr., Sonsteby, A. G., & Sherrill, W. M. (1993), 'A direction finding ionosonde for ionospheric propagation research', *Radio Science*, 28(5), pp. 795–809. doi:10.1029/93RS00690

Boldovskaya, I. G. (1982), 'Doppler frequency shifts of radio waves reflected by parabolic and quasi-parabolic ionospheric layers', *Journal of Atmospheric and Terrestrial Physics*, 44(4), pp. 305–311. doi:10.1016/0021-9169(82)90075-7

Booker, H. G. (1979), 'The role of acoustic gravity waves in the generation of spread-*F* and ionospheric scintillation', *Journal of Atmospheric and Terrestrial Physics*, 41(5), pp. 501–515. doi:10.1016/0021-9169(79)90074-6

Bottone, S. (1993), *Acoustic-Gravity Waves from Low-Altitude Localized Disturbances*, Technical Report DNA-TR-92-140, Defense Nuclear Agency, Alexandria, Virginia. Retrieved from https://apps.dtic.mil/dtic/tr/fulltext/u2/a264804.pdf

Bowman, G. G. (1960), 'Further studies of "spread-F" at Brisbane – I. Experimental', *Planetary and Space Science*, 2(2–3), pp. 133–149. doi:10.1016/0032-0633(60)90008-8

Bowman, G. G. (1964), 'Spread-*F* in the Ionosphere and the Neutral Particle Density of the Upper Atmosphere', *Nature*, 201(4,919), pp. 564–566. doi:10.1038/201564b0

Bowman, G. G. (1968), 'Movements of ionospheric irregularities and gravity waves', *Journal of Atmospheric and Terrestrial Physics*, 30(5), pp. 721–734. doi:10.1016/S0021-9169(68)80028-5

Bowman, G. G. (1981), 'The nature of ionospheric spread-*F* irregularities in mid-latitude regions', *Journal of Atmospheric and Terrestrial Physics*, 43(1), pp. 65–79. doi:10.1016/0021-9169(81)90010-6

Bowman, G. G. (1990), 'A Review of Some Recent Work on Mid-Latitude Spread-*F* Occurrence as Detected by Ionosondes', *Journal of Geomagnetism and Geoelectricity*, 42(2), pp. 109–138. doi:10.5636/jgg.42.109

Bowman, G. G. (1991a), 'Ionospheric frequency spread and its relationship with range spread in mid-latitude regions', *Journal of Geophysical Research*, 96(A6), pp. 9,745–9,753. doi:10.1029/91JA00389

Bowman, G. G. (1991b), 'Nighttime Mid-Latitude Travelling Ionospheric Disturbances Associated with Mild Spread-F Conditions', *Journal of Geomagnetism and Geoelectricity*, 43(11), pp. 899–920. doi:10.5636/jgg.43.899

Bowman, G. G. (1992), 'Some aspects of large-scale travelling ionospheric disturbances', *Planetary and Space Science*, 40(6), pp. 829–845. doi:10.1016/0032-0633(92)90110-A

Bowman, G. G. (1995), 'Multiplicity of travelling disturbances in the nighttime mid-latitude F_2-region ionosphere', *Indian Journal of Radio & Space Physics*, 24(2), pp. 91–96. Retrieved from http://nopr.niscair.res.in/handle/123456789/35790

Bowman, G. G., Clarke, R. H., & Meehan, D. H. (1988), 'Mid-latitude frequency spread and its association with small scale ionospheric stratifications', *Journal of Atmospheric and Terrestrial Physics*, 50(9), pp. 797–809. doi:10.1016/0021-9169(88)90102-X

Bowman, G. G., & Hajkowicz, L. A. (1991), 'Small-scale ionospheric structures associated with mid-latitude spread-*F*', *Journal of Atmospheric and Terrestrial Physics*, 53(5), pp. 447–457. doi:10.1016/0021-9169(91)90038-9

Bowman, G. G., & Monro, P. E. (1988), 'Mid-latitude range spread and travelling ionospheric disturbances', *Journal of Atmospheric and Terrestrial Physics*, 50(3), pp. 215–223. doi:10.1016/0021-9169(88)90070-0

Brace, L. H. (1998), 'Langmuir Probe Measurements in the Ionosphere', in R. F. Pfaff, J. E. Borovsky, & D. T. Young (ed.), *Measurement Techniques in Space Plasmas: Particles*, Geophysical Monograph Series, vol. 102, American Geophysical Union, Washington, DC, pp. 23–35. doi:10.1029/GM102p0023

Bramley, E. N. (1953), 'Direction-finding studies of large-scale ionospheric irregularities', *Proceedings of the Royal Society of London A: Mathematical and Physical Sciences*, 220(1,140), pp. 39–61. doi:10.1098/rspa.1953.0170

Breit, G., & Tuve, M. A. (1925), 'A Radio Method of Estimating the Height of the Conducting Layer', *Nature*, 116(2,914), p. 357. doi:10.1038/116357a0

Breit, G., & Tuve, M. A. (1926), 'A test of the existence of the conducting layer', *Physical Review*, 28(3), pp. 554–575. doi:10.1103/PhysRev.28.554

Briggs, B. H. (1964), 'Observations of radio star scintillations and spread-*F* echoes over a solar cycle', *Journal of Atmospheric and Terrestrial Physics*, 26(1), pp. 1–23. doi:10.1016/0021-9169(64)90104-7

Briggs, B. H. (1980), 'Radar observations of atmospheric winds and turbulence: a comparison of techniques', *Journal of Atmospheric and Terrestrial Physics*, 42(9–10), pp. 823–833. doi:10.1016/0021-9169(80)90086-0

Briggs, B. H. (1984), 'The Analysis of Spaced Sensor Records by Correlation Techniques', *Handbook for Middle Atmosphere Program (MAP)*, 13, pp. 166–186. Retrieved from https://ntrs.nasa.gov/search.jsp?R=19850009143

Briggs, B. H., Phillips, G. J., & Shinn, D. H. (1950), 'The analysis of observations on spaced receivers of the fading of radio signals', *Proceedings of the Physical Society, Section B*, 63(2), pp. 106–121. doi:10.1088/0370-1301/63/2/305

Briggs, B. H., & Vincent, R. A. (1992), 'Spaced-antenna analysis in the frequency domain', *Radio Science*, 27(2), pp. 117–129. doi:10.1029/91RS03051

Brownlie, G. D., Dryburgh, L. G., & Whitehead, J. D. (1973), 'Measurement of the velocity of waves in the ionosphere: A comparison of the ray theory approach and diffraction theory', *Journal of Atmospheric and Terrestrial Physics*, 35(12), pp. 2,147–2,162. doi:10.1016/0021-9169(73)90133-5

Budden, K. G. (1985), *The propagation of radio waves: The theory of radio waves of low power in the ionosphere and magnetosphere*, Cambridge University Press, Cambridge, UK. doi:10.1017/CBO9780511564321

Budden, K. G., & Terry, P. D. (1971), 'Radio ray tracing in complex space', *Proceedings of the Royal Society of London A: Mathematical and Physical Sciences*, 321(1,546), pp. 275–301. doi:10.1098/rspa.1971.0033

Bullett, T. W. (1994), *Mid-latitude ionospheric plasma drift: A comparison of digital ionosonde and incoherent scatter radar measurements at Millstone Hill*, DEng thesis, University of Massachusetts, Lowell, Massachusetts.

Burke, G. J. (2011), *Numerical Electromagnetics Code – NEC-4.2, Method of Moments, Part I: User's Manual*, LLNL-SM-490875, Lawrence Livermore National Laboratory, Livermore, California.

Burtnyk, N., & Wolfe, J. L. (1967), 'Direction-finding performance of a wide-aperture sampling linear array', *Proceedings of the Institution of Electrical Engineers*, 114(5), pp. 581–588. doi:10.1049/piee.1967.0119

Bust, G. S., & Mitchell, C. N. (2008), 'History, current state, and future directions of ionospheric imaging', *Reviews of Geophysics*, 46(1), RG1003. doi:10.1029/2006RG000212

Cameron, A. (1995), 'The Jindalee Operational Radar Network: Its Architecture and Surveillance Capability', *Proceedings of the IEEE International Radar Conference 1995*, Alexandria, Virginia, pp. 692–697. doi:10.1109/RADAR.1995.522633

Cannon, P. S., Reinisch, B. W, Buchau, J., & Bullett, T. W. (1991), 'Response of the Polar Cap *F* Region Convection Direction to Changes in the Interplanetary Magnetic Field: Digisonde Measurements in Northern Greenland', *Journal of Geophysical Research*, 96(A2), pp. 1,239–1,250. doi:10.1029/90JA02128

Capon, J. (1969), 'High-Resolution Frequency-Wavenumber Spectrum Analysis', *Proceedings of the IEEE*, 57(8), pp. 1,408–1,418. doi:10.1109/PROC.1969.7278

Caton, R. G., Colman, J. J., Parris, R. T., Nickisch, L. J., & Bullock, G. (2017), 'HF propagation sensitivity study and system performance analysis with the Air Force Coverage Analysis Program (AFCAP)', paper presented at the *AGU Fall Meeting*, New Orleans, Louisiana.

Cervera, M. A. (2019), *PHaRLAP – Provision of High-frequency Raytracing Laboratory for Propagation studies*, DST Group MATLAB® Toolbox, Intelligence, Surveillance and Space Division, Defence Science and Technology Group, Edinburgh, South Australia, Australia. Retrieved from https://www.dst.defence.gov.au/opportunity/pharlap-provision-high-frequency-raytracing-laboratory-propagation-studies

Cervera, M. A., & Harris, T. J. (2014), 'Modelling ionospheric disturbance features in quasi-vertically incident ionograms using 3D magnetoionic ray tracing and atmospheric gravity waves', *Journal of Geophysical Research: Space Physics*, 119(1), pp. 431–440. doi:10.1002/2013JA019247

Cervera, M. A., Harris, T. J., Holdsworth, D. A., & Netherway, D. J. (in press), 'Ionospheric Effects on HF Radio Wave Propagation', in C. Huang and G. Lu (ed.), *Advances in Ionospheric Research: Current Understanding and Challenges*, Geophysical Monograph Series: Space Physics and Aeronomy, vol. 3, John Wiley & Sons Inc., Hoboken, New Jersey.

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), pp. 26–45. doi:10.1088/0959-5309/43/1/305

Chapman, S. (1950), 'Upper Atmospheric Nomenclature', *Bulletin of the American Meteorological Society*, 31(8), pp. 288–290. doi:10.1175/1520-0477-31.8.288

Chau, J. L., Aponte, N. A., Cabassa, E., Sulzer, M. P., Goncharenko, L. P., & González, S. A. (2010), 'Quiet time ionospheric variability over Arecibo during sudden stratospheric warming events', *Journal of Geophysical Research*, 115(A9), A00G06. doi:10.1029/2010JA015378

Chen, J., Bennett, J. A., & Dyson, P. L. (1990), 'Automatic fitting of quasi-parabolic segments to ionospheric profiles with application to ground range estimation for single-station location', *Journal of Atmospheric and Terrestrial Physics*, 52(4), pp. 277–288. doi:10.1016/0021-9169(90)90095-5

Chen, J., Bennett, J. A., & Dyson, P. L. (1992), 'Synthesis of oblique ionograms from vertical ionograms using quasi-parabolic segment models of the ionosphere', *Journal of Atmospheric and Terrestrial Physics*, 54(3–4), pp. 323–331. doi:10.1016/0021-9169(92)90011-9

Chisham, G., Lester, M., Milan, S. E., Freeman, M. P., Bristow, W. A., Grocott, A., et al. (2007), 'A decade of the Super Dual Auroral Radar Network (SuperDARN): scientific achievements, new techniques and future directions', *Surveys in Geophysics*, 28(1), pp. 33–109. doi:10.1007/s10712-007-9017-8

Cho, H. R., & Yeh, K. C. (1970), 'Neutral winds and the behavior of the ionospheric F_2 region', *Radio Science*, 5(6), pp. 881–894. doi:10.1029/RS005i006p00881

Chum, J., Athieno, R., Baše, J., Burešová, D., Hruška, F., Laštovička, J., McKinnell, L. A., & Šindelářová, T. (2012), 'Statistical investigation of horizontal propagation of gravity waves in the ionosphere over Europe and South Africa', *Journal of Geophysical Research*, 117(A3), A03312. doi:10.1029/2011JA017161

Chum, J., Bonomi, F. A. M., Fišer, J., Cabrera, M. A., Ezquer, R. G., Burešová, D., et al. (2014), 'Propagation of gravity waves and spread F in the low-latitude ionosphere over Tucumán, Argentina, by continuous Doppler sounding: First results', *Journal of Geophysical Research: Space Physics*, 119(8), pp. 6,954–6,965. doi:10.1002/2014JA020184

Chum, J., Laštovička, J., Šindelářová, T., Burešová, D., & Hruška, F. (2008), 'Peculiar transient phenomena observed by HF Doppler sounding on infrasound time scales', *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(6), pp. 866–878. doi:10.1016/j.jastp.2007.06.013

Chum, J., Šindelářová, T., Laštovička, J., Hruška, F., Burešová, D., & Baše, J. (2010), 'Horizontal velocities and propagation directions of gravity waves in the ionosphere over the Czech Republic', *Journal of Geophysical Research*, 115(A11), A11322. doi:10.1029/2010JA015821

Clark, R. M. (1970), *Gravity waves in the ionosphere*, PhD thesis, University of Illinois, Urbana, Illinois. Retrieved from ProQuest Dissertations & Theses Global (document ID 302519386).

Clark, R. M., Yeh, K. C., & Liu, C. H. (1971), 'Interaction of internal gravity waves with the ionospheric *F2*-layer', *Journal of Atmospheric and Terrestrial Physics*, 33(10), pp. 1,567–1,576. doi:10.1016/0021-9169(71)90074-2

Clarke, R. H. (1971), Preliminary observations with a directional ionosonde, *Planetary and Space Science*, 19(10), pp. 1,387–1,390. doi:10.1016/0032-0633(71)90191-7

Clarke, R. H. (1972), *Ground-based radio observations of midlatitude ionospheric irregularities*, PhD thesis, University of Queensland, St Lucia, Queensland, Australia.

Clarke, R. H., & Tibble, D. V. (1978), 'Measurement of the elevation angles of arrival of multicomponent h.f. skywaves', *Proceedings of the Institution of Electrical Engineers*, 125(1), pp. 17–24. doi:10.1049/piee.1978.0005

Colegrove, S. B. (2000), 'Project Jindalee: from bare bones to operational OTHR', *Proceedings of the IEEE International Radar Conference 2000*, Alexandria, Virginia, pp. 825–830. doi:10.1109/RADAR.2000.851942

Coleman, C. J. (1996), 'A model of HF sky wave radar clutter', *Radio Science*, 31(4), pp. 869–875. doi:10.1029/96RS00721

Coleman, C. J. (1997), 'On the simulation of backscatter ionograms', *Journal of Atmospheric and Solar-Terrestrial Physics*, 59(16), pp. 2,089–2,099. doi:10.1016/S1364-6826(97)00038-2

Coleman, C. J. (1998), 'A ray tracing formulation and its application to some problems in over-the-horizon radar', *Radio Science*, 33(4), pp. 1,187–1,197. doi:10.1029/98RS01523

Cowling, D. H., Webb, H. D., & Yeh, K. C. (1971), 'Group Rays of Internal Gravity Waves in a Wind-Stratified Atmosphere', *Journal of Geophysical Research*, 76(1), pp. 213–220. doi:10.1029/JA076i001p00213

Cox, C. R. (1947), 'Mutual Impedance Between Vertical Antennas of Unequal Heights', *Proceedings of the IRE*, 35(11), pp. 1,367–1,370. doi:10.1109/JRPROC.1947.233903

Croft, T. A. (1972), 'Sky-Wave Backscatter: A Means for Observing Our Environment at Great Distances', *Reviews of Geophysics and Space Physics*, 10(1), pp. 73–155. doi:10.1029/RG010i001p00073

Croft, T. A., & Hoogasian, H. (1968), 'Exact Ray Calculations in a Quasi-Parabolic Ionosphere With No Magnetic Field', *Radio Science*, 3(1), pp. 69–74. doi:10.1002/rds19683169

Crowley, G., & Rodrigues, F. S. (2012), 'Characteristics of traveling ionospheric disturbances observed by the TIDDBIT sounder', *Radio Science*, 47(4), RS0L22. doi:10.1029/2011RS004959

Dandenault, P. B., & Richards, P. G. (2015), 'The collapse of the midnight ionosphere and behavior of meridional neutral winds at Townsville over a full solar cycle', *Journal of Geophysical Research: Space Physics*, 120(11), pp. 9,826–9,838. doi:10.1002/2015JA021842

Dao, E. V., McNamara, L. F., & Colman, J. J. (2016), 'Magnetic field effects on the accuracy of ionospheric mirror models for geolocation', *Radio Science*, 51(4), pp. 284–300. doi:10.1002/2015RS005884

Davies, J. A., Lester, M., & Robinson, T. R. (1997), 'Deriving the normalised ion-neutral collision frequency from EISCAT observations', *Annales Geophysicae*, 15(12), pp. 1,557–1,569. doi:10.1007/s00585-997-1557-1

Davies, K. (1990), *Ionospheric Radio*, IEEE Electromagnetic Waves Series, vol. 31, Peter Peregrinus Ltd, London, UK. doi:10.1049/PBEW031E

Davies, K., & Baker, D. M. (1966), 'On Frequency Variations of Ionospherically Propagated HF Radio Signals', *Radio Science*, 1(5), pp. 545–556. doi:10.1002/rds196615545

Davies, K., & Jones, J. E. (1971), 'Three-dimensional observations of traveling ionospheric disturbances', *Journal of Atmospheric and Terrestrial Physics*, 33(1), pp. 39–46. doi:10.1016/0021-9169(71)90180-2

Davies, K., Watts, J. M., & Zacharisen, D. H. (1962), 'A Study of F_2-Layer Effects as Observed with a Doppler Technique', *Journal of Geophysical Research*, 67(2), pp. 601–609. doi:10.1029/JZ067i002p00601

Davis, M. J. (1973), 'The integrated ionospheric response to internal atmospheric gravity waves', *Journal of Atmospheric and Terrestrial Physics*, 35(5), pp. 929–959. doi:10.1016/0021-9169(73)90074-3

Dennis, J. E., Jr., & Schnabel, R. B. (1996), *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, SIAM Classics in Applied Mathematics, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania.

Ding, F., Wan, W., Xu, G., Yu, T., Yang, G., & Wang, J.-S. (2011), 'Climatology of medium-scale traveling ionospheric disturbances observed by a GPS network in central China', *Journal of Geophysical Research*, 116(A9), A09327. doi:10.1029/2011JA016545

Djuth, F. T., Sulzer, M. P., Elder, J. H., & Wickwar, V. B. (1997), 'High-resolution studies of atmosphere-ionosphere coupling at Arecibo Observatory, Puerto Rico', *Radio Science*, 32(6), pp. 2,321–2,344. doi:10.1029/97RS02797

Djuth, F. T., Sulzer, M. P., Gonzáles, S. A., Mathews, J. D., Elder, J. H., & Walterscheid, R. L. (2004), 'A continuum of gravity waves in the Arecibo thermosphere?', *Geophysical Research Letters*, 31(16), L16801. doi:10.1029/2003GL019376

Doerry, A. W. (2013), *Earth Curvature and Atmospheric Refraction Effects on Radar Signal Propagation*, SAND2012-10690, Sandia National Laboratories, Albuquerque, New Mexico. doi:10.2172/1088060

Drob, D. P., Emmert, J. T., Crowley, G., Picone, J. M., Shepherd, G. G., Skinner, W., et al. (2008), 'An empirical model of the Earth's horizontal wind fields: HWM07', *Journal of Geophysical Research*, 113(A12), A12304. doi:10.1029/2008JA013668

Drob, D. P., Emmert, J. T., Meriwether, J. W., Makela, J. J., Doornbos, E., Conde, M., et al. (2015), 'An update to the Horizontal Wind Model (HWM): The quiet time thermosphere', *Earth and Space Science*, 2(7), pp. 301–319. doi:10.1002/2014EA000089

DuCharme, E. D., Petrie, L. E., & Eyfrig, R. (1971), 'A method for predicting the F1 layer critical frequency', *Radio Science*, 6(3), pp. 369–378. doi:10.1029/RS006i003p00369

DuCharme, E. D., Petrie, L. E., & Eyfrig, R. (1973), 'A method for predicting the F1 layer critical frequency based on the Zurich smoothed sunspot number', *Radio Science*, 8(10), pp. 837–839. doi:10.1029/RS008i010p00837

Duly, T. M., Huba, J. D., & Makela, J. J. (2014), 'Self-consistent generation of MSTIDs within the SAMI3 numerical model', *Journal of Geophysical Research: Space Physics*, 119(8), pp. 6,745–6,757. doi:10.1002/2014JA020146

Dyson, P. L. (1975), 'Relationships between the rate of change of phase path (Doppler shift) and angle of arrival', *Journal of Atmospheric and Terrestrial Physics*, 37(8), pp. 1,151–1,154. doi:10.1016/0021-9169(75)90162-2

Dyson, P. L., & Bennett, J. A. (1988), 'A model of the vertical distribution of the electron concentration in the ionosphere and its application to oblique propagation studies', *Journal of Atmospheric and Terrestrial Physics*, 50(3), pp. 251–262. doi:10.1016/0021-9169(88)90074-8

Dyson, P. L., & Bennett, J. A. (1992), 'Exact ray path calculations using realistic ionospheres', *IEE Proceedings H (Microwaves, Antennas and Propagation)*, 139(5), pp. 407–413. doi:10.1049/ip-h-2.1992.0072

Dyson, P. L., Newton, G. P., & Brace, L. H. (1970), 'In Situ Measurements of Neutral and Electron Density Wave Structure from the Explorer 32 Satellite', *Journal of Geophysical Research*, 75(16), pp. 3,200–3,210. doi:10.1029/JA075i016p03200

Dyson, P. L., Parkinson, M. L., Monselesan, D. P., & Morris, R. J. (2001), 'On the determination of vertical profiles of ionospheric velocity from digital ionosonde measurements', *ANARE Report 146*, Australian Antarctic Division, pp. 209–220. Retrieved from https://data.aad.gov.au/aadc/pubs/get_pub.cfm?pub_id=10117

Earl, G. F., & Ward, B. D. (1986), 'Frequency Management Support for Remote Sea-State Sensing Using the JINDALEE Skywave Radar', *IEEE Journal of Oceanic Engineering*, OE-11(2), pp. 164–173. doi:10.1109/JOE.1986.1145165

Earl, G. F., & Ward, B. D. (1987), 'The frequency management system of the Jindalee over-the-horizon backscatter HF radar', *Radio Science*, 22(2), pp. 275–291. doi:10.1029/RS022i002p00275

Earle, G. D., Bhaneja, P., Roddy, P. A., Swenson, C. M., Barjatya, A., Bishop, R. L., et al. (2010), 'A comprehensive rocket and radar study of midlatitude spread *F*', *Journal of Geophysical Research*, 115(A12), A12339. doi:10.1029/2010JA015503

Eckermann, S. D., & Marks, C. J. (1997), 'GROGRAT: A new model of the global propagation and dissipation of atmospheric gravity waves', *Advances in Space Research*, 20(6), pp. 1253–1256. doi:10.1016/S0273-1177(97)00780-1

Elvidge, S., & Angling, M. J. (2014), 'Comparative testing of empirical, data assimilation and physics-based ionospheric/thermospheric models', *Proceedings of the 31st URSI General Assembly and Scientific Symposium*, Beijing, China. doi:10.1109/URSIGASS.2014.6929834

Elvidge, S., Angling, M. J., Codrescu, M., Fedrizzi, M., Galkin, I. A., Hernández-Pajares, M., et al. (2017), 'A Community Wide Ionospheric Modelling Challenge', *Proceedings of the 32nd URSI General Assembly and Scientific Symposium*, Montreal, Canada. Retrieved from http://www.ursi.org/proceedings/procGA17/papers/ Paper_JG8-1(1416).pdf

Elvidge, S., Angling, M. J., & Nava, B. (2014), 'On the use of modified Taylor diagrams to compare ionospheric assimilation models', *Radio Science*, 49(9), pp. 737–745. doi:10.1002/2014RS005435

Emmons, D. J., Dao, E. V., Knippling, K. K., McNamara, L. F., Nava. O. A., Obenberger, K. S., & Colman, J. J. (2020), 'Estimating horizontal phase speeds of a traveling ionospheric disturbance from Digisonde single site vertical ionograms', *Radio Science*, 55(8), e2020RS007089. doi:10.1029/2020RS007089

Fabrizio, G. A. (2000), *Space-Time Characterisation and Adaptive Processing of Ionospherically-Propagated HF Signals*, PhD thesis, University of Adelaide, Adelaide, South Australia, Australia.

Fabrizio, G. A. (2013), *High Frequency Over-the-Horizon Radar*, McGraw-Hill, New York, New York.

Fenwick, R. B., & Barry, G. H. (1966), 'Sweep-Frequency Oblique Ionospheric Sounding at Medium Frequencies', *IEEE Transactions on Broadcasting*, BC-12(1), pp. 25–27. doi:10.1109/TBC.1966.265861

Fitzgerald, R. J. (1974), 'Effects of Range-Doppler Coupling on Chirp Radar Tracking Accuracy', *IEEE Transactions on Aerospace and Electronic Systems*, 10(4), pp. 528–532. doi:10.1109/TAES.1974.307809

Folkestad, K. (1968), 'Exact Ray Computations in a Tilted Ionosphere With No Magnetic Field', *Radio Science*, 3(1), pp. 81–84. doi:10.1002/rds19683181

Forbes, J. M., Palo, S. E., & Zhang, X. (2000), 'Variability of the ionosphere', *Journal of Atmospheric and Solar-Terrestrial Physics*, 62(8), pp. 685–693. doi:10.1016/S1364-6826(00)00029-8

Francey, J. L. A. (1963), 'Ambipolar diffusion in meteor trails', *Australian Journal of Physics*, 16(4), pp. 500–506. doi:10.1071/PH630500

Francis, S. H. (1974), 'A Theory of Medium-Scale Traveling Ionospheric Disturbances', *Journal of Geophysical Research*, 79(34), pp. 5,245–5,260. doi:10.1029/JA079i034p05245

Francis, S. H. (1975), 'Global propagation of atmospheric gravity waves: A review', *Journal of Atmospheric and Terrestrial Physics*, 37(6–7), pp. 1,011–1,054. doi:10.1016/0021-9169(75)90012-4

Frazer, G. J. (2007), 'Forward-based Receiver Augmentation for OTHR', *Proceedings of the IEEE Radar Conference 2007*, Boston, Massachusetts, pp. 373–378. doi:10.1109/RADAR.2007.374245

Frazer, G. J., Heitmann, A. J., & Abramovich, Y. I. (2014), 'Initial Results From An Experimental Skew-Fire Mode-Selective Over-The-Horizon-Radar', *Proceedings of the IEEE International Radar Conference 2014*, Cincinnati, Ohio, pp. 350–353. doi:10.1109/RADAR.2014.6875613

Frazer, G. J., Meehan, D. H., Abramovich, Y. I., & Johnson, B. A. (2010), 'Mode-Selective OTHR: A New Cost-Effective Sensor for Maritime Domain Awareness', *Proceedings of the IEEE Radar Conference 2010*, Washington, DC, pp. 935–940. doi:10.1109/RADAR.2010.5494485

Fridman, S. V., & Nickisch, L. J. (2001), 'Generalization of ionospheric tomography on diverse data sources: Reconstruction of the three-dimensional ionosphere from simultaneous vertical ionograms, backscatter ionograms, and total electron content data', *Radio Science*, 36(5), pp. 1,129–1,139. doi:10.1029/1999RS002405

Fridman, S. V., Nickisch, L. J., Aiello, M., & Hausman, M. (2006), 'Real-time reconstruction of the three-dimensional ionosphere using data from a network of GPS receivers', *Radio Science*, 41(5), RS5S12. doi:10.1029/2005RS003341

Fridman, S. V., Nickisch, L. J., & Hausman, M. (2012), 'Inversion of backscatter ionograms and TEC data for over-the-horizon radar', *Radio Science*, 47(4), RS0L10. doi:10.1029/2011RS004932

Fridman, S. V., Nickisch, L. J., Hausman, M., & Zunich, G. (2016), 'Assimilative model for ionospheric dynamics employing delay, Doppler, and direction of arrival measurements from multiple HF channels', *Radio Science*, 51(3), pp. 176–183. doi:10.1002/2015RS005890

Friedlander, B., & Weiss, A. J. (1991), 'Direction Finding in the Presence of Mutual Coupling', *IEEE Transactions on Antennas and Propagation*, 39(3), pp. 273–284. doi:10.1109/8.76322

From, W. R., Sadler, E. M., & Whitehead, J. D. (1988), 'Measuring ionospheric movements using totally reflected radio waves', *Journal of Atmospheric and Terrestrial Physics*, 50(2), pp. 153–165. doi:10.1016/0021-9169(88)90052-9

Galkin, I. A., & Reinisch, B. W. (2008), 'The New ARTIST 5 for all Digisondes', *Ionosonde Network Advisory Group Bulletin*, vol. 69, Ionospheric Prediction Service, Sydney, New South Wales, Australia.

Galkin, I. A., Reinisch, B. W., Huang, X., & Bilitza, D. (2012), 'Assimilation of GIRO data into a real-time IRI', *Radio Science*, 47(4), RS0L07. doi:10.1029/2011RS004952

Galushko, V. G., Beley, V. S., Koloskov, A. V., Yampolski, Y. M., Paznukhov, V. V., Reinisch, B. W., Foster, J. C., & Erickson, P. (2003), 'Frequency-and-angular HF sounding and ISR diagnostics of TIDs', *Radio Science*, 38(6), 1102. doi:10.1029/2002RS002861

Galushko, V. G., Kashcheyev, A. S., Kashcheyev, S. B., Koloskov, A. V., Pikulik, I. I., Yampolski, Y. M., et al. (2007), 'Bistatic HF diagnostics of TIDs over the Antarctic Peninsula', *Journal of Atmospheric and Solar-Terrestrial Physics*, 69(4–5), pp. 403–410. doi:10.1016/j.jastp.2006.05.010

Galushko, V. G., Kashcheev, A. S., Paznukhov, V. V., Yampolski, Y. M., & Reinisch, B. W. (2008), 'Frequency-and-angular sounding of traveling ionospheric disturbances in the model of three-dimensional electron density waves', *Radio Science*, 43(4), pp. 1–10. doi:10.1029/2007RS003735

Gardiner-Garden, R. S. (2002), *Transforming between the GPS geoid and spherical Earth geometry for OTHR coordinate registration*, internal discussion note, Intelligence, Surveillance and Reconnaissance Division, Defence Science and Technology Organisation, Edinburgh, South Australia, Australia.

Gardiner-Garden, R., Ayliffe, J., Durbridge, L., Frazer, G., Heitmann, A., Praschifka, J., et al. (2008), 'A new high-fidelity oblique HF receiving system', *Proceedings of the 12th International Ionospheric Effects Symposium (IES2008)*, Alexandria, Virginia, pp. 482–491. Retrieved from https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/PB2008112709.xhtml

Gardiner-Garden, R., Cervera, M., Debnam, R., Harris, T., Heitmann, A., Holdsworth, D., et al. (2019), 'A description of the Elevation sensitive Oblique Incidence Sounder Experiment (ELOISE)', *Advances in Space Research*, 64(10), pp. 1,887–1,914. doi:10.1016/j.asr.2019.03.036

Gardiner-Garden, R., Heitmann, A., & Frazer, G. (2011), 'Variability observed in a high-fidelity model of the ionosphere based on a dense network or vertical and oblique incidence sounders', *Proceedings of the 13th International Ionospheric Effects Symposium (IES2011)*, Alexandria, Virginia, pp. 91–102. Retrieved from https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/PB2011113661.xhtml

Gardiner-Garden, R., Heitmann, A., & Northey, B. (2018), *A Parametric Model of the Ionospheric Electron Density Profile for JORN*, DST Group Technical Note TN-1722, National Security and ISR Division, Defence Science and Technology Group, Edinburgh, South Australia, Australia. Retrieved from https://www.dst.defence.gov.au/publication/parametric-model-ionospheric-electron-density-profile-jorn

Gardiner-Garden, R., Heitmann, A., Northey, B., & Turley, M. (2008), 'Modelling uncertainty in a real-time model (nowcast) of the ionosphere', *Proceedings of the 12th International Ionospheric Effects Symposium (IES2008)*, Alexandria, Virginia, pp. 405–414. Retrieved from https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/PB2008112709.xhtml

George, P. L. (1970), 'True height analysis of oblique incidence HF radio wave data', *Journal of Atmospheric and Terrestrial Physics*, 32(5), pp. 905–916. doi:10.1016/0021-9169(70)90035-8

George, P. L., & Bradley, P. A. (1974), 'A new method of predicting the ionospheric absorption of high frequency waves at oblique incidence', *ITU Telecommunication Journal*, 41(5), pp. 307–312.

Georges, T. M. (1967a), 'Evidence for the Influence of Atmospheric Waves on Ionospheric Motions', *Journal of Geophysical Research*, 72(1), pp. 422–425. doi:10.1029/JZ072i001p00422

Georges, T. M. (1967b), *Ionospheric Effects of Atmospheric Waves*, ESSA Technical Report IER 57-ITSA 54, U.S. Government Printing Office, Washington, DC.

Georges, T. M. (1968), 'HF Doppler studies of traveling ionospheric disturbances', *Journal of Atmospheric and Terrestrial Physics*, 30(5), pp. 735–746. doi:10.1016/S0021-9169(68)80029-7

Gething, P. J. D. (1991), *Radio Direction Finding and Superresolution*, 2nd edition, IEEE Electromagnetic Waves Series, vol. 33, Peter Peregrinus, London, UK. doi:10.1049/PBEW033E

Gillmor, C. S. (1982), 'Wilhelm Altar, Edward Appleton, and the magneto-ionic theory', *Proceedings of the American Philosophical Society*, 126(5), pp. 395–440. Retrieved from https://www.jstor.org/stable/986214

Godin, O. A. (2015), 'Wentzel–Kramers–Brillouin approximation for atmospheric waves', *Journal of Fluid Mechanics*, 777, pp. 260–290. doi:10.1017/jfm.2015.367

Gong, W., Cui, X., & Pan, L. (2016), 'Design and application of the digital multifunctional ionosonde', *IET Radar, Sonar & Navigation*, 10(7), pp. 1303–1309. doi:10.1049/iet-rsn.2015.0561

Gossard, E. E. (1962), 'Vertical Flux of Energy into the Lower Ionosphere from Internal Gravity Waves Generated in the Troposphere', *Journal of Geophysical Research*, 67(2), pp. 745–757. doi:10.1029/JZ067i002p00745

Grubb, R. N., Livingston, R., & Bullett, T. W. (2008), 'A new general purpose high performance HF Radar', *Proceedings of the 29th URSI General Assembly and Scientific Symposium*, Chicago, Illinois. Retrieved from http://www.ursi.org/proceedings/procGA08/papers/GHp4.pdf

Gupta, I. J., & Ksienski, A. A. (1983), 'Effect of Mutual Coupling on the Performance of Adaptive Arrays', *IEEE Transactions on Antennas and Propagation*, 31(5), pp. 785–791. doi:10.1109/TAP.1983.1143128

Hajkowicz, L. A. (2007), 'Morphology of quantified ionospheric range spread-F over a wide range of midlatitudes in the Australian longitudinal sector', *Annales Geophysicae*, 25(5), pp. 1,125–1,130. doi:10.5194/angeo-25-1125-2007

Hanselman, D. C., & Littlefield, B. (2005), *Mastering MATLAB® 7*, Pearson Prentice Hall, Upper Saddle River, New Jersey.

Härdle, W. K., & Simar, L. (2015), *Applied Multivariate Statistical Analysis*, 4th edition, Springer, Berlin, Germany.

Harris, T. J., Cervera, M. A., & Meehan, D. H. (2012), 'SpICE: A program to study small-scale disturbances in the ionosphere', *Journal of Geophysical Research*, 117(A6), A06321. doi:10.1029/2011JA017438

Harris, T. J., Cervera, M. A., Pederick, L. H., & Quinn, A. D. (2017), 'Separation of *O/X* Polarization Modes on Oblique Ionospheric Soundings', *Radio Science*, 52(12), pp. 1522–1533. doi:10.1002/2017RS006280

Harris, T. J., & Frazer, G. J. (2005), 'High Resolution Radar Measurements of Ionospheric Variations on an Oblique Path', *Proceedings of the 11th International Ionospheric Effects Symposium (IES2005)*, Alexandria, Virginia, pp. 290–297. Retrieved from https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/ PB2007100059.xhtml

Harris, T. J., Quinn, A. D., & Pederick, L. H. (2016), 'The DST group ionospheric sounder replacement for JORN', *Radio Science*, 51(6), pp. 563–572. doi:10.1002/2015RS005881

Haselgrove, C. B., & Haselgrove, J. (1960), 'Twisted Ray Paths in the Ionosphere', *Proceedings of the Physical Society*, 75(3), pp. 357–363. doi:10.1088/0370-1328/75/3/304

Haselgrove, J. (1963), 'The Hamiltonian ray path equations', *Journal of Atmospheric and Terrestrial Physics*, 25(7), pp. 397–399. doi:10.1016/0021-9169(63)90173-9

He, L.-S., Dyson, P. L., Parkinson, M. L., & Wan, W. (2004), 'Studies of medium scale travelling ionospheric disturbances using TIGER SuperDARN radar sea echo observations', *Annales Geophysicae*, 22(12), pp. 4,077–4,088. doi:10.5194/angeo-22-4077-2004

Headrick, J. M., & Anderson, S. J. (2008), 'HF Over-the-Horizon Radar', in M. I. Skolnik (ed.), *Radar Handbook*, 3rd edition, McGraw-Hill, New York, New York, ch. 20.

Headrick, J. M., & Thomason, J. F. (1996), 'Naval Applications of High Frequency Over-the-Horizon Radar', *Naval Engineers Journal*, 108(3), pp. 353–362. doi:10.1111/j.1559-3584.1996.tb01572.x

Hearn, A. L., & Yeh, K. C. (1977), 'Medium Scale TID's and Their Associated Internal Gravity Waves as Seen Through Height-Dependent Electron Density Power Spectra', *Journal of Geophysical Research*, 82(32), pp. 4,983–4,990. doi:10.1029/JA082i032p04983

Heitmann, A. J., Cervera, M. A., Gardiner-Garden, R. S., Holdsworth, D. A., MacKinnon, A. D., Reid, I. M., & Ward, B. D. (2017), 'Observations and modeling of traveling ionospheric disturbance signatures from an Australian network of oblique angle-of-arrival sounders' (summary paper), *Proceedings of the 32nd URSI General Assembly and Scientific Symposium*, Montreal, Canada. doi:10.23919/URSIGASS.2017.8105329

Heitmann, A. J., Cervera, M. A., Gardiner-Garden, R. S., Holdsworth, D. A., MacKinnon, A. D., Reid, I. M., & Ward, B. D. (2018), 'Observations and modeling of traveling ionospheric disturbance signatures from an Australian network of oblique angle-of-arrival sounders', *Radio Science*, 53(9), pp. 1,089–1,107. doi:10.1029/2018RS006613

Heitmann, A. J., & Gardiner-Garden, R. S. (2019), 'A robust feature extraction and parameterized fitting algorithm for bottom-side oblique and vertical incidence ionograms', *Radio Science*, 54(1), pp. 115–134. doi:10.1029/2018RS006682

Herman, J. R. (1966), 'Spread *F* and ionospheric *F*-region irregularities', *Reviews of Geophysics*, 4(2), pp. 255–299. doi:10.1029/RG004i002p00255

Hill, J. R. (1979), 'Exact ray paths in a multisegment quasi-parabolic ionosphere', *Radio Science*, 14(5), pp. 855-861. doi:10.1029/RS014i005p00855

Hilsenrath, J., Beckett, C. W., Benedict, W. S., Fano, L., Hoge, H. J., Masi, J. F., et al. (1955), 'Tables of Thermal Properties of Gases', *NBS Circular 564*, National Bureau of Standards Circular, Government Printing Office, Washington, DC. Retrieved from https://archive.org/details/circularofbureau564hils

Hines, C. O. (1960), 'Internal atmospheric gravity waves at ionospheric heights', *Canadian Journal of Physics*, 38(11), pp. 1,441–1,481. doi:10.1139/p60-150

Hines, C. O. (1965), 'Atmospheric Gravity Waves: A New Toy for the Wave Theorist', *Journal of Research of the National Bureau of Standards, Section D: Radio Science*, 69D(3), pp. 375–380.

Hines, C. O. (1968), 'An effect of ohmic losses in upper atmospheric gravity waves', *Journal of Atmospheric and Terrestrial Physics*, 30(5), pp. 851–856. doi:10.1016/S0021-9169(68)80037-6

Hines, C. O., et al. (1974), *The Upper Atmosphere in Motion*, Geophysical Monograph Series, vol. 18, American Geophysical Union, Washington, DC. doi:10.1029/GM018

Hines, C. O., & Rao, R. R. (1968), 'Validity of three-station methods of determining ionospheric motions', *Journal of Atmospheric and Terrestrial Physics*, 30(5), pp. 979–993. doi:10.1016/S0021-9169(68)80048-0

Hines, C. O., & Reddy, C. A. (1967), 'On the propagation of atmospheric gravity waves through regions of wind shear', *Journal of Geophysical Research*, 72(3), pp. 1,015–1,034. doi:10.1029/JZ072i003p01015

Hocke, K., & Schlegel, K. (1996), 'A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995', *Annales Geophysicae*, 14(9), pp. 917–940. doi:10.1007/s00585-996-0917-6

Hoctor, R. T., & Kassam, S. A. (1990), 'The Unifying Role of the Coarray in Aperture Synthesis for Coherent and Incoherent Imaging', *Proceedings of the IEEE*, 78(4), pp. 735–752. doi:10.1109/5.54811

Holdsworth, D. A., & Reid, I. M. (2004), 'Comparisons of full correlation analysis (FCA) and imaging Doppler interferometry (IDI) winds using the Buckland Park MF radar', *Annales Geophysicae*, 22, pp. 3,829–3,842. doi:10.5194/angeo-22-3829-2004

Hooke, W. H. (1968), 'Ionospheric irregularities produced by internal atmospheric gravity waves', *Journal of Atmospheric and Terrestrial Physics,* 30(5), pp. 795–823. doi:10.1016/S0021-9169(68)80033-9

Hooke, W. H. (1970), 'The Ionospheric Response to Internal Gravity Waves – 1. The F_2 Region Response', *Journal of Geophysical Research*, 75(28), pp. 5,535–5,544. doi:10.1029/JA075i028p05535

Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., Yen, N.-C., Tung, C. C., & Liu, H. H. (1998), 'The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis', *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 454(1,971), pp. 903–995. doi:10.1098/rspa.1998.0193

Huang, X., Reinisch, B. W., Sales, G. S., Paznukhov, V. V., & Galkin, I. A. (2016), 'Comparing TID simulations using 3-D ray tracing and mirror reflection', *Radio Science*, 51(4), pp. 337–343. doi:10.1002/2015RS005872

Huba, J. D., Joyce, G., & Krall, J. (2008), 'Three-dimensional equatorial spread *F* modeling', *Geophysical Research Letters*, 35(10), L10102. doi:10.1029/2008GL033509

Hui, H. T. (2007), 'Decoupling Methods for the Mutual Coupling Effect in Antenna Arrays: A Review', *Recent Patents on Engineering*, 1(2), pp. 187–193. doi:10.2174/187221207780832200

Hunsucker, R. D. (1982), 'Atmospheric Gravity Waves Generated in the High-Latitude Ionosphere: A Review', *Reviews of Geophysics and Space Physics*, 20(2), pp. 293–315. doi:10.1029/RG020i002p00293

Hunsucker, R. D. (1991), *Radio Techniques for Probing the Terrestrial Ionosphere*, Physics and Chemistry in Space Planetology, vol. 22, Springer-Verlag, Berlin, Germany. doi:10.1007/978-3-642-76257-4

Hysell, D. L., Milla, M. A., & Vierinen, J. (2016), 'A multistatic HF beacon network for ionospheric specification in the Peruvian sector', *Radio Science*, 51(5), pp. 392–401. doi:10.1002/2016RS005951

Ippolito, A., Scotto, C., Francis, M., Settimi, A., & Cesaroni, C. (2015), 'Automatic interpretation of oblique ionograms', *Advances in Space Research*, 55(6), pp. 1,624–1,629. doi:10.1016/j.asr.2014.12.025

Ishida, T., Hosokawa, K., Shibata, T., Suzuki, S., Nishitani, N., & Ogawa, T. (2008), 'SuperDARN observations of daytime MSTIDs in the auroral and mid-latitudes: Possibility of long-distance propagation', *Geophysical Research Letters*, 35(13), L13102. doi:10.1029/2008GL034623

Ivanov, V. A., Kurkin, V. I., Nosov, V. E., Uryadov, V. P., & Shumaev, V. V. (2003), 'Chirp ionosonde and its application in the ionospheric research', *Radiophysics and Quantum Electronics*, 46(11), pp. 821–851. doi:10.1023/B:RAQE.0000028576.51983.9c

Jacobson, A. R. (1994), 'Comments on "Middle and upper atmosphere radar observations of ionospheric density gradients produced by gravity wave packets" by W. L. Oliver et al.', *Journal of Geophysical Research*, 99(A11), pp. 21,411–21,413. doi:10.1029/94JA02208

Jacobson, A. R., & Carlos, R. C. (1991), 'A study of apparent ionospheric motions associated with multiple Traveling Ionospheric Disturbances', *Journal of Atmospheric and Terrestrial Physics*, 53(1–2), pp. 53–62. doi:10.1016/0021-9169(91)90020-8

Jakowski, N., Béniguel, Y., De Franceschi, G., Hernandez Pajares, M., Jacobsen, K. S., Stanislawska, I., et al. (2012), 'Monitoring, tracking and forecasting ionospheric perturbations using GNSS techniques', *Journal of Space Weather and Space Climate*, 2, A22. doi:10.1051/swsc/2012022

Jakowski, N., Wehrenpfennig, A., Heise, S., Reigber, Ch., Lühr, H., Grunwaldt, L., & Meehan, T. K. (2002), 'GPS radio occultation measurements of the ionosphere from CHAMP: Early results', *Geophysical Research Letters*, 29(10), 1,457. doi:10.1029/2001GL014364

Jiang, C. H., Wei, L. H., Yang, G. B., Zhou, C., & Zhao, Z. Y. (2020), 'Numerical simulation of the propagation of electromagnetic waves in ionospheric irregularities', *Earth and Planetary Physics*, 4(6), pp. 565–570. doi:10.26464/epp2020059

Johnson, B. A., Amey, P. F., Braendler, S., & Turley, M. D. E. (2008), 'In-Service Enhancements of the Jindalee Over-the-Horizon Radars', *Proceedings of the IEEE International Radar Conference 2008*, Adelaide, South Australia, Australia, pp. 218–223. doi:10.1109/RADAR.2008.4653921

Jones, J., & Lyon, G. F. (1974), 'The focussing of radio waves by large scale ionospheric irregularities', *Journal of Atmospheric and Terrestrial Physics*, 36(3), pp. 461–473. doi:10.1016/0021-9169(74)90126-3

Jones, T. B., & Reynolds, J. S. B. (1975), 'Ionospheric perturbations and their effect on the accuracy of h.f. direction finders', *The Radio and Electronic Engineer*, 45(1–2), pp. 63–73. doi:10.1049/ree.1975.0008

Kelley, M. C. (2009), *The Earth's Ionosphere: Plasma Physics and Electrodynamics*, 2nd edition, International Geophysics Series, vol. 96, Academic Press, Burlington, Massachusetts.

Kelley, M. C. (2011), 'On the origin of mesoscale TIDs at midlatitudes', *Annales Geophysicae*, 29(2), pp. 361–366. doi:10.5194/angeo-29-361-2011

Kelley, M. C., & Fukao, S. (1991), 'Turbulent Upwelling of the Mid-Latitude Ionosphere: 2. Theoretical Framework', *Journal of Geophysical Research*, 96(A3), pp. 3,747–3,753. doi:10.1029/90JA02252

Kelley, M. C., Larsen, M. F., LaHoz, C., & McClure, J. P. (1981), 'Gravity wave initiation of equatorial spread F: A case study', *Journal of Geophysical Research*, 87(A11), pp. 9,087–9,100. doi:10.1029/JA086iA11p09087

Kelley, M. C., & Makela, J. J. (2001), 'Resolution of the discrepancy between experiment and theory of midlatitude *F*-region structures', *Geophysical Research Letters*, 28(13), pp. 2,589–2,592. doi:10.1029/2000GL012777

Kelley, M. C., & Miller, C. A. (1997), 'Electrodynamics of midlatitude spread *F*: 3. Electrohydrodynamic waves? A new look at the role of electric fields in thermospheric wave dynamics', *Journal of Geophysical Research*, 102(A6), pp. 11,539–11,547. doi:10.1029/96JA03841

King, G. A. M. (1970), 'Spread-*F* on ionograms', *Journal of Atmospheric and Terrestrial Physics*, 32(2), pp. 209–221. doi:10.1016/0021-9169(70)90192-3

Kirchengast, G. (1996), 'Elucidation of the physics of the gravity wave-TID relationship with the aid of theoretical simulations', *Journal of Geophysical Research*, 101(A6), pp. 13,353–13,368. doi:10.1029/96JA00750

Kirchengast, G., Hocke, K., & Schlegel, K. (1995), 'Gravity waves determined by modelling of traveling ionospheric disturbances in incoherent-scatter radar measurements', *Radio Science*, 30(5), pp. 1,551–1,567. doi:10.1029/95RS02080

Kirchengast, G., Hocke, K., & Schlegel, K. (1996), 'The gravity wave-TID relationship: insight via theoretical model-EISCAT data comparison', *Journal of Atmospheric and Terrestrial Physics*, 58(1–4), pp. 233–243. doi:10.1016/0021-9169(95)00032-1

Klostermeyer, J. (1969), 'Gravity waves in the *F*-region', *Journal of Atmospheric and Terrestrial Physics*, 31(1), pp. 25–45. doi:10.1016/0021-9169(69)90079-8

Knippling, K. L. (2018), *A Comparison of High Frequency Angle of Arrival and Ionosonde Data During a Traveling Ionospheric Disturbance*, masters thesis, AFIT-ENP-MS-18-M-087, Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio. Retrieved from https://apps.dtic.mil/sti/citations/AD1056258

Kohl, H., & King, J. W. (1967), 'Atmospheric winds between 100 and 700 km and their effects on the ionosphere', *Journal of Atmospheric and Terrestrial Physics*, 29(9), pp. 1,045–1,062. doi:10.1016/0021-9169(67)90139-0

Kotake, N., Otsuka, Y., Ogawa, T., Tsugawa, T., & Saito, A. (2007), 'Statistical study of medium-scale traveling ionospheric disturbances observed with the GPS networks in Southern California', *Earth, Planets and Space*, 59(2), pp. 95–102. doi:10.1186/BF03352681

Kouris, S. S., & Muggleton, L. M. (1973), 'Diurnal variation in the E-layer ionization', *Journal of Atmospheric and Terrestrial Physics*, 35(1), pp. 133–139. doi:10.1016/0021-9169(73)90221-3

Kozick, R. J., & Kassam, S. A. (1992), 'Coarray synthesis with circular and elliptical boundary arrays', *IEEE Transactions on Image Processing*, 1(3), pp. 391–405. doi:10.1109/83.148611

Kozlov, A. V, & Paznukhov, V. V. (2008), 'Digisonde Drift Analysis Software', *AIP Conference Proceedings*, 974(1), pp. 167–174. doi:10.1063/1.2885026

Kraus, J. D. (1988), *Antennas*, 2nd edition, McGraw-Hill, New York, New York.

Krim, H., & Viberg, M. (1996), 'Two Decades of Array Signal Processing Research: The Parametric Approach', *IEEE Signal Processing Magazine*, 13(4), pp 67–94. doi:10.1109/79.526899

Kudeki, E. (1988), 'Radar Interferometer Observations of Mesospheric Echoing Layers at Jicamarca', *Journal of Geophysical Research*, 93(D5), pp. 5,413–5,421. doi:10.1029/JD093iD05p05413

Laštovička, J. (2006), 'Forcing of the ionosphere by waves from below', *Journal of Atmospheric and Solar-Terrestrial Physics*, 68(3–5), pp. 479–497. doi:10.1016/j.jastp.2005.01.018

Lei, J., Liu, L., Wan, W., & Zhang, S.-R. (2005), 'Variations of electron density based on long-term incoherent scatter radar and ionosonde measurements over Millstone Hill', *Radio Science*, 40(2), RS2008. doi:10.1029/2004RS003106

Leitinger, R., Zhang, M.-L., & Radicella, S. M. (2005), 'An improved bottomside for the ionospheric electron density model NeQuick', *Annals of Geophysics*, 48(3), pp. 525–534. doi:10.4401/ag-3217

Liu, C. H., & Yeh, K. C. (1969), 'Effect of Ion Drag on Propagation of Acoustic-Gravity Waves in the Atmospheric *F* Region', *Journal of Geophysical Research*, 74(9), pp. 2,248–2,255. doi:10.1029/JA074i009p02248

Lobb, R. J., & Titheridge, J. E. (1977), 'The effects of travelling ionospheric disturbances on ionograms', *Journal of Atmospheric and Terrestrial Physics*, 39(2), pp. 129–138. doi:10.1016/0021-9169(77)90106-4

Lynn, K. J. W. (2000), *The Range Conversion of Ionograms*, Surveillance Systems Division, Defence Science and Technology Organisation, Edinburgh, South Australia, Australia.

Lynn, K. J. W. (2008), 'Ionospheric observations made by a time-interleaved Doppler ionosonde', *Advances in Space Research*, 42(7), pp. 1,218–1,230. doi:10.1016/j.asr.2007.06.007

Lynn, K. J. W. (2009), 'A technique for calculating ionospheric Doppler shifts from standard ionograms suitable for scientific, HF communication, and OTH radar applications', *Radio Science*, 44(6), RS6002. doi:10.1029/2009RS004210

Lynn, K. J. W., Gardiner-Garden, R. S., & Heitmann, A. (2014), 'The spatial and temporal structure of twin peaks and midday bite out in f_oF_2 (with associated height changes) in the Australian and South Pacific low midlatitude ionosphere', *Journal of Geophysical Research: Space Physics*, 119(12), pp. 10,294–10,304. doi:10.1002/2014JA020617

Lynn, K. J. W., Gardiner-Garden, R. S., & Heitmann, A. (2016), 'The observed compression and expansion of the F_2 ionosphere as a major component of ionospheric variability', *Radio Science*, 51(5), pp. 538–552. doi:10.1002/2016RS006036

Lynn, K. J. W., Otsuka, Y., & Shiokawa, K. (2013), 'Ionogram-based range-time displays for observing relationships between ionosonde satellite traces, spread F and drifting optical plasma depletions', *Journal of Atmospheric and Solar-Terrestrial Physics*, 98, pp. 105–112. doi:10.1016/j.jastp.2013.03.020

Lyon, G. F. (1979), 'The corrugated reflector model for one-hop oblique propagation', *Journal of Atmospheric and Terrestrial Physics*, 41(1), pp. 5–9. doi:10.1016/0021-9169(79)90041-2

MacDougall, J., Abdu, M. A., Batista, I., Buriti, R., Medeiros, A. F., Jayachandran, P. T., & Borba, G. (2011), 'Spaced transmitter measurements of medium scale traveling ionospheric disturbances near the equator', *Geophysical Research Letters*, 38(16), L16806. doi:10.1029/2011GL048598

MacDougall, J. W., Grant, I. F., & Shen, X. (1995), 'The Canadian Advanced Digital Ionosonde: Design and Results', *Ionosonde Networks and Stations*, Report UAG-104, World Data Center A for Solar-Terrestrial Physics, Boulder, Colorado, pp. 21–27. Retrieved from https://www.ngdc.noaa.gov/stp/space-weather/online-publications/ stp_uag/uag-104.pdf

MacLeod, M. A. (1966), 'Sporadic E Theory. I. Collision-Geomagnetic Equilibrium', *Journal of the Atmospheric Sciences*, 23(1), pp. 96–109. doi:10.1175/1520-0469(1966)023<0096:SETICG>2.0.CO;2

Maeda, S., & Handa, S. (1980), 'Transmission of large-scale TIDs in the ionospheric *F2*-region', *Journal of Atmospheric and Terrestrial Physics*, 42(9–10), pp. 853–859. doi:10.1016/0021-9169(80)90089-6

Makela, J. J. (2006), 'A review of imaging low-latitude ionospheric irregularity processes', *Journal of Atmospheric and Solar-Terrestrial Physics*, 68(13), pp. 1,441–1,458. doi:10.1016/j.jastp.2005.04.014

Mannucci, A. J., Wilson, B. D., Yuan, D. N., Ho, C. H., Lindqwister, U. J., & Runge, T. F. (1998), 'A global mapping technique for GPS-derived ionospheric total electron content measurements', *Radio Science*, 33(3), pp. 565–582. doi:10.1029/97RS02707

Maraun, D., & Kurths, J. (2004), 'Cross wavelet analysis: significance testing and pitfalls', *Nonlinear Processes in Geophysics*, 11(4), pp. 505–514. doi:10.5194/npg-11-505-2004

Martyn, D. F. (1935), 'The propagation of medium radio waves in the ionosphere', *Proceedings of the Physical Society*, 47(2), pp. 323–339. doi:10.1088/0959-5309/47/2/311

Martyn, D. F. (1950), 'Cellular atmospheric waves in the ionosphere and troposphere', *Proceedings of the Royal Society of London A: Mathematical and Physical Sciences*, 201(1,065), pp. 216–234. doi:10.1098/rspa.1950.0055

McNamara, L. F. (1991), *The Ionosphere: Communications, Surveillance, and Direction Finding*, Krieger Publishing Company, Malabar, Florida.

McNamara, L. F., Angling, M. J., Elvidge, S., Fridman, S. V., Hausman, M. A., Nickisch, L. J., & McKinnell, L.-A. (2013), 'Assimilation procedures for updating ionospheric profiles below the F2 peak', 48(2), pp. 143–157. doi:10.1002/rds.20020

McNamara, L. F., Bishop, G. J., & Welsh, J. A. (2011), 'Assimilation of ionosonde profiles into a global ionospheric model', *Radio Science*, 46(2), RS2006. doi:10.1029/2010RS004457

McNamara, L. F., Cooke, D. L., Valladares, C. E., & Reinisch, B. W. (2007), 'Comparison of CHAMP and Digisonde plasma frequencies at Jicamarca, Peru', *Radio Science*, 42(2), RS2005. doi:10.1029/2006RS003491

McNamara, L. F., Decker, D. T., Welsh, J. A., & Cole, D. G. (2007), 'Validation of the Utah State University Global Assimilation of Ionospheric Measurements (GAIM) model predictions of the maximum usable frequency for a 3000 km circuit', *Radio Science*, 42(3), RS3015. doi:10.1029/2006RS003589

McNicol, R. W. E., Webster, H. C., & Bowman, G. G. (1956), 'A study of "spread-*F*" ionospheric echoes at night at Brisbane: I. Range spreading (experimental)', *Australian Journal of Physics*, 9(2), pp. 247–271. doi:10.1071/PH560247

Medeiros, A. F., Taylor, M. J., Takahashi, H., Batista, P. P., & Gobbi, D. (2003), 'An investigation of gravity wave activity in the low-latitude upper mesosphere: Propagation direction and wind filtering', *Journal of Geophysical Research*, 108(D14), 4411. doi:10.1029/2002JD002593

Meek, C. E., & Manson, A. H. (1987), 'Mesospheric motions observed by simultaneous medium-frequency interferometer and spaced antenna experiments', *Journal of Geophysical Research*, 92(D5), pp. 5,627–5,639. doi:10.1029/JD092iD05p05627

Miller, W. (1951), 'Effective Earth's Radius for Radiowave Propagation Beyond the Horizon', *Journal of Applied Physics*, 22(1), pp. 55–62. doi:10.1063/1.1699820

Millman, G. H. (1977), 'The effect of tropospheric refraction on HF propagation', R77EMH5, General Electric Company, Syracuse, New York.

Mills, B. Y., & Little, A. G. (1953), 'A high-resolution aerial system of a new type', *Australian Journal of Physics*, 6(3), pp. 272–278. doi:10.1071/PH530272

Millward, G. H., Moffett, R. J., Quegan, S., & Fuller-Rowell, T. J. (1993), 'Effects of an Atmospheric Gravity Wave on the Midlatitude Ionospheric *F* Layer', *Journal of Geophysical Research*, 98(A11), pp. 19,173–19,179. doi:10.1029/93JA02093

Mitchell, C. N., Rankov, N. R., Bust, G. S., Miller, E., Gaussiran, T., Calfas, R., et al. (2017), 'Ionospheric data assimilation applied to HF geolocation in the presence of traveling ionospheric disturbances', *Radio Science*, 52(7), pp. 829–840. doi:10.1002/2016RS006187

Miyoshi, Y., Jin, H., Fujiwara, H., & Shinagawa, H. (2017), 'Numerical Study of Traveling Ionospheric Disturbances Generated by an Upward Propagating Gravity Wave', *Journal of Geophysical Research: Space Physics*, 123(3), pp. 2,141–2,155. doi:10.1002/2017JA025110

Morgan, M. G., & Calderón, C. H. J. (1978), 'Testing Hooke's Perturbation Formula for the Production of TID's by Gravity Waves', *Journal of Geophysical Research*, 83(A12), pp. 5,737–5,740. doi:10.1029/JA083iA12p05737

Morgan, M. G., Calderón, C. H. J., & Ballard, K. A. (1978), 'Techniques for the study of TID's with multi-station rapid-run ionosondes', *Radio Science*, 13(4), pp. 729–741. doi:10.1029/RS013i004p00729

Morgan, M. G., & Tedd, B. L. (1983), 'The Dispersion of Traveling Ionospheric Disturbances', *Journal of Geophysical Research*, 88(A12), pp. 10,253–10,258. doi:10.1029/JA088iA12p10253

Muggleton, L. M. (1975), 'A method of predicting foE at any time and place', *ITU Telecommunication Journal*, 42(7), pp. 413–418.

Munro, G. H. (1948), 'Short-Period Changes in the *F* Region of the Ionosphere', *Nature*, 162(4127), pp. 886–887. doi:10.1038/162886a0

Munro, G. H. (1950), 'Travelling disturbances in the ionosphere', *Proceedings of the Royal Society of London A: Mathematical and Physical Sciences*, 202(1,069), pp. 208–223. doi:10.1098/rspa.1950.0095

Munro, G. H. (1958), 'Travelling ionospheric disturbances in the *F* region', *Australian Journal of Physics*, 11(1), pp. 91–112. doi:10.1071/PH580091

Munro, G. H., & Heisler, L. H. (1956a), 'Cusp Type Anomalies in Variable Frequency Ionospheric Records', *Australian Journal of Physics*, 9(3), pp. 343–358. doi:10.1071/PH560343

Munro, G. H., & Heisler, L. H. (1956b), 'Divergence of Radio Rays in the Ionosphere', *Australian Journal of Physics*, 9(3), pp. 359–372. doi:10.1071/PH560359

Munton, D., Rainwater, D., Gaussiran, T., Calfas, R., Reay, M., Schofield, J., & Fleischmann, A. (2019), 'A mid-latitude HF propagation experiment over New Mexico', *Radio Science*, 54(3), pp. 298–313. doi:10.1029/2018RS006740

Narayanan, V. L., Shiokawa, K., Otsuka, Y., & Neudegg, D. (2018), 'On the Role of Thermospheric Winds and Sporadic *E* Layers in the Formation and Evolution of Electrified MSTIDs in Geomagnetic Conjugate Regions', *Journal of Geophysical Research: Space Physics*, 123(8), pp. 6,957–6,980. doi:10.1029/2018JA025261

Nava, B., Coïsson, P., & Radicella, S. M. (2008), 'A new version of the NeQuick ionosphere electron density model', *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(15), pp. 1,856–1,862. doi:10.1016/j.jastp.2008.01.015

Negrea, C., & Zabotin, N. A. (2016), 'Mean spectral characteristics of acoustic gravity waves in the thermosphere-ionosphere determined from Dynasonde data', *Radio Science*, 51(3), pp. 213–222. doi:10.1002/2015RS005823

Negrea, C., Zabotin, N., Bullett, T., Fuller-Rowell, T., Fang, T.-W., & Codrescu, M. (2016), 'Characteristics of acoustic gravity waves obtained from Dynasonde data', *Journal of Geophysical Research: Space Physics*, 121(4), pp. 3,665–3,680. doi:10.1002/2016JA022495

Nelson, G. J., & Cogger, L. L. (1971), 'Dynamical behaviour of the nighttime ionosphere at Arecibo', *Journal of Atmospheric and Terrestrial Physics*, 33(11), pp. 1,711–1,726. doi:10.1016/0021-9169(71)90219-4

Netherway, D. J., & Gardiner-Garden, R. S. (2015), 'On Power Measured and Predicted for Oblique HF Propagation', *Proceedings of the 15th Australian Space Research Conference 2015*, Canberra, Australian Capital Territory, Australia, pp. 183–194. Retrieved from https://marssociety.org.au/sites/default/files/15ASRC-proceedings.pdf

Newton, M. (2018), *Australian Atmospheric Sounding Information*. Accessed 4 June 2018 from http://slash.dotat.org/cgi-bin/atmos

Nguyen, V. K., & Turley, M. D. E. (2015), 'Bandwidth Extrapolation of LFM Signals for Narrowband Radar Systems', *IEEE Transactions on Aerospace and Electronic Systems*, 51(1), pp. 702–712. doi:10.1109/TAES.2014.130525

Nickisch, L. J. (1988), 'Focusing in the stationary phase approximation', *Radio Science*, 23(2), pp. 171–182. doi:10.1029/RS023i002p00171

Nickisch, L. J. (1992), 'Non-uniform motion and extended media effects on the mutual coherence function: An analytic solution for spaced frequency, position, and time', *Radio Science*, 27(1), pp. 9–22. doi:10.1029/91RS02356

Nickisch, L. J. (2008), 'Practical Applications of Haselgrove's Equations for HF Systems', *The Radio Science Bulletin*, 325, pp. 36–48. doi:10.23919/URSIRSB.2008.7909584

Nickisch, L. J., Fridman, S., Hausman, M., Kraut, S., & Zunich, G. (2016), 'Assimilative modeling of ionospheric dynamics for nowcasting of HF propagation channels in the presence of TIDs', *Radio Science*, 51(3), pp. 184–193. doi:10.1002/2015RS005902

Nickisch, L. J., Hausman, M. A., & Fridman, S. V. (2006), 'Range rate–Doppler correlation for HF propagation in traveling ionospheric disturbance environments', *Radio Science*, 41(5), RS5S39. doi:10.1029/2005RS003358

Nickisch, L. J., Hausman, M. A., & Fridman, S. (2007), 'Traveling Ionospheric Disturbance Mitigation for OTH Radar', *Proceedings of the IEEE Radar Conference 2007*, Boston, Massachusetts, pp. 362–366. doi:10.1109/RADAR.2007.374243

Norman, R. J., & Cannon, P. S. (1999), 'An evaluation of a new two-dimensional analytic ionospheric ray-tracing technique: Segmented method for analytic ray tracing (SMART)', *Radio Science*, 34(2), pp. 489–499. doi:10.1029/98RS01788

Norman, R. J., Le Marshall, J., Carter, B. A., Wang, C.-S., Gordon, S., & Zhang, K. (2012), 'A New Pseudo Three-Dimensional Segment Method Analytical Ray Tracing (3-D SMART) Technique', *IEEE Transactions on Antennas and Propagation*, 60(12), pp. 5,818–5,824. doi:10.1109/TAP.2012.2214194

Obenberger, K. S., Dao., E., & Dowell, J. D. (2019), 'Experimenting with frequency-and-angular sounding to characterize traveling ionospheric disturbances using the LWA-SV radio telescope and a DPS4D', *Radio Science*, 54(2), pp. 181–193. doi:10.1029/2018RS006690

Ogawa, T., Balan, N., Otsuka, Y., Shiokawa, K., Ihara, C., Shimomai, T., & Saito, A. (2002), 'Observations and modeling of 630 nm airglow and total electron content associated with traveling ionospheric disturbances over Shigaraki, Japan', *Earth, Planets and Space*, 54(1), pp. 45–56. doi:10.1186/BF03352420

Ogawa, T., Nishitani, N., Otsuka, Y., Shiokawa, K., Tsugawa, T., & Hosokawa, K. (2009), 'Medium-scale traveling ionospheric disturbances observed with the SuperDARN Hokkaido radar, all-sky imager, and GPS network and their relation to concurrent sporadic *E* irregularities', *Journal of Geophysical Research*, 114(A3), A03316. doi:10.1029/2008JA013893

Oliver, W. L., Fukao, S., Yamamoto, Y., Takami, T., Yamanaka, M. D., Yamamoto, M., et al. (1994), 'Middle and upper atmosphere radar observations of ionospheric density gradients produced by gravity wave packets', *Journal of Geophysical Research*, 99(A4), pp. 6,321–6,329. doi:10.1029/94JA00171

Oliver, W. L., Otsuka, Y., Sato, M., Takami, T., & Fukao, S. (1997), 'A climatology of *F* region gravity wave propagation over the middle and upper atmosphere radar', *Journal of Geophysical Research*, 102(A7), pp. 14,499–14,512. doi:10.1029/97JA00491

Otsuka, Y., Ogawa, T., Saito, A., Tsugawa, T., Fukao, S., & Miyazaki, S. (2002), 'A new technique for mapping of total electron content using GPS network in Japan', *Earth, Planets and Space*, 54(1), pp. 63–70. doi:10.1186/BF03352422

Otsuka, Y., Shiokawa, K., Ogawa, T., & Wilkinson, P. (2004), 'Geomagnetic conjugate observations of medium-scale traveling ionospheric disturbances at midlatitude using all-sky airglow imagers', *Geophysical Research Letters*, 31(15), L15803. doi:10.1029/2004GL020262

Park, J., Lühr, H., Lee, C., Kim, Y. H., Jee, G., & 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), pp. 2,187–2,196. doi:10.1002/2013JA019705

Parkinson, M. L., Breed, A. M., Dyson, P. L., & Morris, R. J. (1999), 'Application of the Dopplionogram to Doppler-sorted interferometry measurements of ionospheric drift velocity', *Radio Science*, 34(4), pp. 899–912. doi:10.1029/1999RS900029

Parkinson, M. L., Dyson, P. L., & Smith, P. R. (1997), 'Analysis of direction-of-arrival aliasing for MF/HF Doppler-sorted interferometry measurements of ionospheric drift', *Radio Science*, 32(3), pp. 999–1,009. doi:10.1029/97RS00051

Paul, A. K. (1985), '*F* region tilts and ionogram analysis', *Radio Science*, 20(4), pp. 959–971. doi:10.1029/RS020i004p00959

Paul, A. K., Wright, J. W., & Fedor, L. S. (1974), 'The interpretation of ionospheric radio drift measurements – VI. Angle-of-arrival and group path [echolocation] measurements from digitized ionospheric soundings: The group path vector', *Journal of Atmospheric and Terrestrial Physics*, 36(2), pp. 193–214. doi:10.1016/0021-9169(74)90040-3

Paul, K. S., Haralambous, H., Oikonomou, C., & Paul, A. (2019), 'Long-term aspects of nighttime spread F over a low mid-latitude European station', *Advances in Space Research*, 64(6), pp. 1,199–1,216. doi:10.1016/j.asr.2019.06.020

Paznukhov, V. V. (2004), *Atmospheric gravity waves/traveling ionospheric disturbances study with digital ionosondes and incoherent scatter radar*, PhD thesis, University of Massachusetts, Lowell, Massachusetts. Retrieved from ProQuest Dissertations & Theses Global (document ID 305172830).

Paznukhov, V. V., Galushko, V. G., & Reinisch, B. W. (2012), 'Digisonde observations of TIDs with frequency and angular sounding technique', *Advances in Space Research*, 49(4), pp. 700–710. doi:10.1016/j.asr.2011.11.012

Paznukhov, V., Altadill, D., Juan, J. M., & Blanch, E. (2020), 'Ionospheric Tilt Measurements: Application to Traveling Ionospheric Disturbances Climatology Study', *Radio Science*, 55(2), e2019RS007012. doi:10.1029/2019RS007012

Pederick, L. H., & Cervera, M. A. (2014), 'Semiempirical Model for Ionospheric Absorption based on the NRLMSISE-00 atmospheric model', *Radio Science*, 49(2), pp. 81–93. doi:10.1002/2013RS005274

Pederick, L. H., Cervera, M. A., & Harris, T. J. (2017), 'Interpreting Observations of Large-Scale Traveling Ionospheric Disturbances by Ionospheric Sounders', *Journal of Geophysical Research: Space Physics*, 122(12), pp. 12,556–12,569. doi:10.1002/2017JA024337

Perkins, F. (1973), 'Spread *F* and Ionospheric Currents', *Journal of Geophysical Resesarch*, 78(1), pp. 218–226. doi:10.1029/JA078i001p00218

Pezzopane, M., Pietrella, M., Pignatelli, A., Zolesi, B., & Cander, L. R. (2011), 'Assimilation of autoscaled data and regional and local ionospheric models as input sources for real-time 3-D International Reference Ionosphere modeling', *Radio Science*, 46(5), RS5009. doi:10.1029/2011RS004697

Pezzopane, M., & Scotto, C. (2004), 'Software for the automatic scaling of critical frequency *foF2* and *MUF(3000)F2* from ionograms applied at the Ionospheric Observatory of Gibilmanna', *Annals of Geophysics*, 47(6), pp. 1783–1790. doi:10.4401/ag-3375

Pezzopane, M., & Scotto, C. (2007), 'Automatic scaling of critical frequency *foF2* and *MUF(3000)F2*: A performance comparison between Autoscala and ARTIST 4.5 on Rome data', *Radio Science*, 42(4), RS4003. doi:10.1029/2006RS003581

Pezzopane, M., & Scotto, C. (2008), 'A method for automatic scaling of F1 critical frequencies from ionograms', *Radio Science*, 43(2), RS2S91. doi:10.1029/2007RS003723

Pfister, W. (1971), 'The wave-like nature of inhomogeneities in the *E*-region', *Journal of Atmospheric and Terrestrial Physics*, 33(7), pp. 999–1,025. doi:10.1016/0021-9169(71)90123-1

Pickering, L. W. (1975), 'The Calculation of Ionospheric Doppler Spread on HF Communication Channels', *IEEE Transactions on Communications*, 23(5), pp. 526–537. doi:10.1109/TCOM.1975.1092843

Picone, J. M., Hedin, A. E., Drob, D. P., & Aikin, A. C. (2002), 'NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues', *Journal of Geophysical Research*, 107(A12), 1468. doi:10.1029/2002JA009430

Pierce, J. A., & Mimno, H. R. (1940), 'The Reception of Radio Echoes From Distant Ionospheric Irregularities', *Physical Review*, 57(2), pp. 95–105. doi:10.1103/PhysRev.57.95

Piggott, W. R., & Rawer, K. (1978), *U.R.S.I. handbook of ionogram interpretation and reduction*, Report UAG-23A, 2nd edition (revised), World Data Centre A for Solar-Terrestrial Physics, NOAA, Boulder, Colorado. Retrieved from https://www.sws.bom.gov.au/IPSHosted/INAG/uag_23a/UAG_23A_indexed.pdf

Pillai, S. U. (1989), *Array Signal Processing*, Signal Processing and Digital Filtering, Springer-Verlag, New York, New York. doi:10.1007/978-1-4612-3632-0

Pitteway, M. L. V., & Hines, C. O. (1963), 'The viscous damping of atmospheric gravity waves', *Canadian Journal of Physics*, 41(12), pp. 1,935–1,948. doi:10.1139/p63-194

Pitteway, M. L. V., & Hines, C. O. (1965), 'The reflection and ducting of atmospheric acoustic-gravity waves', *Canadian Journal of Physics*, 43(12), pp. 2,222–2,245. doi:10.1139/p65-217

Poole, A. W. V., & Evans, G. P. (1985), 'Advanced sounding – 2. First results from an advanced chirp ionosonde', *Radio Science*, 20(6), pp. 1,617–1,623. doi:10.1029/RS020i006p01617

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (2007), *Numerical Recipes: The Art of Scientific Computing*, 3rd edition, Cambridge University Press, Cambridge, UK.

Psiaki, M. L. (2019), 'Ionosphere Ray Tracing of Radio-Frequency Signals and Solution Sensitivities to Model Parameters', *Radio Science*, 54(8), pp. 738–757. doi:10.1029/2019RS006792

Qian, L., Burns, A. G., Emery, B. A., Foster, B., Lu, G., Maute, A., et al. (2014), 'The NCAR TIE-GCM: A Community Model of the Coupled Thermosphere/Ionosphere System', in J. Huba, R. Schunk and G. Khazanov (ed.), *Modeling the Ionosphere–Thermosphere System*, American Geophysical Union, Washington, DC, ch. 7, pp. 73–83. doi:10.1002/9781118704417.ch7

Radicella, S. M., & Leitinger, R. (2001), 'The evolution of the DGR approach to model electron density profiles', *Advances in Space Research*, 27(1), pp. 35–40. doi:10.1016/S0273-1177(00)00138-1

Ratcliffe, J. A. (1956), 'Some Aspects of Diffraction Theory and their Application to the Ionosphere', *Reports on Progress in Physics*, 19(1), pp. 188–267. doi:10.1088/0034-4885/19/1/306

Ratcliffe, J. A. (1959), 'Ionizations and Drifts in the Ionosphere', *Journal of Geophysical Research*, 64(12), pp. 2,102–2,111. doi:10.1029/JZ064i012p02102

Ratovsky, K. G., Medvedev, A. V., Tolstikov, M. V., & Kushnarev, D. S. (2008), 'Case studies of height structure of TID propagation characteristics using cross-correlation analysis of incoherent scatter radar and DPS-4 ionosonde data', *Advances in Space Research*, 41(9), pp. 1,454–1,458. doi:10.1016/j.asr.2007.03.008

Rawer, K. (1993), *Wave Propagation in the Ionosphere*, Developments in Electromagnetic Theory and Applications, vol. 5, Springer Science+Business Media, Dordrecht, Netherlands.

Rehman, N., & Mandic, D. P. (2010), 'Multivariate empirical mode decomposition', *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 466(2,117), pp. 1,291–1,302. doi:10.1098/rspa.2009.0502

Reilly, M. H. (1985), 'Ionospheric true height profiles from oblique ionograms', *Radio Science*, 20(3), pp. 280–286. doi:10.1029/RS020i003p00280

Reinisch, B. W. (1986), 'New techniques in ground-based ionospheric sounding and studies', *Radio Science*, 21(3), pp. 331–341. doi:10.1029/RS021i003p00331

Reinisch, B., Galkin, I., Belehaki, A., Paznukhov, V., Huang, X., Altadill, D., et al. (2018), 'Pilot ionosonde network for identification of traveling ionospheric disturbances', *Radio Science*, 53(3), pp. 365–378. doi:10.1002/2017RS006263

Reinisch, B. W., Galkin, I. A., Khmyrov, G. M., Kozlov, A. V., Bibl, K., Lisysyan, I. A., et al. (2009), 'New Digisonde for research and monitoring applications', *Radio Science*, 44(1), RS0A24. doi:10.1029/2008RS004115

Reinisch, B. W., Gamache, R. R., Huang, X., & McNamara, L. F. (1988), 'Real-time electron density profiles from ionograms', *Advances in Space Research*, 8(4), pp. 63–72. doi:10.1016/0273-1177(88)90208-6

Reinisch, B. W., Haines, D. M., Bibl, K., Galkin, I., Huang, X., Kitrosser, D. F., et al. (1997), 'Ionospheric sounding in support of over-the-horizon radar', *Radio Science*, 32(4), pp. 1,681–1,694. doi:10.1029/97RS00841

Reinisch, B. W., & Huang, X. (1983), 'Automatic calculation of electron density profiles from digital ionograms. 3. Processing of bottomside ionograms', *Radio Science*, 18(3), pp. 477–492. doi:10.1029/RS018i003p00477

Reinisch, B. W., Huang, X., Galkin, I. A., Paznukhov, V., & 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), pp. 1,054–1,062. doi:10.1016/j.jastp.2005.01.009

Reinisch, B. W., Scali, J. L., & Haines, D. M. (1998), 'Ionospheric drift measurements with ionosondes', *Annals of Geophysics*, 41(5–6), pp. 695–702. doi:10.4401/ag-3812

Rice, D. W. (1973), 'Phase characteristics of ionospherically propagated radio waves', *Nature Physical Science*, 244(136), pp. 86–88.

Richmond, A. D. (1978), 'Gravity Wave Generation, Propagation, and Dissipation in the Thermosphere', *Journal of Geophysical Research*, 83(A9), pp. 4,131–4,145. doi:10.1029/JA083iA09p04131

Ridley, A. J., Deng, Y., & Toth, G. (2006), 'The global ionosphere–thermosphere model', *Journal of Atmospheric and Solar-Terrestrial Physics*, 68(8), pp. 839–864. doi:10.1016/j.jastp.2006.01.008

Rietveld, M. T., Wright, J. W., Zabotin, N., & Pitteway, M. L. V. (2008), 'The Tromsø dynasonde', *Polar Science*, 2(1), pp. 55–71. doi:10.1016/j.polar.2008.02.001

Rishbeth, H. (1988), 'Basic physics of the ionosphere: a tutorial review', *Journal of the Institution of Electronic and Radio Engineers*, 58(6S), pp. S207–S223. doi:10.1049/jiere.1988.0060

Rishbeth, H. (2000), 'The equatorial F-layer: progress and puzzles', *Annales Geophysicae*, 18(7), pp. 730–739. doi:10.1007/s00585-000-0730-6

Rishbeth, H., & Garriott, O. K. (1969), *Introduction to Ionospheric Physics*, International Geophysics Series, vol. 14, Academic Press, New York, New York.

Rishbeth, H., & Mendillo, M. (2001), 'Patterns of F2-layer variability', *Journal of Atmospheric and Solar-Terrestrial Physics*, 63(15), pp. 1,661–1,680. doi:10.1016/S1364-6826(01)00036-0

Rogers, N. C., Mitchell, C. N., Heaton, J. A. T., Cannon, P. S., & Kersley, L. (2001), 'Application of radio tomographic imaging to HF oblique incidence ray tracing', *Radio Science*, 36(6), pp. 1,591–1,598. doi:10.1029/2000RS002505

Rogers, N. C., Warrington, E. M., & Jones, T. B. (2003), 'Oblique ionogram features associated with off-great circle HF propagation at high and sub-auroral latitudes', *IEE Proceedings – Microwaves, Antennas and Propagation*, 150(4), pp. 295–300. doi:10.1049/ip-map:20030552

Ross, W., Bramley, E. N., & Ashwell, G. E. (1951), 'A phase-comparison method of measuring the direction of arrival of ionospheric radio waves', *Proceedings of the IEE - Part III: Radio and Communication Engineering*, 98(54), pp. 294–302. doi:10.1049/pi-3.1951.0061

Saito, A., Iyemori, T., Blomberg, L. G., Yamamoto, M., & Takeda, M. (1998), 'Conjugate observations of the mid-latitude electric field fluctuations with the MU radar and the Freja satellite', *Journal of Atmospheric and Solar-Terrestrial Physics*, 60(1), pp. 129–140. doi:10.1016/S1364-6826(97)00094-1

Saryo, T., Takeda, M., Araki, T., Sato, T., Tsuda, T., Fukao, S., & Kato, S. (1989), 'A Midday Bite-Out Event of the *F*2-Layer Observed by MU Radar', *Journal of Geomagnetism and Geoelectricity*, 41(8), pp. 727–734. doi:10.5636/jgg.41.727

Scali, J. L., Reinisch, B. W., Heinselman, C. J., and Bullett, T. W. (1995), 'Coordinated digisonde and incoherent scatter radar *F* region drift measurements at Sondre Stromfjord', *Radio Science*, 30(5), pp. 1,481–1,498. doi:10.1029/95RS01730

Schelleng, J. C., Burrows, C. R., & Ferrell, E. B. (1933), 'Ultra-short-wave propagation', *Proceedings of the IRE*, 21(3), pp. 427–463. doi:10.1109/JRPROC.1933.227639

Scherliess, L., Schunk, R. W., Sojka, J. J., & Thompson, D. C. (2004), 'Development of a physics-based reduced state Kalman filter for the ionosphere', 39(1), RS1S04. doi:10.1029/2002RS002797

Scherliess, L., Schunk, R. W., Sojka, J. J., Thompson, D. C., & Zhu, L. (2006), 'Utah State University Global Assimilation of Ionospheric Measurements Gauss-Markov Kalman filter model of the ionosphere: Model description and validation', *Journal of Geophysical Research*, 111(A11), A11315. doi:10.1029/2006JA011712

Schmidt, R. O. (1986), 'Multiple Emitter Location and Signal Parameter Estimation', *IEEE Transactions on Antennas and Propagation*, AP-34(3), pp. 276–280. doi:10.1109/TAP.1986.1143830

Schreiner, W. S., Sokolovskiy, S. V., Rocken, C., & Hunt, D. C. (1999), 'Analysis and validation of GPS/MET radio occultation data in the ionosphere', *Radio Science*, 34(4), pp. 949–966. doi:10.1029/1999RS900034

Schulkin, M. (1952), 'Average Radio-Ray Refraction in the Lower Atmosphere', *Proceedings of the IRE*, 40(5), pp. 554–561. doi:10.1109/JRPROC.1952.273820

Schunk, R. W., & Nagy, A. F. (2009), *Ionospheres: Physics, Plasma Physics, and Chemistry*, 2nd edition, Cambridge Atmospheric and Space Science Series, Cambridge University Press, Cambridge, UK. doi:10.1017/CBO9780511635342

Schunk, R. W., Scherliess, L., Sojka, J. J., Thompson, D. C., Anderson, D. N., Codrescu, M., et al. (2004), 'Global Assimilation of Ionospheric Measurements (GAIM)', *Radio Science*, 39(1), RS1S02. doi:10.1029/2002RS002794

Scotto, C. (2009), 'Electron density profile calculation technique for Autoscala ionogram analysis', *Advances in Space Research*, 44(6), pp. 756–766. doi:10.1016/j.asr.2009.04.037

Scotto, C., & Pezzopane, M. (2007), 'A method for automatic scaling of sporadic *E* layers from ionograms', *Radio Science*, 42(2), RS2012. doi:10.1029/2006RS003461

Sherrill, W. M., & Brown, B. C. (1997), 'Azimuth spread of two-hop ionospheric modes', *Radio Science*, 32(5), pp. 1,875–1,880. doi:10.1029/97RS01267

Sherrill, W. M., & Smith, G. A. (1977), 'Directional dispersion of sporadic-*E* observations between 9 and 14 MHz', *Radio Science*, 12(5), pp. 773–778. doi:10.1029/RS012i005p00773

Shiokawa, K., Ihara, C., Otsuka, Y., & Ogawa, T. (2003), 'Statistical study of nighttime medium-scale traveling ionospheric disturbances using midlatitude airglow images', *Journal of Geophysical Research*, 108(A1), 1052. doi:10.1029/2002JA009491

Shiokawa, K., Otsuka, Y., & Ogawa, T. (2009), 'Propagation characteristics of nighttime mesospheric and thermospheric waves observed by optical mesosphere thermosphere imagers at middle and low latitudes', *Earth, Planets and Space*, 61(4), pp. 479–491. doi:10.1186/BF03353165

Shiokawa, K., Otsuka, Y., Tsugawa, T., Ogawa, T., Saito, A., Ohshima, K., et al. (2005), 'Geomagnetic conjugate observation of nighttime medium-scale and large-scale traveling ionospheric disturbances: FRONT3 campaign', *Journal of Geophysical Research*, 110(A5), A05303. doi:10.1029/2004JA010845

Singh, H., Sneha, H. L., & Jha, R. M. (2013), 'Mutual Coupling in Phased Arrays: A Review', *International Journal of Antennas and Propagation*, 2013, 348123. doi:10.1155/2013/348123

Sinnott, D. H. (2015), 'Over the horizon radar down-under', *Proceedings of the IEEE Radar Conference 2015*, Arlington, Virginia, pp. 1,761–1,764. doi:10.1109/RADAR.2015.7131284

Smith, E. K., Jr., & Weintraub, S. (1953), 'The Constants in the Equation for Atmospheric Refractive Index at Radio Frequencies', *Proceedings of the IRE*, 41(8), pp. 1,035–1,037. doi:10.1109/JRPROC.1953.274297

Solomon, I. S. D. (1998), *Over-the-Horizon Radar Array Calibration*, PhD thesis, University of Adelaide, Adelaide, South Australia, Australia.

Space Weather Services (2018), *Monthly Sunspot Numbers*. Accessed 7 June 2018 from https://www.sws.bom.gov.au/Solar/1/6

Space Weather Services (2020), *SWS Software Introduction*. Accessed 28 April 2020 from http://www.sws.bom.gov.au/asaps

Sultan, P. J. (1996), 'Linear theory and modeling of the Rayleigh-Taylor instability leading to the occurrence of equatorial spread *F*', *Journal of Geophysical Research*, 101(A12), pp. 26,875–26,891. doi:10.1029/96JA00682

Sweeney, L. E., Jr. (1970), *Spatial Properties of Ionospheric Radio Propagation as Determined With Half-Degree Azimuthal Resolution*, Technical Report No. 155, SU-SEL-70-034, Stanford Electronics Laboratories, Stanford, California.

Swingler, D. N., & Walker, R. S. (1989), 'Line-array beamforming using linear prediction for aperture interpolation and extrapolation', *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(1), pp. 16–30. doi:10.1109/29.17497

Taylor, M. J., Ryan, E. H., Tuan, T. F., & Edwards, R. (1993), 'Evidence of Preferential Directions for Gravity Wave Propagation Due to Wind Filtering in the Middle Atmosphere', *Journal of Geophysical Research*, 98(A4), pp. 6,047–6,057. doi:10.1029/92JA02604

Tedd, B. L., Morgan, M. G., & Ballard, K. A. (1984), 'The Height Dependence of TID and Gravity Wave Parameters', *Journal of Geophysical Research*, 89(A10), pp. 9,023–9,033. doi:10.1029/JA089iA10p09023

Thébault, E., Finlay, C. C., Beggan, C. D., Alken, P., Aubert, J., Barrois, O., et al. (2015), 'International Geomagnetic Reference Field: the 12th generation', *Earth, Planets and Space*, 67(1), 79. doi:10.1186/s40623-015-0228-9

Thome, G. D. (1964), 'Incoherent Scatter Observations of Traveling Ionospheric Disturbances', *Journal of Geophysical Research*, 69(19), pp. 4,047–4,049. doi:10.1029/JZ069i019p04047

Thome, G. (1968), 'Long-Period Waves Generated in the Polar Ionosphere during the Onset of Magnetic Storms', *Journal of Geophysical Research*, 73(19), pp. 6,319–6,336. doi:10.1029/JA073i019p06319

Torrence, C., & Compo, G. P. (1998), 'A Practical Guide to Wavelet Analysis', *Bulletin of the American Meteorological Society*, 79(1), pp. 61–78. doi:10.1175/1520-0477(1998)079<0061:APGTWA>2.0.CO;2

Tsagouri, I., Zolesi, B., Belehaki, A., & Cander, L. R. (2005), 'Evaluation of the performance of the real-time updated simplified ionospheric regional model for the European area', *Journal of Atmospheric and Solar-Terrestrial Physics*, 67(12), pp. 1,137–1,146. doi:10.1016/j.jastp.2005.01.012

Tsugawa, T., Otsuka, Y., Coster, A. J., & Saito, A. (2007), 'Medium-scale traveling ionospheric disturbances detected with dense and wide TEC maps over North America', *Geophysical Research Letters*, 34(22), L22101. doi:10.1029/2007GL031663

Turley, M. D. E. (2008), 'Signal processing techniques for maritime surveillance with skywave radar', *Proceedings of the IEEE International Radar Conference 2008*, Adelaide, South Australia, Australia, pp. 241–246. doi:10.1109/RADAR.2008.4653925

Turley, M. D. E., Heitmann, A. J., & Gardiner-Garden, R. S. (2019), 'Ionogram RFI rejection using an auto-regressive interpolation process', *Radio Science*, 54(1), pp. 135–150. doi:10.1029/2018RS006683

Turley, M. D. E., & Voigt, S. M. (1992), 'The use of a hybrid AR/classical spectral analysis technique with application to HF radar', *Proceedings of the International Symposium on Signal Processing and its Applications (ISSPA-1992)*, Gold Coast, Queensland, Australia, pp. 467–470.

Unewisse, A. M., Cervera, M. A., Pederick, L. H., Harris, T. J., & Cool, A. D. (2016), 'Airglow Observations from ELOISE', *Proceedings of the 16th Australian Space Research Conference 2016*, Melbourne, Victoria, Australia, pp. 139–150. Retrieved from https://marssociety.org.au/sites/default/files/16asrc%20proceedings.pdf

Unewisse, A. M., & Cool, A. D. (2016), 'Imagers for Ionospheric Airglow Observations', *Proceedings of the 16th Australian Space Research Conference 2016*, Melbourne, Victoria, Australia, pp. 127–138. Retrieved from https://marssociety.org.au/sites/default/files/16asrc%20proceedings.pdf

Vadas, S. L., & Fritts, D. C. (2001), 'Gravity Wave Radiation and Mean Responses to Local Body Forces in the Atmosphere', *Journal of the Atmospheric Sciences*, 58(16), pp. 2,249–2,279. doi:10.1175/1520-0469(2001)058<2249:GWRAMR>2.0.CO;2

Vadas, S. L., & Fritts, D. C. (2005), 'Thermospheric responses to gravity waves: Influences of increasing viscosity and thermal diffusivity', *Journal of Geophysical Research*, 110(D15), D15103. doi:10.1029/2004JD005574

Vadas, S. L., & Fritts, D. C. (2006), 'Influence of solar variability on gravity wave structure and dissipation in the thermosphere from tropospheric convection', *Journal of Geophysical Research*, 111(A10), A10S12. doi:10.1029/2005JA011510

Vadas, S. L., Fritts, D. C., & Alexander, M. J. (2003), 'Mechanism for the Generation of Secondary Waves in Wave Breaking Regions', *Journal of the Atmospheric Sciences*, 60(1), pp. 194–214. doi:10.1175/1520-0469(2003)060<0194:MFTGOS>2.0.CO;2

Valladares, C. E., & Sheehan, R. (2016), 'Observations of conjugate MSTIDs using networks of GPS receivers in the American sector', *Radio Science*, 51(9), pp. 1,470–1,488. doi:10.1002/2016RS005967

Van Trees, H. L. (2002), *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*, John Wiley & Sons Inc., New York, New York. doi:10.1002/0471221104

Västberg, A., & Lundborg, B. (1996), 'Signal intensity in the geometrical optics approximation for the magnetized ionosphere', *Radio Science*, 31(6), pp. 1,579–1,588. doi:10.1029/96RS02630

Västberg, A., & Lundborg, B. (1997), 'Ray tracing studies of multipath fading in the ionospheric HF channel', *Proceedings of the 7th International Conference on HF Radio Systems and Techniques*, Nottingham, UK, pp. 30–34. doi:10.1049/cp:19970755

Vertogradov, G. G., Uryadov, V. P., Vertogradova, E. G., Vertogradov, V. G., & Kubatko, S. V. (2013), 'Chirp ionosonde-radiodirection finder as a new tool for studying the ionosphere and radio-wave propagation', *Radiophysics and Quantum Electronics*, 56(5), pp. 259–275. doi:10.1007/s11141-013-9431-3

Wagner, L. S., Goldstein, J. A., & Meyers, W. D. (1988), 'Wideband probing of the transauroral HF channel: Solar minimum', *Radio Science*, 23(4), pp. 555–568. doi:10.1029/RS023i004p00555

Waldock, J. A., & Jones, T. B. (1984), 'The effects of neutral winds on the propagation of medium-scale atmospheric gravity waves at mid-latitudes', *Journal of Atmospheric and Terrestrial Physics*, 46(3), pp. 217–231. doi:10.1016/0021-9169(84)90149-1

Waldock, J. A., & Jones, T. B. (1987), 'Source regions of medium scale travelling ionospheric disturbances observed at mid-latitudes', *Journal of Atmospheric and Terrestrial Physics*, 49(2), pp. 105–114. doi:10.1016/0021-9169(87)90044-4

Warrington, E. M. (1996), 'Measurement of the angular spread of HF signals reflected from the high latitude ionosphere by means of super-resolution direction finding', *Proceedings of the IEE Colloquium on Remote Sensing of the Propagation Environment*, London, UK. doi:10.1049/ic:19961175

Waynick, A. H. (1975), 'The early history of ionospheric investigations in the United States', *Philosophical Transactions of the Royal Society of London A: Mathematical and Physical Sciences*, 280(1,293), pp. 11–25. Retrieved from www.jstor.org/stable/74681

Weber, E. J., Buchau, J., Eather, R. H., & Mende, S. B. (1978), 'North-South Aligned Equatorial Airglow Depletions', *Journal of Geophysical Research*, 83(A2), pp. 712–716. doi:10.1029/JA083iA02p00712

Wheadon, N. S., Whitehouse, J. C., Milsom, J. D., & Herring, R. N. (1994), 'Ionospheric modelling and target coordinate registration for HF sky-wave radars', *Proceedings of the Sixth International Conference on HF Radio Systems and Techniques*, York, UK, pp. 258–266. doi:10.1049/cp:19940504

Whitehead, J. D. (1989), 'Recent work on mid-latitude and equatorial sporadic-E', *Journal of Atmospheric and Terrestrial Physics*, 51(5), pp. 401–424. doi:10.1016/0021-9169(89)90122-0

Whitehead, J. D., From, W. R., Jones, K. L., & Monro, P. E. (1983), 'Measurement of movements in the ionosphere using radio reflections', *Journal of Atmospheric and Terrestrial Physics*, 45(5), pp. 345–351. doi:10.1016/S0021-9169(83)80039-7

World Meteorological Organization (2012), *Guide to Meteorological Instruments and Methods of Observations*, 2008 edition, updated in 2010, WMO-No. 8, World Meteorological Organization, Geneva, Switzerland. Retrieved from https://library.wmo.int/pmb_ged/wmo_8_en-2012.pdf

Wright, J. W., Argo, P. E., & Pitteway, M. L. V. (1996), 'On the radiophysics and geophysics of ionogram spread *F*', *Radio Science*, 31(2), pp. 349–366. doi:10.1029/95RS03104

Wright, J. W., & Kressman, R. I. (1983), 'First bistatic oblique-incidence ionograms between digital ionosondes', *Radio Science*, 18(4), pp. 608–614. doi:10.1029/RS018i004p00608

Wright, J. W., & Pitteway, M. L. V. (1979), 'Real-time data acquisition and interpretation capabilities of the Dynasonde – 1. Data acquisition and real-time display', *Radio Science*, 14(5), pp. 815–825. doi:10.1029/RS014i005p00815

Wright, J. W., & Pitteway, M. L. V. (1982), 'Data Processing for the Dynasonde: The Dopplionogram', *Journal of Geophysical Research*, 87(A3), pp. 1,589–1,598. doi:10.1029/JA087iA03p01589

Wright, J. W., & Pitteway, M. L. V. (1994), 'High-resolution vector velocity determinations from the dynasonde', *Journal of Atmospheric and Terrestrial Physics*, 56(8), pp. 961–977. doi:10.1016/0021-9169(94)90157-0

Yau, K. S. B. (2003), 'A simple polarisation fading model for HF propagation in the ionosphere', *Proceedings of the 9th International Conference on HF Radio Systems and Techniques*, Bath, UK, pp. 131–135. doi:10.1049/cp:20030443

Yeh, K. C., & Liu, C. H. (1974), 'Acoustic-Gravity Waves in the Upper Atmosphere', *Reviews of Geophysics and Space Physics*, 12(2), pp. 193–216. doi:10.1029/RG012i002p00193

Yue, X., Schreiner, W. S., Pedatella, N., Anthes, R. A., Mannucci, A. J., Straus, P. R., & Liu, J.-Y. (2014), 'Space Weather Observations by GNSS Radio Occultation: From FORMOSAT-3/COSMIC to FORMOSAT-7/COSMIC-2', *Space Weather*, 12(11), pp 616–621. doi:10.1002/2014SW001133

Zabotin, N. A., Wright, J. W., & Zhbankov, G. A. (2006), 'NeXtYZ: Three-dimensional electron density inversion for dynasonde ionograms', *Radio Science*, 41(6), RS6S32. doi:10.1029/2005RS003352

Zawdie, K. (2015), *3D Ionospheric Effects on HF Propagation and Heating*, PhD thesis, University of Maryland, College Park, Maryland.

Zawdie, K. A., Drob, D. P., Huba, J. D., & Coker, C. (2016), 'Effect of time-dependent 3-D electron density gradients on high angle of incidence HF radiowave propagation', *Radio Science*, 51(7), pp. 1,131–1,141. doi:10.1002/2015RS005843

Zawdie, K. A., Drob, D. P., Siskind, D. E., & Coker, C. (2017), 'Calculating the absorption of HF radio waves in the ionosphere', *Radio Science*, 52(6), pp. 767–783. doi:10.1002/2017RS006256

Zolesi, B., Belehaki, A., Tsagouri, I., & Cander, L. R. (2004), 'Real-time updating of the Simplified Ionospheric Regional Model for operational applications', *Radio Science*, 39(2), RS2011. doi:10.1029/2003RS002936

Zolesi, B., Cander, L. R., & De Franceschi, G. (1993), 'Simplified ionospheric regional model for telecommunication applications', *Radio Science*, 28(4), pp. 603–612. doi:10.1029/93RS00276

Zolesi, B., Cander, L. R., & De Franceschi, G. (1996), 'On the potential applicability of the simplified ionospheric regional model to different midlatitude areas', *Radio Science*, 31(3), pp. 547–552. doi:10.1029/95RS03817

Zuccheretti, E., Tutone, G., Sciacca, U., Bianchi, C., & Arokiasamy, B. J. (2003), 'The new AIS-INGV digital ionosonde', *Annals of Geophysics*, 46(4), pp. 647–659. doi:10.4401/ag-4377