CHALMERS
UNIVERSITY OF TECHNOLOGY

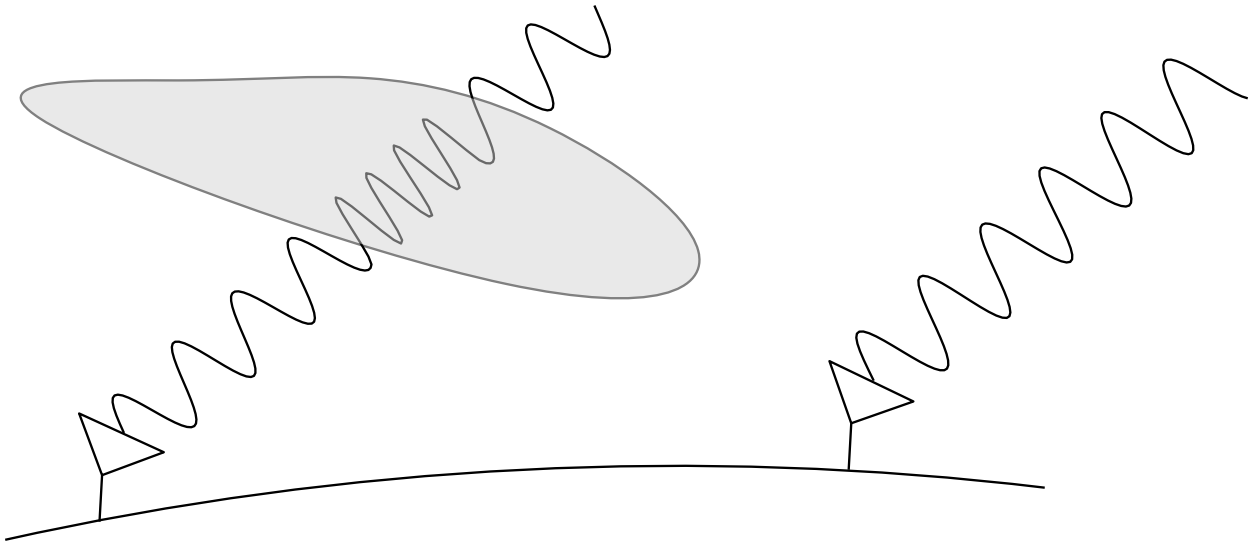

Modelling and Simulation of Tropospheric Water Vapour With Gaussian Random Fields

Time dependence beyond the frozen flow hypothesis

Master's thesis in Applied Physics

HENRIK GINGSJÖ

Department of Space, Earth and Environment
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2018

Modelling and Simulation of Tropospheric Water Vapour With Gaussian Random Fields

Time dependence beyond the frozen flow hypothesis

HENRIK GINGSJÖ

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Space, Earth and Environment
Chalmers University of Technology
Gothenburg, Sweden 2018

Modelling and Simulation of Tropospheric Water Vapour With Gaussian Random Fields
Time dependence beyond the frozen flow hypothesis
HENRIK GINGSJÖ

Supervisor: Thomas Hobiger, Department of Space, Earth and Environment
Examiner: Gunnar Elgered, Department of Space, Earth and Environment

Department of Space, Earth and Environment
Chalmers University of Technology
SE-412 96 Gothenburg

Cover: Sketch of a two antenna radio interferometer. The left antenna is affected by more atmospheric delay than the right.

Typeset in LaTeX on 2018-06-07
Gothenburg, Sweden 2018

Modelling and Simulation of Tropospheric Water Vapour With Gaussian Random Fields
Time dependence beyond the frozen flow hypothesis
HENRIK GINGSJÖ
Department of Space, Earth and Environment
Chalmers University of Technology

Abstract

One of the major sources of error in Very Long Baseline Interferometry (VLBI) is signal delay due to tropospheric water vapour. Turbulent convection makes it inherently unpredictable and it must therefore be measured directly or modelled stochastically. In particular, realizations of delay signals are necessary to simulate the performance of existing and future VLBI networks which, in turn, is needed to optimize them and reduce errors.

In previous work, modelling of tropospheric delay has been performed only on the spatial structure of refractivity through phenomenological second order statistics derived from Kolmogorov theory. Time dependence has been introduced exclusively through the frozen-flow hyporthesis.

In this thesis, refractivity fields are modelled as Gaussian random fields. Efficient software is implemented to generate realizations of such fields sampled on a 3D grid. To achieve realistic time evolution of such gridded fields, it turns out to be both necessary and natural to introduce intrinsic time dependence beyond the frozen-flow hypothesis. Such time dependence can easily be made compatible with the temporal structure of Kolmogorov turbulence.

The novel contributions of this thesis are methods of obtaining two kinds of time dependence for refractivity fields beyond the frozen-flow hypothesis. Firstly: Intrinsic time dependence compatible with Kolmogorov theory. Secondly: Translation by horizontal wind with arbitrary height and time dependence. The latter may provide a more realistic description of the planetary boundary layer which has strong wind shear and contains about 15 % of the total water vapour; corresponding to delays of several centimetres.

Keywords: tropospheric turbulence, wet tropospheric delay, frozen-flow hypothesis, Gaussian random fields, FFT

Contents

Chapter 1

Introduction

All Earth-based astronomy is influenced by the atmosphere. It is transparent essentially only to visible light and microwaves and therefore prevents observations at other wavelengths. But even at wavelengths where air is transparent, fluctuations in the refractive index is problematic. In optical telescopes this is the resolution-limiting factor [1] and it is one of the most significant error sources in Very Long Baseline Interferometry (VLBI) [2] and Global Navigation Satellite Systems (GNSS).

VLBI is one of the most precise astronomical techniques. For in-depth information and a large number of references, see the comprehensive review by Sovers et al. from 1998 [2]. A greatly simplified description follows. VLBI consists of using multiple radio telescopes, thousands of kilometres apart, to simultaneously record microwave signals from the same astronomical radio-source[1]. Next, the recorded signals are cross-correlated in order to obtain differences in arrival time between telescopes to ~10 ps uncertainty. The arrival time differences are then used for tasks such as maintaining the celestial and terrestrial reference frames [3, 4] (see also `www.iers.org`), measuring continental drift, imaging astronomical radio-sources and testing the validity of general relativity [2]. An interesting combination of the last two is an ongoing project [5] that aims to resolve the event horizon of the black hole at the centre of our galaxy (see also `https://eventhorizontelescope.org`).

However, since radio signals are delayed by varying amounts by the atmosphere, the differences in arrival time do not correspond exactly to the distances between telescopes. The atmospheric delay has three main components:

1. Dispersion in the ionosphere. Has a strong frequency dependence ($\propto f^{-2}$) and can therefore be accurately compensated for using simultaneous observations at two frequencies. Total zenith ionospheric delay[2]: a few metres[3] at single digit GHz frequencies, but it is highly dependent on frequency and solar radiation.
2. Delay due non-polar gases (which forms the majority of the atmosphere). Can be accurately estimated using the air pressure at the observation site. Total *zenith hydrostatic delay* (ZHD): $\approx 2\,\mathrm{m}$.

[1] Often radio bright, extragalactic objects called quasars that are so far away that they appear stationary and pointlike.

[2] The equivalent time delay is obtained by dividing by the speed of light.

[3] See [6] for details. For example, $10\,\mathrm{TECU} = 10^{17}\,\mathrm{m}^{-2}$, results in a 4 m delay at 1 GHz.

3. Delay due to tropospheric water vapour. Cannot be well estimated by surface meteorological observables and fluctuates on all timescales between seconds and days due to turbulence. Total *zenith wet delay* (ZWD): about 0–0.5 m depending on local weather and climate conditions.

Clearly, to reach centimetre and millimetre accuracy, all three effects must be compensated for in some manner. As noted, there are simple methods to accurately account for the first two, but not for the third. Furthermore, in a 2010 simulation study [7] of a state of the art VLBI network, the fluctuations in tropospheric water vapour were found to be the largest error source for geodetic applications.

The aim of this thesis is to develop software that can stochastically simulate wet tropospheric delays in a manner consistent with Kolmogorov turbulence theory. This is achieved by modelling refractivity fields as quasi-homogeneous Gaussian random fields which can be accurately and efficiently realized on regular grids using a standard FFT-based algorithm. The delay along a given signal path is found by numerical integration over the grid.

This is an alternate approach to that described in [8] which is used in the previously cited VLBI simulation [7], where a covariance matrix for a sampled delay signal is first computed by numerical integration. After a matrix inversion (Cholesky decomposition), realizations of the stochastic process can be generated from a matrix multiplication of a vector of independent normal distributed numbers. A very similar method seems to have been developed independently [9], with GPS as the indended application.

Simulating the refractivity distribution directly on a grid enables the structure models to be tested in greater detail. In particular, time evolution can be done in a more general manner than with the frozen-flow assumption. Other features include:

- Computation time scales linearly with number of timesteps and concurrent signals.
- Could conceivably be combined with ray tracing tools such as [10].
- Can be used for efficient stochastic simulation of water vapour radiometer signals.
- Enables the atmospheric effects to be directly visualized.

However, the computation of covariances for use in data filtering is expensive and inaccurate compared to more direct methods.

Chapter 2

Atmospheric structure

To realistically model the atmosphere, some understanding of is needed. How high is it? How fast does it move? Where is the water vapour located? How much and how fast does it fluctuate? What measurement methods are available? The purpose of this chapter is to answer those questions.

For a more thorough and authoritative description of the atmosphere see a textbook such as [11]. Here only the local structure of the lower atmosphere will be considered. Up to heights of about 10 km and horizontal distances at most an order of magnitude larger.

2.1 Equilibrium descriptions

In the absence of unusual weather, the atmosphere is slowly varying and it can be assumed to be in equilibrium as first approximation. This enables the vertical structure to be reasonably well understood. Hydrostatic balance dictates

$$\frac{\mathrm{d}p}{\mathrm{d}z} = -\rho g, \tag{2.1}$$

where z is the height coordinate, p is pressure, ρ is air density and g is the gravitational acceleration (approximately constant with height). Additionally, a constitutive relation between pressure and density is needed. The ideal gas law is a good approximation

$$p = \frac{R}{m_{\text{eff}}} T\rho, \tag{2.2}$$

where T is the absolute temperature, R is the molar gas constant and $m_{\text{eff}} = \sum_i f_i m_i$ is the effective molar mass. Here f_i is the molar fraction of component i and m_i is the corresponding molar mass.

The temperature remains a free parameter that can be modelled in various ways. The simplest assumption is a constant temperature, in which case the solution of (2.1) is an exponentially decreasing pressure

$$p(z) = p_0 \mathrm{e}^{-z/H}, \tag{2.3}$$

where $H = \frac{RT}{gm_{\text{eff}}}$. Using $g = 9.8\,\mathrm{N/kg}$, $m_{\text{eff}} = 29\,\mathrm{g/mol}$ (80 % N_2, 20 % O_2) and $T = 270\,\mathrm{K} \approx 0\ ^\circ\mathrm{C}$, the scale height becomes $H = 7.9\,\mathrm{km}$.

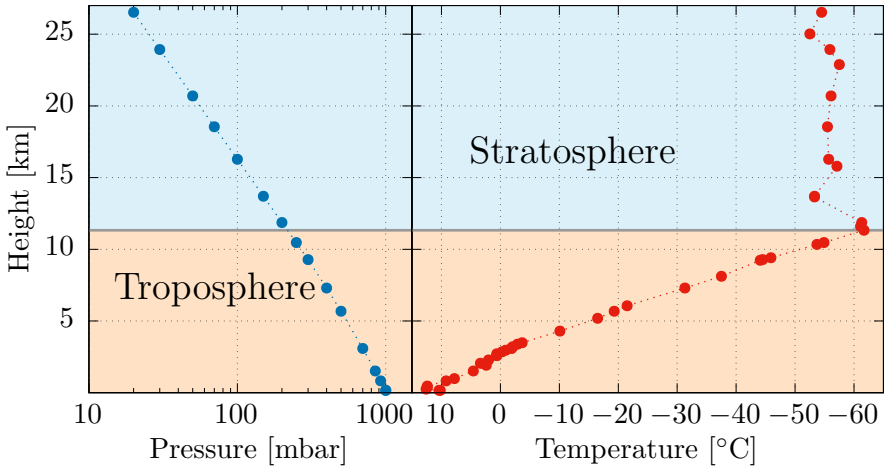

Figure 2.1: Pressure and temperature height profiles measured by a radiosonde launched from Landvetter Airport on 2016-09-23. The tropopause is indicated by the horizontal line at 11.3 km.

However, since the atmosphere is transparent to much of solar radiation, it is not heated directly by the sun, but indirectly by the ground. This makes the lowest lying air buoyant and mixes the atmosphere. Since the air is locally in equilibrium, we can expect the atmosphere to be adiabatic. This implies that the temperature and pressure profiles are such that if an air parcel is moved through the atmosphere while allowed to expand adiabatically to the surrounding pressure, its temperature will match the new surroundings.

In an adiabatic atmosphere the quantity $p^{1-\gamma}T^{\gamma}$ is constant, where $\gamma \approx 7/5$ for air (since it consists mainly of diatomic molecules). So $T = T_0 \left(p_0/p\right)^{1/\gamma - 1}$. Plugging this into (2.2) and (2.1) results in a power law for the pressure

$$p(z) = p_0 \left(1 - \frac{\gamma - 1}{\gamma} z/H\right)^{\gamma/(\gamma-1)} \tag{2.4}$$

and a linear temperature relation

$$T(z) = T_0(1 - \frac{\gamma - 1}{\gamma} z/H), \tag{2.5}$$

where H is the same as above with T replaced by the surface temperature T_0. The predicted temperature lapse rate is $-11\,\mathrm{K/km}$. Interestingly this predicts a sharp end of the atmosphere at height $\frac{\gamma}{\gamma-1}H = 3.5 \cdot H \approx 27\,\mathrm{km}$ (since the pressure and temperature are both zero there). That does not happen in reality.

The real lower atmosphere shows both behaviours. This can be seen in Figure 2.1, where the height profiles of temperature and pressure measured by a radiosonde (carried by weather balloon) are shown. In the lowest part (called the troposphere), the temperature drops approximately linearly with height. However, in this case with a lapse rate of $-7\,\mathrm{K/km}$; significantly lower than predicted. This means that the atmosphere is not quite adiabatic. The higher up air is positively buoyant which stabilizes the atmosphere. A more transparent way of showing this is to use

the *potential temperature*, which is the temperature an air parcel would get if it were adiabatically compressed to a reference pressure.

At the tropopause (by definition), the temperature suddenly stops decreasing and becomes constant. This is the boundary to the stratosphere, which is heated from above by ultraviolet radiation (to which air is opaque). The potential temperature therefore increases very rapidly with height, preventing significant vertical stirring.

2.2 Water vapour

The picture is complicated by the existence of water vapour in varying amounts. Figure 2.2 shows typical water vapour density profiles in the troposphere over western Sweden. The amount of water is mainly determined by the temperature due to the strong temperature dependence of its vapour pressure ($\sim \mathrm{e}^{-T_0/T}$). For this reason, most of the water exists in the lower troposphere where the temperature is highest. Furthermore, the amount of water varies throughout the year due to the variation of temperature. In the northern hemisphere, the lowest amount typically occurs around January–February and the maximal amount around July–August.

Regardless of the season, most of the water vapour (about 95 %) is contained in the lowest 5 km of the troposphere. Therefore, a reasonable height for the simulation volume is about 5 km. This also sets the necessary width of the volume. If observations are needed down to elevations of 30°, then a the volume must be at least $5\,\mathrm{km}/\tan(30°) \approx 10\,\mathrm{km}$ wide. Furthermore, this might be needed both "forwards" and "backwards" increasing the necessary width to 20 km. Finally, the simulation method used in this thesis generates spatially periodic fields. Hence, 40 km is advisable to avoid artificial correlations.

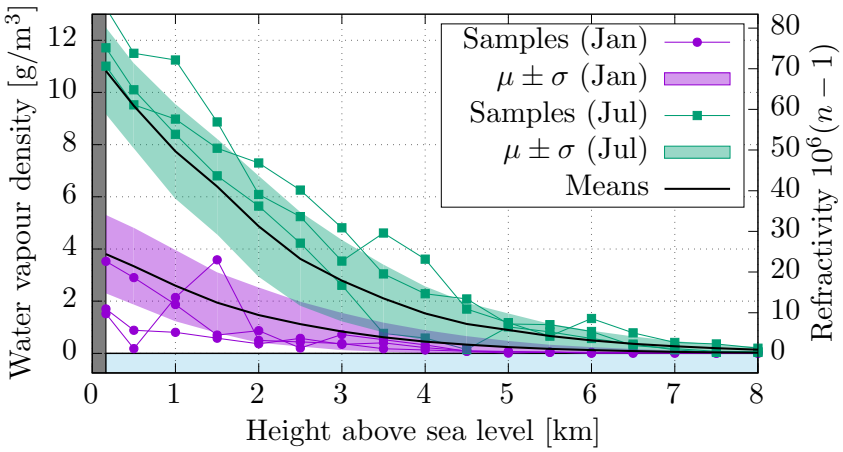

Figure 2.2: Mean water vapour density profiles and its standard deviation computed from radiosonde measurements over Landvetter Airport during six years 2010–2016. Also shown are profiles from individual soundings on three consecutive days. On the right axis an approximate refractivity scale using effective temperature 270 K (see Section 2.4) is provided.

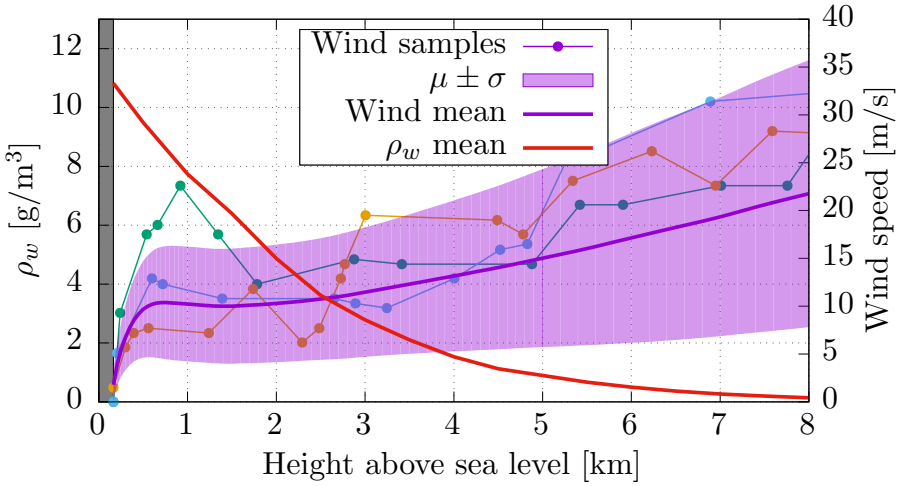

Figure 2.3: The average vertical wind profile and its standard deviation obtained from radiosondes released from Landvetter Airport during 2010–2016. Also shown are three randomly chosen individual soundings from that period. Furthermore, the July mean water vapour profile from the same years is overlaid. This is the same curve as in Figure 2.2.

2.3 The frozen-flow hypothesis

In previous work on statistical modelling of refractive index fluctuations, time dependence has been derived from the *frozen-flow hypothesis*. In this context, it states that the refractive index field evolves so slowly that it can be considered to just be translated by a constant horizontal wind. More precisely, the time evolution of the refractive index field n is assumed to be

$$n(\vec{r}, t) = n(\vec{r} - \vec{v}t, 0), \tag{2.6}$$

where $\vec{v}$ is horizontal, independent of height and constant in time. It is worthwhile to investigate the validity of this approximation.

As a first approximation it is very reasonable. It suffices to look at clouds to see this. Much of the time they seem to move steadily across the sky without perceptively changing shape. See Figure 2.4 for an example of a cloud moving past. Although not visible at this temporal resolution, most of the apparent changes come from the changed viewing angle.

But it is possible to be more quantitative. Figure 2.3 shows the mean wind speed and its standard deviation during 6 years as a function of height over western Sweden. In a large part of the troposphere, the mean wind speed is roughly constant at 10 m/s. This is a small justification for the frozen-flow hypothesis. Furthermore, it immediately provides a typical timescale for the simulation since it means all the air in a 40 km wide simulation volume will be replaced in about 1 h.

Contrary to the frozen-flow hypothesis, the mean wind speed is not constant the first few hundred metres. In this boundary layer, the mean wind speed rapidly increases. Using data from Figure 2.2 reveals that the first 500 m contains about 15 % of the total water vapour. This corresponds to a delay of several centimetres of

delay. In addition, the boundary layer is more turbulent than the free atmosphere, so noticeable deviations from the frozen-flow hypothesis can be expected.

Although the long term mean wind speed is roughly constant in the free atmosphere, wind shear occurs there as well. This can be guessed already from the large standard deviation in Figure 2.3 and is confirmed by the variations within individual soundings. Using the same data that was used to generate Figure 2.3, it is found that the mean wind shear (here defined as $\left\langle \left| \frac{\mathrm{d}\vec{v}}{\mathrm{d}z} \right| \right\rangle$) in the height range 1–8 km is about 6 $\frac{\mathrm{m/s}}{\mathrm{km}}$ with an equally large standard deviation. In comparison, just above the ground the mean wind shear is 70 $\frac{\mathrm{m/s}}{\mathrm{km}}$ and at 500 m it is 15 $\frac{\mathrm{m/s}}{\mathrm{km}}$.

Furthermore, there are some intrinsic variations in the atmosphere. A simulation that translates a field with height dependent wind, needs instrinsic time dependence to have stationary statistics. Another more practical reason for needing intrinsic time dependence will be given in Chapter 3.

Again, clouds are the most easily accessible way to see the phenomenon. See Figure 2.5 for a view of the edge of a cloud during 50 s. Though the cloud moved noticeably in that time, the edge was kept in the centre of the frame. The large scale features clearly remain the same, but the smallest scales on the edge visibly change. To some extent, the changes involve evaporation and condensation of water vapour which does not necessarily affect the density of water vapour very much. But, it is still evidence that small scale change can be observed even during 10 s.

2.4 Refractive index and delay

This description is mainly based on [12] and [13]. The extra time taken for a signal to propagate through the atmosphere due to a refractive index n (>1) is $\Delta t = (n-1)L/c$ with c being the speed of light. This quantity is conveniently expressed as the equivalent, excess optical path length $\Delta L = c\Delta t = (n-1)L$. At radio frequencies, the excess refractivity $\Delta n \coloneqq (n-1)$ can be decomposed into the effect from the dry components of air and water vapour separately

$$\Delta n \approx k_1 \rho_d + k_2 \frac{\rho_w}{T} + k_3 \rho_w, \tag{2.7}$$

where $\rho_{d/w}$ is the density of the "dry air"/water vapour, T is the temperature and k_i are constants.

Since dry air consist of non-polar molecules, their electromagnetic interaction is limited to induced polarization and therefore the refractive index depends only density (the first term of (2.7)). Induced polarization happens in water molecules as well (third term), but the dominant interaction is with the static dipole moment (second term). Due to the competition between field alignment and thermally induced rotations, this effect has a temperature dependence. However, the absolute temperature only varies on the order of 30 % in the troposphere, so using an effective constant temperature in lieu of T in (2.7) is a relatively small error. This is the basis behind the refractivity scale on the right axis in Figure 2.2.

As mentioned in Chapter 1, the hydrostatic delay can be accurately inferred from the ground pressure since

$$\int_0^\infty g\rho(z)\mathrm{d}z = p(0) \tag{2.8}$$

Figure 2.4: Wide-angle photographs of the sky taken at 10 min intervals in Kungsbacka, Sweden on 2018-05-25. Read left to right and top to bottom. The clouds move from the lower right to the upper left

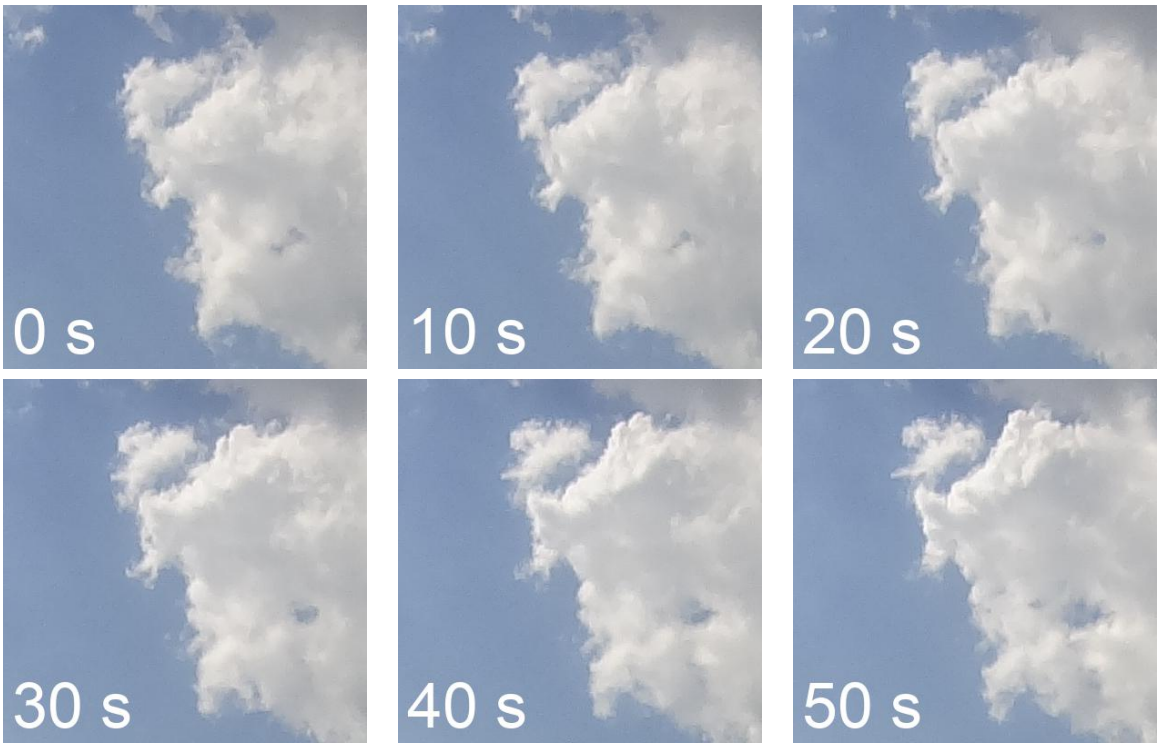

Figure 2.5: Comoving photographs of the edge of a cloud taken at 10 s intervals in Kungsbacka, Sweden on 2018-06-04.

according to (2.1), where $\rho = \rho_d + \rho_w$ is the total density. Hence a zenith hydrostatic delay can be defined as

$$\text{ZHD} := \int k_1 \rho \mathrm{d}z = k_1 p(0)/g. \tag{2.9}$$

Notice that (2.9) includes some of the contributions of water vapour. It is defined this way so that it can be accurately inferred from ground measurements of air pressure. The zenith wet delay can then be defined as the remaining contributions of water vapour

$$\text{ZWD} = \int k_2 \rho_w / T + (k_3 - k_1)\rho_w \mathrm{d}z. \tag{2.10}$$

These are delays in the zenith direction. The standard way of obtaining delays in other directions is through mapping functions [14]. The simplest mapping function is that of a non-curved, horizontally homogeneous atmosphere. Then Δn is only a function of altitude z. In the absence of bending, the slant delay $\Delta L(\theta)$ observed at elevation θ is

$$\Delta L(\theta) = \int \Delta n(s \sin(\theta)) \mathrm{d}s = \frac{1}{\sin(\theta)} \int \Delta n(z) \mathrm{d}z = \frac{1}{\sin(\theta)} \Delta L(90^\circ). \tag{2.11}$$

Therefore, the mapping function between zenith and slant delay is $1/\sin(\theta)$. This mapping function used universally in this thesis to obtain *equivalent zenith delays* from simulated *slant delays*.

2.4.1 Measurement methods

There are many methods that can estimate the wet delay. Some are used here to get some bearings on time and length scales involved. See [15] for further references and long-term comparison of the methods.

By measuring pressure, temperature and humidity as a function of height, the refractivity Δn can be computed from (2.7). Integrating Δn with respect to height then provides an indirect estimate of the zenith delay. For meteorological purposes, such measurements are performed daily by radiosondes carried by weather balloons. This is one of few methods that can be used to obtain the vertical refractivity profile. On the other hand, the horizontal and temporal resolution is limited.

Figure 2.6 shows the wet and hydrostatic zenith delays at Landvetter Airport during three years obtained by this method. The maximum variations in both types of delay are of the same magnitude. There is no clear trend in the hydrostatic delay, but there is a clear seasonal trend in wet delay due to varying temperatures. Large variations from one day to the next are also seen.

The presence of water vapour can be directly measured using water vapour radiometry. The principle is to measure the intensity of thermal radiation at two frequencies close to 22.2 GHz; where water vapour has a weak emission/absorption peak. The measured quantity is a temperature weighted integral of the water vapour density within the antenna beam. Unfortunately, this is a different weighting than that determining the delay (2.10), so the vertical profiles must be modelled. These models are typically obtained from nearby radiosonde launches. Water vapour radiometers can typically be sampled once per 1–10 s and measure in narrow (~5°) beams allowing local structure to be mapped.

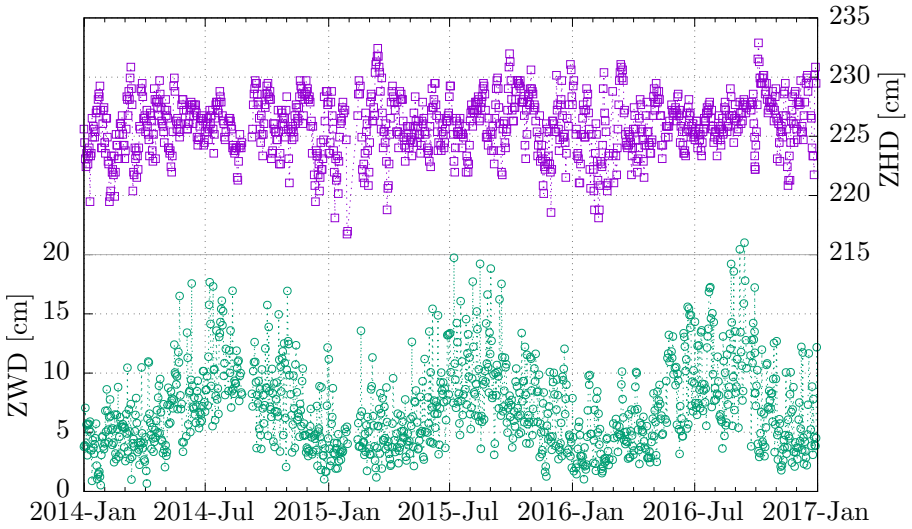

Figure 2.6: The hydrostatic (ZHD) and wet (ZWD) zenith delays during three years at Landvetter Airport 57° N computed from radiosonde data obtained from the IGRA[16] (station SWM00002527) using equations (5.14) and (5.18) from [12].

The total tropospheric delay can be directly estimated using GNSS receivers. By subtracting the hydrostatic delay (obtained from measurements of ground pressure (2.9)) from the total delay, an estimate for the wet delay is obtained. This also allows mapping of local structure, but the directions are fixed by satellite orbits.

A much less direct, but also simpler, approach is to measure ρ_w at surface (via temperature and humidity). It can be expected to be proportional to the total amount of water vapour in the local atmosphere on sufficiently long timescales. The relation between delay and ρ_w can be expected to have an additional seasonal dependence due to varying vertical temperature profiles.

Lastly, meteorological data from many sources are combined with fluid dynamics in numerical weather models for forecasts and interpolation. Ray tracing through the output of such a model can be used to compute delays. For instance that was done in [10].

2.4.2 Measured characteristics

In Figure 2.7, the zenith wet delay over Onsala Space Observatory as estimated from the aforementioned methods during July 2016 are shown. Additionally, a smoothed ground based model is included. The estimated delay from that method is just $\rho_w \times 1.39\,\mathrm{cm}/(\mathrm{g/m^3})$, where ρ_w is measured at a ground based weather station. The numerical coefficient was obtained from linear regression between daily means of ZWD from GPS measurements and ρ_w. It predicts the wet delay to near centimetre accuracy on some days, but deviates wildly on others. For instance on day 6 and 23 in Figure 2.7a. Interestingly, GPS and radiometer estimates often disagree on the centimetre level.

A more quantitative way to show the fluctuations in zenith delay is through prob-

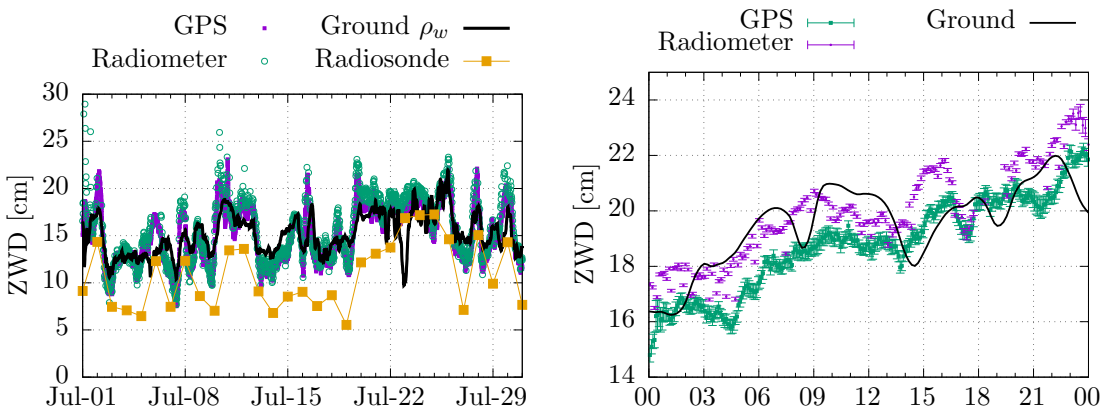

(a) 2016-07 **(b)** 2016-07-25

Figure 2.7: Zenith wet delay during July 2016 estimated from radiometry, GPS and a ground based model over Onsala Space Observatory. Additionally, delays obtained from radiosonde launches at the nearby Landvetter Airport are included. In (a) the full month is presented while in (b) one full day is presented, the horizontal unit being hours.

ability density and structure function of delay increments

$$\Delta\mathrm{ZWD}(\Delta t) \coloneqq \mathrm{ZWD}(t + \Delta t) - \mathrm{ZWD}(t). \tag{2.12}$$

For data with a wide range of time scales, it is preferable to work with ΔZWD instead of ZWD directly, since this naturally filters out longer time-scales. An alternative is spectral methods, but for that to be well defined, several of the longest periods must be contained in the data. By observing many t-values, an empirical distribution of $\Delta\mathrm{ZWD}(\Delta t)$ is obtained. We can also define an empirical structure function

$$S(\Delta t) \coloneqq \left\langle \Delta\mathrm{ZWD}(\Delta t)^2 \right\rangle, \tag{2.13}$$

where the expectation is the mean over all available t-values.

Both empirical probability distributions and structure functions measured by GPS are presented in Figure 2.8 at Onsala Space Observatory (57° N) as well as the Brazilian city Recife at latitude 8° S during June–September 2016. Since the RECF site is located near the equator, there is a much less pronounced seasonal dependence. The delays are significant and of similar magnitude throughout the year.

The probability density is Gaussian only for the smallest Δt. It should be noted that in this regime, ΔZWD might be dominated by noise, since the uncertainty estimate is 0.25 cm (for both datasets). For larger time differences, the tails of the distribution grow increasingly fat; appearing to fall off linearly instead of quadratically (in logarithmic scale). This is characteristic of PDFs of turbulent quantities. Interestingly, the deviation from normality is smaller for measurements from the tropical RECF station than the continental[1] ONSA station.

The structure functions display the same behaviour for $\Delta t < 1$ h, but with about 1.5 times larger fluctuations at RECF. For longer periods the fluctuations are larger

[1] According to the Köppen climate classification.

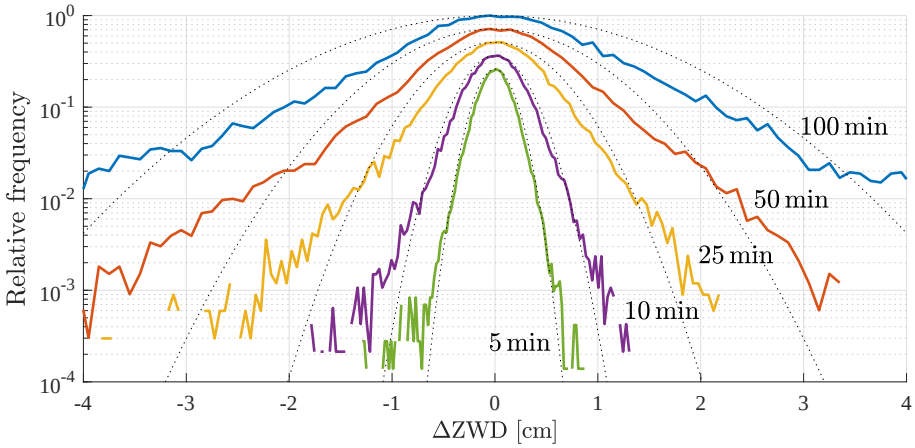

(a) Observed pdf of $\Delta\mathrm{ZWD}(\Delta t)$ at ONSA

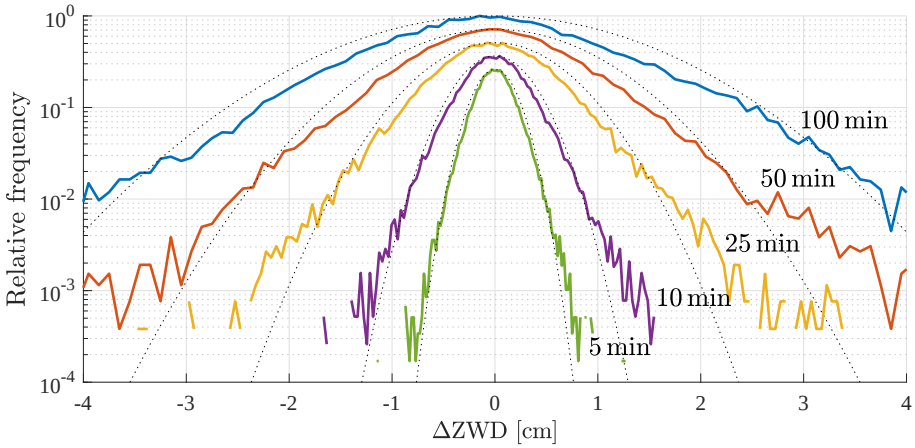

(b) Observed pdf of $\Delta\mathrm{ZWD}(\Delta t)$ at RECF

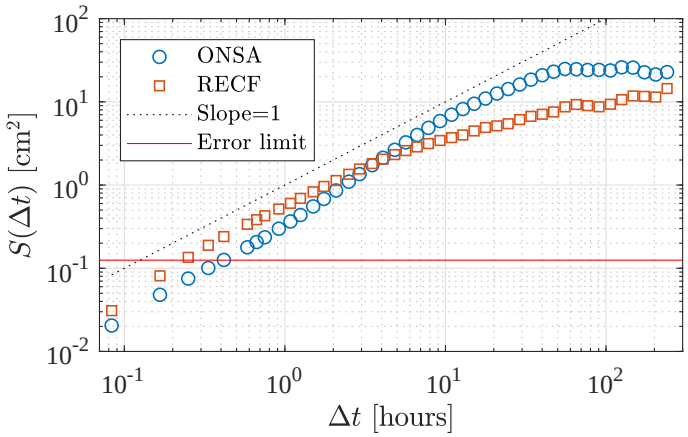

(c) Empirical structure functions.

Figure 2.8: Statistics of increments of ZWD estimated from GPS during June–September 2016 at Onsala, Sweden and Recife, Brazil. In (a) & (b) the unnormalized probability density function of $\Delta\mathrm{ZWD}(\Delta t)$ are plotted for various Δt indicated by text next to the curves on a logarithmic scale. The dotted lines are Gaussians with the same variance as the observations. In (c) the second order structure function is plotted as a function of Δt on a log-log scale. The dotted line has log-log slope 1. Here $\mathrm{ZWD}(t)$ is sampled once per 5 min.

at ONSA. This can be interpreted as RECF having more consistent turbulence while ONSA has larger day to day variations. When Δt approaches one week, the ONSA structure function has saturated, but the RECF-curve still has an upward trend. At and below the 1 h timescale the structure function follows an approximate power law; nothing spectacular happens. So 1 h is not a special physical timescale, it is only an important timescale in simulations. Just like 40 km is not a physically significant length scale.

Chapter 3

Tropospheric turbulence

Turbulence is ubiquitous in macroscopic fluid dynamics, and the most macroscopic fluid on the Earth is the atmosphere. The Reynolds number $\mathrm{Re} = uL/\nu$ (where u, L and ν are respectively a characteristic velocity, length and kinematic viscosity) is a dimensionless quantity that can be used to predict if a flow is turbulent. The threshold Reynolds number above which turbulence occurs is usually around 10^3–10^5 depending on geometry. In the lower troposphere $u \approx 10\,\mathrm{m/s}$ (see Figure 2.3) and $\nu \approx 10^{-5}\,\mathrm{m^2/s}$ at standard conditions [17, T-1.5]. There are many length scales involved, but even with L as small as 1 m, the Reynolds number is

$$\mathrm{Re} = \frac{uL}{\nu} \sim 10^6. \tag{3.1}$$

This is well into the turbulent regime, so turbulence must have a central role in any microscale description of the atmosphere. Particularly when dealing with rapid fluctuations.

There is no precise definition of what a turbulent flow is; only a collection of symptoms that together makes a flow fall under that label [18]. The central idea behind turbulence is a cascade process wherein energy is extracted at large scales forming coherent, but unstable, flow patterns which are continuously broken up into smaller flow features. This process is sketched in Figure 3.1. The breakup process continues until the features are small enough for viscous forces to dominate over inertia. However, a fractal-like cascade process is not enough for a flow to be called turbulent. It must also be irregular in both space and time. This irregularity is the basis of statistical descriptions of turbulence.

3.1 Kolmogorov turbulence

The most widely used theory of turbulence is a scaling analysis originally brought to light by Kolmogorov in 1941[1]. The version presented here is different and less rigorous than the original, but also more accessible.

[1] The original articles are in Russian. References to the original papers as well as some expansion of the old ideas are provided in [19, 20].

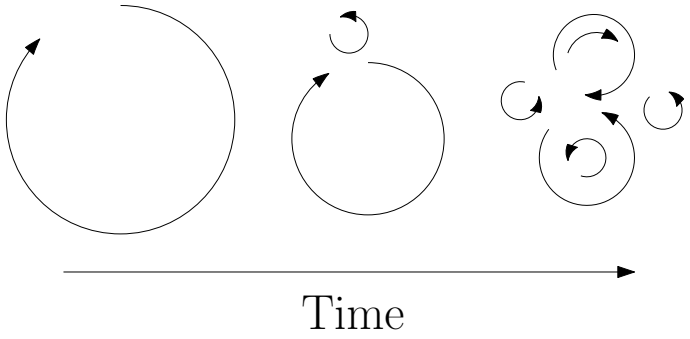

Figure 3.1: Sketch of how large scale coherent structures recursively break up into smaller structures due to instability; eventually forming intricate fractal-like patterns; one of the characteristics of turbulence.

3.1.1 Microscales

The assumptions are that an incompressible fluid is put into statistically homogeneous and isotropic turbulent motion through some unspecified method. This enables dimensional analysis to be employed[2]. Of course, for the atmosphere neither of these assumptions is valid due to the strong vertical dependence. For instance, the pressure at 5 km is only half that at sea level. Therefore, at best, the theory can only be applied locally. The behaviour at large scales has to be obtained by some phenomenological extension.

The first question to ask is at what length scale the fluid becomes smooth (dissipation dominated). These small scales are far removed from the forces driving the turbulence, so the only available quantities is the fluids kinematic viscosity ν [m^2/s] and the amount of power dissipated per unit *mass* ε [$W/kg = m^2/s^3$] (since the density is constant, this it is directly proportional energy dissipation per *volume*). The only length and time scales that can be formed from these quantities are

$$\ell = \nu^{3/4}\varepsilon^{-1/4} \text{ and } \tau = \nu^{1/2}\varepsilon^{-1/2}. \tag{3.2}$$

To estimate the scales (3.2), the dissipation ε is needed. It can be estimated from a characteristic velocity U and length L as $\varepsilon \sim U^3/L$. As shown in Figure 2.3, velocity gradients as large as $5\ \frac{m/s}{km}$ can sometimes be sustained over vertical distances as large as 1 km. Hence a reasonable range of velocity scales U is 1–5 m/s on a length scale of $L = 1$ km. This range inserted into (3.2) yields $\ell = 0.3$–1 mm and $\tau = 0.01$–0.1 s. Clearly the ℓ-scale cannot be resolved in a simulation stretching 5 km into the air; which is the height needed to contain most of the water vapour (see Figure 2.2).

The significance of the characteristic time τ is that only for periods shorter than that is the frozen flow hypothesis is valid on all length scales. However, on periods of order τ, the field can only change significantly on length scales of order ℓ. Longer length scales should change more slowly. At this point it is unclear how to quantify that relation. This will be considered in some detail in Section 3.1.3. Still, notice that τ is somewhat closer to observational timescales than ℓ.

[2]I will here denote dimensions with the symbols of SI units that have such a dimension. Firstly to avoid confusion with actual lengths and times. Secondly because there are many short forms for relevant combinations of base dimensions such as W, Pa, N etc.

3.1.2 Spatial structure of the inertial subrange

At larger scales (but still small enough to be in the isotropic regime) we can ask how much energy resides at each scale. Suppose there is such an energy distribution $E(k)\mathrm{d}k$ (where k [m^{-1}] denotes wavenumber) with the property that the total energy contained in turbulence per unit mass is $\int_0^\infty E(k)\mathrm{d}k$ [$\mathrm{m}^2/\mathrm{s}^2$]. Since the dimension of k is m^{-1}, the dimension of E is $\mathrm{m}^3/\mathrm{s}^2$. At scales larger than (3.2), viscous forces have a negligible impact on the dynamics of the breakup process. Therefore ν should not influence $E(k)$ for $k \gg \ell^{-1}$. Thus the only available quantities to form E are ε and k, which implies

$$E(k) \sim \varepsilon^{2/3} k^{-5/3}. \tag{3.3}$$

This is the famous Kolmogorov 5/3-law.

One way to to relate $E(k)$ to flow quantities is through the spatial power spectral density $\Phi(\vec{k})$ of the instantaneous velocity field $\vec{u}(\vec{r})$. It can be defined as

$$\Phi(\vec{k}) \coloneqq (2\pi)^{-3} \frac{1}{V} \left| \int_V \vec{u}(\vec{r}) \mathrm{e}^{-\mathrm{i}\vec{k}\cdot\vec{r}} \mathrm{d}^3\vec{r} \right|^2, \tag{3.4}$$

in the limit of large volume V. A version of Plancherel's theorem then states

$$\frac{1}{V} \int_V \frac{1}{2} \left|\vec{u}(\vec{r})\right|^2 \mathrm{d}^3\vec{r} = \frac{1}{2} \int \Phi(\vec{k}) \mathrm{d}^3\vec{k}, \tag{3.5}$$

where the left hand side can be recognized as the average kinetic energy per mass (to see this, divide and multiply with the density). Due to isotropy, Φ is a function only of $k = \left|\vec{k}\right|$ and hence we can identify

$$E(k)\mathrm{d}k \sim \Phi(k) k^2 \mathrm{d}k. \tag{3.6}$$

This results in the scaling

$$\Phi(k) \sim \varepsilon^{2/3} k^{-11/3}. \tag{3.7}$$

There is an equivalent description in real space in terms of the structure function

$$D(\vec{r}) \coloneqq \left\langle \left|\vec{u}(\vec{x}+\vec{r}) - \vec{u}(\vec{x})\right|^2 \right\rangle \sim \varepsilon^{2/3} \left|\vec{r}\right|^{2/3}, \tag{3.8}$$

the average $\langle\cdot\rangle$ being taken over ensembles, space ($\vec{x}$) and/or time. This can be seen directly from a dimensional analysis. An alternative viewpoint is that $\left|\vec{r}\right|^{2/3}$ is, apart from some constants, the distributional Fourier transform of $\left|\vec{k}\right|^{-11/3}$.

3.1.3 Temporal structure of the inertial subrange

As described in Section 2.3, previous statistical treatments have used the frozen flow hypothesis to describe time-evolution according to (2.6). But as demonstrated by Figure 2.5, visible changes in cloud shape occur in 10 s on short length scales. Thankfully, for Kolmogorov turbulence, it is straightforward to derive a temporal structure function/spectrum using the same method as in the last section to account for such intrinsic change.

As before, the most elementary quantity to work with is the structure function

$$D(t) \coloneqq \left\langle \left|\vec{u}(T+t) - \vec{u}(T)\right|^2 \right\rangle. \tag{3.9}$$

With the same reasoning as above, $D(t)$ can only depend on ε and t. The only possibility is

$$D(t) \sim \varepsilon t. \tag{3.10}$$

Interestingly this scaling is exactly the same as for Brownian motion.

The same result can be expressed in terms of the temporal power spectral density

$$\Phi(\omega) \coloneqq (2\pi)^{-1} \frac{1}{T} \left| \int_0^T \vec{u}(t) \mathrm{e}^{-\mathrm{i}\omega t} \mathrm{d}t \right|^2 \tag{3.11}$$

which satisfies $\int_{-\infty}^{\infty} \Phi(\omega) \mathrm{d}\omega = \frac{1}{T} \int_0^T \left|\vec{u}(t)\right|^2 \mathrm{d}t$. The only combination of ε and ω producing the correct dimension is

$$\Phi(\omega) \sim \varepsilon \omega^{-2}. \tag{3.12}$$

A combined spatio-temporal structure function $D(\vec{r}, t) = \left\langle \left|\vec{u}(\vec{r}, t) - u(0,0)\right|^2 \right\rangle$ and corresponding spectrum $\Phi(\vec{k}, \omega)$ (which has units $(\mathrm{m/s})^2 \cdot \mathrm{m}^3 \cdot \mathrm{s} = \mathrm{m}^5\mathrm{s}^{-1}$) can be defined, but cannot be determined from dimensional analysis alone. This is because a dimensionless quantity

$$\alpha(k, \omega) = \varepsilon k^2 \omega^{-3} \tag{3.13}$$

can be formed. Therefore Φ can have an arbitrary dependence on α. There are many ways to parametrise the spectrum. One which will prove natural in Section 3.4.2 is

$$\Phi(\vec{k}, \omega) = \varepsilon^{2/3} k^{-11/3} \omega^{-1} \varphi(\alpha), \tag{3.14}$$

where φ is an arbitrary function. It is straightforward to verify that the marginal spectra $\int \Phi(\vec{k}, \omega) \mathrm{d}\omega$ and $\int \Phi(\vec{k}, \omega) \mathrm{d}^3\vec{k}$ reduce to the scalings (3.7) and (3.10) respectively. Simply use (3.14) and change integration variable to α.

Equivalently, (3.13) can be thought of as defining a dispersion relation relating spatial scales to temporal scales. In frequency space $k \propto \varepsilon^{1/2} \omega^{3/2}$ or in real space $L \propto \varepsilon^{-1/2} T^{3/2}$. L can be interpreted at the characteristic length scale of intrinssic changes during periods of order T.

3.1.4 Convection by turbulence

The results of the previous sections were strictly for the *flow velocity* of incompressible, homogeneous, isotropic turbulence. But in this thesis it is the scalar field of water vapour density that is of interest. In the literature, the same spatial structure $\propto r^{2/3}$ is universally used, but I have not been able to access a reference for that.

3.2 Existing applications to signal delay

In most existing work, it is assumed that the refractivity field can be described using Kolmogorov theory for the convection of a passive scalar by isotropic turbulence. Sometimes this is used as a starting point for phenomenological structure functions and spectra that don't diverge on long scales (as do (3.8) and (3.7)). See for instance Ishimaru (1972) [21] for a spectral description and Treuhaft & Lanyi (1987) [22] for a "structural" description. The temporal structure is without exception obtained from the frozen flow hypothesis. This is used to obtain a temporal spectrum, covariance or structure function of the signal of interest by integrating over the 4D statistics.

Using covariance matrices computed in this way to generate random signals has been done by at least two authors [8, 9]. However in most existing literature, the predicted signal statistics are compared directly to measured statistics. Without generating any stochastic signals in between.

On short time and length scales (orthogonal to the signal path) excellent agreement between model and specific experiments has been found for some observables. At microwave frequencies, see [21] which analyzes data from an experiment [23] where microwave signals travelled along a 65 km slanted path between two Hawaiian islands. In the optical regime, an experiment in similar spirit [24] was carried out with a laser beam propagating over flat land or a lake for 0.25–16 km. Agreement was found here as well, at least for second order structure functions.

Comparing [23] with typical GNSS and radiometer data, their lowest frequency data point has a period of 100 s which is not much slower than the sampling rate achievable with GNSS and radiometers. Furthermore, in these experiments the length scales probed orthogonal to the path were very limited. Therefore the success of these experiments say little about the applicability of the turbulence models to refractivity fields of size 10 km.

At longer time and length scales the most relevant articles might be [25, 26] that apply the theory to GPS signals and [27] that applies it to water vapour radiometer observations. In both cases, qualitative agreement between model and experiment is found. Quantitatively it is much less certain. In part due to limited statistics, which is a consequence of measuring on long timescales. Furthermore, for measurements lasting days, the model parameters cannot be expected to remain constant.

3.3 Gaussian random fields

In Section 3.1, scaling relations for structure functions and spectra were derived in the simplest existing turbulence model. The next task is to find some way to compute random fields having a given structure. One natural way is through *Gaussian random fields*. The main reason is practical; Gaussian fields have several properties that make them analytically and computationally tractable. Furthermore, errors and disturbances are very often modelled that way. However, as will be discussed in Chapter 5, there are many reasons to distrust this assumption. Especially in light of Figure 2.8 which shows directly that observed wet delays are non-Gaussian.

A d-dimensional Gaussian random field f is a random function $f : \quad D \to \mathbb{R}$

with $D \subset \mathbb{R}^d$ such that for any $\vec{r}_1, \cdots, \vec{r}_n \in D$, the quantities $f(\vec{r}_1), \cdots, f(\vec{r}_n)$ are random variables with a multivariate normal distribution. This imposes strong restrictions on the statistics of the field. The most important in this context is that it is completely characterised by its mean $\mu(\vec{r})$ and covariance $C(\vec{r}_1, \vec{r}_2)$ functions

$$\mu(\vec{r}) \coloneqq \langle f(\vec{r}) \rangle, \quad C(\vec{r}_1, \vec{r}_2) \coloneqq \langle (f(\vec{r}_1) - \mu(\vec{r}_1))(f(\vec{r}_2) - \mu(\vec{r}_2)) \rangle \tag{3.15}$$

since this determines the probability distribution of any finite sample. In the sequel, all fields will be assumed to have $\mu \equiv 0$ since this simplifies the equations. If some field of interest has non-zero mean, it can be transformed away with $f(\vec{r}) \to f(\vec{r}) - \mu(\vec{r})$, so there is no loss of generality.

The properties of interest follow directly from two properties of multivariate normal distributions: they are closed under *linear combinations* and *limits*. To be precise, let $\vec{\xi} = (\xi_1, \cdots \xi_n)$ be a Gaussian random vector meaning that it follows a multivariate normal distribution; for instance it could be a sampled Gaussian random field. Then any vector $(\eta_1, \cdots \eta_m)$ formed from linear combinations of the ξ:s

$$\eta_i = \sum_{j=1}^{n} A_{ij} \xi_j, \tag{3.16}$$

also follows a multivariate normal distribution. Secondly, any sequence $\vec{\xi}_1, \vec{\xi}_2, \cdots$ of Gaussian random vectors whose mean vectors $\vec{\mu}_i \coloneqq \left\langle \vec{\xi}_i \right\rangle$ and covariance matrices $C_i \coloneqq \left\langle \vec{\xi}_i \vec{\xi}_i^{\mathrm{T}} \right\rangle$ converge (in $\mathbb{R}^n$ and $\mathbb{R}^{n^2}$ respectively), converges in distribution to a multivariate normal distribution with mean $\mu = \lim_{i\to\infty} \mu_i$ and covariance $C = \lim_{i\to\infty} C_i$. These properties ensure that any sums and integrals (for instance of the Fourier type) of Gaussian fields remain Gaussian fields/vectors.

Not all functions can be covariance functions. From the definition (3.15), it follows directly that a covariance function must be symmetric $C(\vec{r}_1, \vec{r}_2) = C(\vec{r}_2, \vec{r}_1)$. Furthermore, for any finite set of points $\vec{r}_1, \cdots, \vec{r}_n$, the covariance matrix $C_{ij} = \langle f(\vec{r}_i) f(\vec{r}_j) \rangle = C(\vec{r}_i, \vec{r}_j)$ must be positive semi-definite. In general it is not easy to characterize these functions, but it is possible in the special case of *homogeneous* fields.

A homogeneous Gaussian field is invariant under translations. More precisely its covariance function has the property

$$C(\vec{r}_1, \vec{r}_2) = C(\vec{r}_1 - \vec{r}_2). \tag{3.17}$$

Define the Fourier transform of the d-dimensional Gaussian field f as

$$\hat{f}(\vec{k}) \coloneqq \int f(\vec{r}) \mathrm{e}^{-\mathrm{i}\vec{k}\cdot\vec{r}} \mathrm{d}^d\vec{r}. \tag{3.18}$$

Using the properties mentioned above, it can be proved that $\hat{f}$ is a complex-valued Gaussian field[3] . Since $\hat{f}$ is complex, both the covariance $\left\langle \hat{f}(\vec{k}_1) \hat{f}^*(\vec{k}_2) \right\rangle$ and pseudo-covariance $\left\langle \hat{f}(\vec{k}_1) \hat{f}(\vec{k}_2) \right\rangle$ are needed to completely specify it. But due to $f(\vec{r})$ being

[3]Note however that (3.18) doesn't converge in the normal sense[4] for homogeneous fields. Still, it is meaningful in a distributional sense.

[4]With probability 1, $f(\vec{r})$ will be such that the integral is not be absolutely convergent.

real, its Fourier transform satisfies $\hat{f}(-\vec{k}) = \hat{f}^*(\vec{k})$ so both functions can be computed at the same time:

$$\begin{aligned}
\left\langle \hat{f}(\vec{k}_1)\hat{f}(\vec{k}_2)\right\rangle &= \int\int \langle f(\vec{r}_1)f(\vec{r}_2)\rangle \, \mathrm{e}^{-\mathrm{i}\vec{r}_1\cdot\vec{k}_1-\mathrm{i}\vec{r}_2\cdot\vec{k}_2}\mathrm{d}^d\vec{r}_1\mathrm{d}^d\vec{r}_2 \\
\{(3.15) \text{ and } (3.17)\} &= \int\int C(\vec{r}_1-\vec{r}_2)\mathrm{e}^{-\mathrm{i}\vec{r}_1\cdot\vec{k}_1-\mathrm{i}\vec{r}_2\cdot\vec{k}_2}\mathrm{d}^d\vec{r}_1\mathrm{d}^d\vec{r}_2 \\
\left\{\vec{r}=\vec{r}_1-\vec{r}_2,\ \vec{R}=(\vec{r}_1+\vec{r}_2)/2\right\} &= \int\int C(\vec{r})\mathrm{e}^{-\mathrm{i}\vec{R}\cdot(\vec{k}_1+\vec{k}_2)}\mathrm{e}^{-\mathrm{i}\vec{r}\cdot\frac{1}{2}(\vec{k}_1-\vec{k}_2)}\mathrm{d}^d\vec{r}\mathrm{d}^d\vec{R} \\
\left\{\int \mathrm{e}^{-\mathrm{i}\vec{x}\cdot\vec{k}}\mathrm{d}^d\vec{x} = (2\pi)^d\delta(\vec{k})\right\} &= (2\pi)^d\,\delta(\vec{k}_1+\vec{k}_2)\int C(\vec{r})\mathrm{e}^{-\mathrm{i}\vec{r}\cdot\vec{k}_1}\mathrm{d}^d\vec{r}. \qquad (3.19)
\end{aligned}$$

Hence, to the extent allowed by the identity $\hat{f}(-\vec{k}) = \hat{f}^*(\vec{k})$, the Fourier-transformed field consists of circular[5] Gaussians independent for each $\vec{k}$. This is the motivation for spectral simulation methods.

Due to the δ-function in (3.19), the frequency domain variance is infinite in some sense. That makes it more convenient to work with the spectrum $\Phi(\vec{k})$ defined as

$$\Phi(\vec{k}) = \frac{1}{(2\pi)^d}\int C(\vec{r})\mathrm{e}^{-\mathrm{i}\vec{r}\cdot\vec{k}}\mathrm{d}^d\vec{r} \iff C(\vec{r}) = \int \Phi(\vec{k})\mathrm{e}^{\mathrm{i}\vec{r}\cdot\vec{k}}\mathrm{d}^d\vec{k}. \qquad (3.20)$$

The symmetry $C(-\vec{r}) = C(\vec{r})$ implies that Φ is real. Furthermore, C is positive semi-definite so for any real function f

$$\begin{aligned}
0 &\le \int C(\vec{r}_1-\vec{r}_2)f(\vec{r}_1)f(\vec{r}_2)\mathrm{d}^d\vec{r}_1\mathrm{d}^d\vec{r}_2 = \int C(\vec{r})f(\vec{R}+\vec{r})f(\vec{R})\mathrm{d}^d\vec{R}\mathrm{d}\vec{r} \\
&= \int C(\vec{r})F(\vec{r})\mathrm{d}^d\vec{r}, \qquad (3.21)
\end{aligned}$$

where $F(\vec{r}) \coloneqq \int f(\vec{R})f(\vec{R}+\vec{r})\mathrm{d}^d\vec{R}$ is the autocorrelation of f. Writing F in terms of the Fourier transform $\hat{f} = \mathcal{F}[f]$ we get

$$\begin{aligned}
F(\vec{r}) &= \int \hat{f}(\vec{k}_1)\hat{f}(\vec{k}_2)\mathrm{e}^{\mathrm{i}\vec{R}\cdot\vec{k}_1+\mathrm{i}(\vec{R}+\vec{r})\cdot\vec{k}_2}\mathrm{d}^d\vec{R}\mathrm{d}^d\vec{k}_1\mathrm{d}^d\vec{k}_2 \\
&= (2\pi)^d\int \delta(\vec{k}_1+\vec{k}_2)\hat{f}(\vec{k}_1)\hat{f}(\vec{k}_2)\mathrm{e}^{\mathrm{i}\vec{r}\cdot\vec{k}_2}\mathrm{d}^d\vec{k}_1\mathrm{d}^d\vec{k}_2 \\
&= (2\pi)^d\int \hat{f}(\vec{k})\hat{f}(-\vec{k})\mathrm{e}^{-\mathrm{i}\vec{r}\cdot\vec{k}}\mathrm{d}^d\vec{k}.
\end{aligned}$$

Since f is real we have that $\hat{f}(\vec{k})\hat{f}(-\vec{k}) = \left|\hat{f}(\vec{k})\right|^2 \ge 0$. Plugging this into (3.21) we immediately obtain

$$0 \le (2\pi)^d\int C(\vec{r})\mathrm{e}^{-\mathrm{i}\vec{r}\cdot\vec{k}}\left|\hat{f}(\vec{k})\right|^2\mathrm{d}^d\vec{r}\mathrm{d}^d\vec{k} = (2\pi)^{2d}\int \Phi(\vec{k})\left|\hat{f}(\vec{k})\right|^2\mathrm{d}^d\vec{k}. \qquad (3.22)$$

Since f can be chosen arbitrarily, $\left|\hat{f}\right|^2$ can be arbitrarily well localized[6]. Therefore the statement that C is positive semi-definite is equivalent to the spectrum being non-negative $\Phi(\vec{k}) \ge 0$. All under some unspecified regularity conditions.

[5] A circular Gaussian random variable has independent and identically distributed Gaussian real and imaginary parts.

[6] Apart from the $\pm\vec{k}$ symmetry, but that is possessed by Φ as well.

3.4 Simulation methods for Gaussian fields

A very general and direct method to generate a sample $f(\vec{r}_1), \dots f(\vec{r}_n)$ of a zero mean Gaussian field is the following:

1. Assemble the discrete covariance matrix $C_{ij} = C(\vec{r}_i, \vec{r}_j)$. $\mathcal{O}(n^2)$
2. Compute a Cholesky decomposition $C = RR^T$. $\mathcal{O}(n^3)$
3. Generate a vector $\xi = (\xi_1, \dots \xi_n)$ of independent unit variance normal variables. $\mathcal{O}(n)$
4. Compute $X = R\xi$. $\mathcal{O}(n^2)$

The vector X now has the same distribution as $(f(\vec{r}_1), \dots f(\vec{r}_n))$ since it's a zero mean multivariate normal distribution with covariance matrix

$$\left\langle XX^T \right\rangle = R \left\langle \xi\xi^T \right\rangle R^T = C. \tag{3.23}$$

The last equality uses that $\left\langle \xi\xi^T \right\rangle$ is the identity matrix.

Due to the $\mathcal{O}(n^3)$ cost, this is infeasible for generating samples on regular grids in 2D and 3D. If the number of gridpoints in one direction of a d-dimensional grid is N, then the cost of step 2 is $\mathcal{O}(N^{3d})$.

3.4.1 Direct FFT method

This is the method used to generate refractivity fields in this thesis. As shown above, the Fourier transform of a homogeneous Gaussian field is independent noise with a particular variance. The same turns out to be true for homogeneous periodic fields in a discrete setting, allowing the efficiency of the Fast Fourier Transform (FFT) to be leveraged.

To show this, let $\hat{f}_{\vec{k}}$ be a regular d-dimensional array of shape $N_1 \times \cdots \times N_d$ containing independent circular Gaussians with variance $2\Phi_{\vec{k}}$ (the imaginary and real parts are independent and each has variance $\Phi_{\vec{k}}$). Then set[7]

$$f_{\vec{r}} = \mathrm{Re}\left(\sum_{\vec{k}} \hat{f}_{\vec{k}} \mathrm{e}^{2\pi \mathrm{i} \tilde{k} \cdot \vec{r}} \right) = \frac{1}{2} \sum_{\vec{k}} \hat{f}_{\vec{k}} \mathrm{e}^{2\pi \mathrm{i} \tilde{k} \cdot \vec{r}} + \text{c.c.} \tag{3.24}$$

Here $\vec{k} = (k_1, \dots k_d)$ is an integer vector, $\tilde{k} = (k_1/N_1, \dots k_d/N_d)$ and *c.c.* is shorthand for "complex conjugate of the preceding expression". The sum goes over all unique values of $\vec{k}$ with $0 \leq k_i < N_i$ and to make the sum interpretable as a DFT (Discrete Fourier Transform), $\vec{r}$ is any integer vector from the same set as $\vec{k}$. The covariance of $f_{\vec{r}}$ is

$$\begin{aligned} \langle f_{\vec{r}_1} f_{\vec{r}_2} \rangle &= \frac{1}{4} \left\langle \left(\sum_{\vec{k}} \hat{f}_{\vec{k}} \mathrm{e}^{2\pi \mathrm{i} \tilde{k} \cdot \vec{r}_1} + \text{c.c.} \right) \left(\sum_{\vec{k}} \hat{f}_{\vec{k}} \mathrm{e}^{2\pi \mathrm{i} \tilde{k} \cdot \vec{r}_2} + \text{c.c.} \right) \right\rangle \\ &= \frac{1}{4} \sum_{\vec{k}_1, \vec{k}_2} \left(\left\langle \hat{f}_{\vec{k}_1} \hat{f}_{\vec{k}_2} \right\rangle \mathrm{e}^{2\pi \mathrm{i} \left(\tilde{k}_1 \cdot \vec{r}_1 + \tilde{k}_2 \cdot \vec{r}_2 \right)} + \left\langle \hat{f}_{\vec{k}_1} \hat{f}^*_{\vec{k}_2} \right\rangle \mathrm{e}^{2\pi \mathrm{i} \left(\tilde{k}_1 \cdot \vec{r}_1 - \tilde{k}_2 \cdot \vec{r}_2 \right)} \right) + \text{c.c.}, \end{aligned} \tag{3.25}$$

[7] It turns out that the imaginary part of (3.24) (before the "Re") has the same distribution and is independent of the real part, but that is not shown here.

where the identity

$$(z_1 + z_1^*)(z_2 + z_2^*) = z_1 z_2 + z_1 z_2^* + z_1^* z_2 + z_1^* z_2^* = (z_1 z_2 + z_1 z_2^*) + \text{c.c.} \tag{3.26}$$

was used. The first term of (3.25) is zero since the $\hat{f}_{\vec{k}}$ were taken to be circular, while the second is $\left\langle f_{\vec{k}_1} f_{\vec{k}_2}^* \right\rangle = 2\delta_{\vec{k}_1 \vec{k}_2} \Phi_{\vec{k}_1}$. Thus

$$\langle f_{\vec{r}_1} f_{\vec{r}_2} \rangle = \frac{1}{4} \sum_{\vec{k}} 2\Phi_{\vec{k}} \mathrm{e}^{-2\pi \mathrm{i} \vec{k} \cdot (\vec{r}_1 - \vec{r}_2)} + \text{c.c.} = \sum_{\vec{k}} \Phi_{\vec{k}} \mathrm{e}^{-2\pi \mathrm{i} \vec{k} \cdot (\vec{r}_1 - \vec{r}_2)}, \tag{3.27}$$

where the last step makes use of the symmetry $\Phi_{\vec{k}} = \Phi_{-\vec{k}}$ (this also ensures that the expression is real). This shows that the array $f_{\vec{r}}$ forms a multivariate normal distribution which is homogeneous (in a periodic sense) since $\langle f_{\vec{r}_1} f_{\vec{r}_2} \rangle = C_{\vec{r}_1 - \vec{r}_2}$ only depends on the difference of the indices (interpreted periodically).

For $\vec{r}$ on an integer grid, (3.27) shows that $C_{\vec{r}}$ is the DFT of $\Phi_{\vec{k}}$. To produce a field with given covariance $C_{\vec{r}}$, the spectrum should be chosen as

$$\Phi_{\vec{k}} = \frac{1}{N_1 \cdots N_d} \sum_{\vec{r}} \mathrm{e}^{2\pi \mathrm{i} \vec{k} \cdot \vec{r}}, \tag{3.28}$$

which can be computed efficiently with FFT. In practice, $\Phi_{\vec{k}}$ computed in this way is sometimes not exactly positive and real. Taking the absolute value is advisable and is a negligible error.

It is also possible to simulate a field with a given spectrum $\Phi(\vec{\kappa})$. In principle it can be put it in the place of $\Phi_{\vec{k}}$, but to get the absolute scale right for all choices of grid resolution $N_1 \times \ldots \times N_d$ and size $L_1 \times \ldots \times L_d$ it is necessary to be a bit more precise. An approximation of the Fourier integral is

$$\begin{aligned} C(\vec{\rho}) = &\int \Phi(\vec{\kappa}) \mathrm{e}^{\mathrm{i} \vec{\kappa} \cdot \vec{\rho}} \mathrm{d}^d \vec{\kappa} \\ \approx &\Delta\kappa_1 \cdots \Delta\kappa_d \sum_{\vec{k}} \Phi(k_1 \Delta\kappa_1, \cdots, k_d \Delta\kappa_d) \mathrm{e}^{\mathrm{i}(\Delta\kappa_1 k_1 \rho_1, \cdots, \Delta\kappa_d k_d \rho_d)}, \end{aligned} \tag{3.29}$$

where $\Delta\kappa_i = \frac{2\pi}{L_i}$. So for (3.27) to have the right size we must set

$$\Phi_{\vec{k}} = \frac{1}{\Delta\kappa_1 \cdots \Delta\kappa_d} \Phi(k_1 \Delta\kappa_1, \cdots, k_d \Delta\kappa_d). \tag{3.30}$$

Note that this diverges as L^d so that the δ-function in (3.19) is reproduced as $L \to \infty$.

Convergence of the discretization (3.29) is a bit tricky. Increasing the number of grid points N only helps resolve large $\vec{k}$. Increasing the size of the volume L increases resolution for small $\vec{k}$ while cutting it off earlier for large $\vec{k}$. Satisfying any of these requirements is impossible when C and Φ are given by power laws since their transforms don't converge in the normal sense.

The computational complexity of this method is the same as for a d-dimensional FFT: $\mathcal{O}(N^d \log N)$. But, the FFT comes with a prize: the generated fields are automatically periodic. The covariance is only exact for relative distances smaller than $L/2$ in any direction.

3.4.2 Time evolution

In all previous work referenced in this thesis, the frozen flow hypothesis was used to introduce time dependence. The interpretation being that a fixed spatial structure is moved over a "measurement site" by a constant wind $\vec{v}$. The refractivity field will then satisfy

$$n(\vec{r}, t) = n(\vec{r} - t\vec{v}, 0). \tag{3.31}$$

But, that is not compatible with a finite volume. A natural solution is to use periodic fields, but then a highly artificial periodicity is introduced. The approach taken here is to use a periodic field with intrinsic time dependence.

In principle, the time variable is no different from the space variables. The refractivity can be considered to be a 4-dimensional random field instead of a 3-dimensional one. However, keeping $\mathcal{O}(N^4)$ numbers in memory at the same time is not an option. Instead it would be better if the next time step could be computed by updating the current field. With such a time evolution, the field is a Markov process; its future evolution depends only on its last known value. The approach taken in this thesis is a combination of (3.31) (with periodic boundary conditions) with a gradual randomization to gradually decorrelate the field enough to avoid periodicity.

Since each wavenumber mode $\hat{f}_{\vec{k}}$ is independent, it is reasonable to assume that they evolve independently as well. The evolution will have to be some bounded random walk so that the variance of $\hat{f}_{\vec{k}}$ remains $\Phi_{\vec{k}}$. The simplest Gaussian process with this property is the Ornstein-Uhlenbeck process (in discrete time it can be thought of as an AR(1)-process). It has covariance

$$C(\Delta t) = \Phi_{\vec{k}} \mathrm{e}^{-\gamma|\Delta t|}. \tag{3.32}$$

and its time evolution is given exactly by

$$\hat{f}_{\vec{k}}(t + \Delta t) = \beta \hat{f}_{\vec{k}}(t) + \sqrt{\Phi_{\vec{k}}(1 - \beta^* \beta)}\xi \tag{3.33}$$

where $\beta = \mathrm{e}^{-\gamma \Delta t}$ and ξ is a circular normal random variable with zero mean and variance 2. If γ is given a non-zero imaginary part (3.33) still works (but the covariance function must be modified). Then $\hat{f}_{\vec{k}}$ will resemble a noisy driven harmonic oscillator. A purely imaginary γ results in deterministic time evolution $\hat{f}_{\vec{k}}(t) = \mathrm{e}^{-\mathrm{i}\,\mathrm{Im}(\gamma)t} \hat{f}_{\vec{k}}(t)$. This kind of time evolution was used in 1970 by Kraichnan [28] for computational study of turbulent diffusion. But the timescale was chosen arbitrarily.

The simplest choice of γ is to make the field decorrelate in the time it takes the volume to move one period. A constant windspeed v_x in the x-direction requires $\gamma \propto v_x / L_x \frac{v_x}{L_x}$. Is this consistent with Kolmogorov theory? If γ is a real function of $\vec{k}$, then the spatiotemporal spectrum can be written as

$$\Phi(k, \omega) \propto \int_{-\infty}^{\infty} \mathrm{e}^{-\gamma(\vec{k})|t| + \mathrm{i}\omega t} \mathrm{d}t = \frac{-1}{\mathrm{i}\omega - \gamma(\vec{k})} + \frac{1}{\mathrm{i}\omega + \gamma(\vec{k})} = \frac{2\gamma(\vec{k})}{\gamma(\vec{k})^2 + \omega^2}. \tag{3.34}$$

Supposing that the variance of $\hat{f}_{\vec{k}}$ scales as in Kolmogorov theory

$$\Phi(k, \omega) \propto \varepsilon^{2/3} k^{-11/3} \frac{2\gamma(k)}{\gamma(k)^2 + \omega^2} = \varepsilon^{2/3} k^{-11/3} \omega^{-1} \frac{2\,(\gamma(k)/\omega)}{(\gamma(k)/\omega)^2 + 1}. \tag{3.35}$$

This is of the allowed form (3.14) only when $\gamma(k)/\omega$ is some function of α from (3.13). The only possibility is $\gamma(k) = \varepsilon^{1/3}k^{2/3}$. This choice causes small scales (large k) to decorrelate more rapidly than large scales (small k) as would be intuitively expected.

Moving the field can easily be done at the same time with a phase shift

$$\hat{f}_{\vec{k}}(t + \Delta t) = \hat{f}_{\vec{k}}(t)\mathrm{e}^{\mathrm{i}\vec{k}\cdot\vec{v}(t)\Delta t} \tag{3.36}$$

which can be done to to no additional cost when applying (3.33). Note that the velocity $\vec{v}$ is allowed to be time dependent.

However, more general wind phenomena can emulated. In the spatial domain, it is possible to translate each horizontal layer independently allowing for a height dependent wind, but it can be done more easily by mixed frequency and real space descriptions. Let $\tilde{f}_{\vec{k}_\parallel,z}$ be the 1D FFT along the z-direction of $\hat{f}_{\vec{k}}$ (roughly equivalent to a 2D FFT along the x- and y-axes of $f_{\vec{r}}$). Now,each layer can be moved separately using the update rule[8]

$$\tilde{f}_{\vec{k}_\parallel,z}(t + \Delta t) = \tilde{f}_{\vec{k}_\parallel,z}(t)\mathrm{e}^{\mathrm{i}\vec{k}_\parallel\cdot\vec{v}(z,t)\Delta t}. \tag{3.37}$$

Then $\tilde{f}_{\vec{k}_\parallel,z}$ can be transformed back to frequency space or to real space with an additional FFT.

There are two advantages of doing this in mixed space rather than real space: Firstly, it provides automatic, non-degrading interpolation when the field is translated a non-integer number of grid points. Secondly, it simplifies the programming. All the heavy lifting involving movement of data is performed by an efficient and reliable FFT implementation. When programming, it is sufficient to do element-wise multiplications and call the FFT routine.

This removes another assumption of the frozen flow hypothesis and enables measured wind profiles to be used, even into the planetary boundary layer.

Note that although (3.37) and (3.33) are both linear, they not commutative for arbitrary Δt. They do however commute assymptotically in the limit of small Δt.

3.4.3 Other methods

There are two other methods than the FFT-scheme presented above that would be feasible to generate 3D fields: *circulant embedding* and the *turning bands method*.

3.4.3.1 Circulant embedding

Computationally this is somewhat similar to the FFT method since it involves using FFTs to diagonalize covariance matrices. But the theoretical justification is very different. While the FFT method works directly with pairwise covariances, the circulant embedding approach works with covariance matrices. See [29] and [30] for mathematical and algorithmic details. Below follows a brief descripton.

When the grid coordinates are oriented in a natural way, covariance matrices have the Toeplitz symmetry property (in higher dimension, block Toeplitz). The

[8]Strictly, in the time-dependent case $\vec{v}(z,t)\Delta t$ should be replaced by $\int_t^{t+\Delta t}\vec{v}(z,t')\mathrm{d}t'$.

construction is then based on two theorems. The first being that any (block) Toeplitz matrix can be embedded in a (block) circulant matrix. A circulant matrix can be represented by a 1d sequence c_i of the same length as the matrix size. In the block circulant case, representation is a d-dimensional array $c_{i,j,\ldots k}$. The second theorem is that the d-dimensional DFT, written in unitary matrix form Q, diagonalizes the circulant matrices; i.e. $\Lambda = Q^H C Q$ is diagonal. Here Q^H is the hermitian transpose of Q or, equivalently, a matrix representation of the inverse DFT. Independent normal random variables can be generated in the diagonalized space and then transformed back to real space using an FFT.

It appears this method has only small advantages over the simplistic FFT method. It generates *aperiodic* fields, but so does the FFT method if the volume is truncated. Furthermore, there is an added complexity in the fact that the circulant embedding is sometimes not exactly positive definite. However, the biggest problem with this method for the intended application is that time evolution is difficult to achieve. For that periodicity is a *desired* property since translations have no artificial boundaries.

3.4.3.2 Turning bands method

The turning bands method seems to be common in geological applications. See [31] for an example implementation and further references. It can be seen as a special case of more general methods where a random field $f(\vec{r})$ is approximated as a sum of a large number of independent random basis functions

$$f(\vec{r}) = \sum_{i=1}^{n} \xi_i(\vec{r}). \tag{3.38}$$

The basis functions are chosen such that their structure functions have the sought after form in an ensemble sense

$$\left\langle \left(\xi_i(\vec{R} + \vec{r}) - \xi_i(\vec{R}) \right)^2 \right\rangle = D(\vec{r}). \tag{3.39}$$

But each individual ξ_i is usually highly regular, for example sinusoidal plane waves $\xi_i(\vec{r}) = A_i \cos(\vec{k_i} \cdot \vec{r} + \varphi_i)$ where A_i, $\vec{k_i}$ and φ_i are random variables. In the limit of large n, f converges towards a Gaussian random field with structure function $D(\vec{r})$.

The FFT-method detailed in Section 3.4.1 can be seen as a variant of this method where only the amplitudes and phases are randomized and the wavenumber takes all values on a grid. There are two significant differences: Firstly that the FFT-method is exact. Secondly that the sum in (3.38) can be evaluated very efficiently using FFT.

Generating regular d-dimensional grids with a turning bands method is relatively expensive since the cost is $\mathcal{O}(nN^d)$ and n needs to be much larger than $\log(N)$. The potential advantage of this method is that no grid is needed, only the parameters of the basis functions. If basis functions are chosen such that line integrals can be evaluated analytically, this would be a strong competitor to grid based methods for the application considered in this thesis.

3.5 Log-normal random fields

The use of zero mean fields above rests on the interpretation that the fields are a fluctuation from some mean $\mu(\vec{r})$ that varies on a much longer timescale. Keeping in mind that refractivity is a positive quantity, another aesthetically pleasing way to add a mean level is by exponentiation

$$n(\vec{r}) = \mu(z)\mathrm{e}^{\delta(\vec{r})} \tag{3.40}$$

where $\mu(z)$ is an expected mean refractivity profile and $\delta(\vec{r})$ is some random relative deviation. To emulate Figure 2.2, the variance of $\delta(\vec{r})$ can be chosen as $\frac{\sigma(z)}{\mu(z)}$, where σ is the standard deviation as a function of height. For nearby points the structure function of n is approximately

$$\left\langle \left(n(\vec{R}+\vec{r}) - n(\vec{r}) \right)^2 \right\rangle \approx \mu(Z) \left\langle \left(\mathrm{e}^{\delta(\vec{R}+\vec{r})} - \mathrm{e}^{\delta(\vec{r})} \right)^2 \right\rangle. \tag{3.41}$$

If the variance of δ is much smaller than 1 this can be Taylor expanded to obtain

$$\left\langle \left(n(\vec{R}+\vec{r}) - n(\vec{r}) \right)^2 \right\rangle \approx \mu(Z) \left\langle \left(\delta(\vec{R}+\vec{r}) - \delta(\vec{r}) \right)^2 \right\rangle. \tag{3.42}$$

Therefore n will have a structure function very similar to that of δ.

If $\delta(\vec{r})$ is simulated as a Gaussian random field, then $n(\vec{r})$ will be a log-normal random field. This is attractive for two reasons: Firstly because $n(\vec{r})$ is strictly positive. Secondly because the simulated signals will no longer be normal distributed, but might be able to replicate the non-Gaussian tails seen in Figure 2.8.

Chapter 4

Implementation

The python software developed in this thesis to generate atmospheric signal delay is provided at `https://github.com/hengin/StochasticTurbulenceGeneration`. It splits naturally into two modules: *generation* of refractivity fields and *integration* over such fields to produce signal delays. They are described in their own sections below. Furthermore a simple plotting utility for visualizing 3D fields was created using the open source python module PyQtGraph (`www.pyqtgraph.org`). The last section in this chapter features some "applications" of the software to situations characteristic of delay due to atmospheric water vapour.

4.1 Generation of refractivity fields

Two kinds of generators of 3D stochastic fields were implemented. The first handles quasi-homogeneous covariances of the form[1]

$$\langle n(\vec{r}_1)n(\vec{r}_2)\rangle = \sqrt{C_v(z_1)C_v(z_2)}C_h\left(\vec{r}_1 - \vec{r}_2\right). \tag{4.1}$$

For instance, this can be used to realize the function

$$\langle n(\vec{r}_1)n(\vec{r}_2)\rangle = \exp(-(z_1+z_2)/H)\frac{C_n^2 L^{4/3}}{L^{2/3} + \left|\vec{r}_1 - \vec{r}_2\right|^{2/3}} \tag{4.2}$$

used by Nilsson & Haas in [7] by setting $C_v(z) = \exp(-2z/H)$. Since the homogeneous part C_h can depend separately on all three components of the difference $\vec{r}_1 - \vec{r}_2$, some non-isotropic covariance functions can also be realized. One example is

$$\langle n(\vec{r}_1)n(\vec{r}_2)\rangle = \frac{C_n^2 L^{4/3}}{L^{2/3} + \left(\left|\vec{r}_1 - \vec{r}_2\right|^2 + (C_0 - 1)(z_1 - z_2)^2\right)^{1/3}} \tag{4.3}$$

which was mentioned in [13, eq. (4.10)] as a way of introducing height dependence. Note however, that (4.3) can be simulated even with only an isotropic generator by rescaling the z-coordinate. The implementation has a cost of $\mathcal{O}(N^3 \log N)$ operations per realization and uses $\mathcal{O}(N^3)$ memory.

[1] Actually, C_v can be an arbitrary function of the 3D position with little additional cost, but that has no motivation in this context.

The second implementation handles horizontally homogeneous covariances with arbitrary vertical dependence

$$\langle n(\vec{r}_1)n(\vec{r}_2)\rangle = C\left(z_1, z_2, \left|\vec{r}_{1\parallel} - \vec{r}_{2\parallel}\right|\right). \tag{4.4}$$

Very few such functions occur in the literature that cannot be written on the form (4.1). Therefore it is only described in Appendix A. Note that using the horizontal distance $\left|\vec{r}_{1\parallel} - \vec{r}_{2\parallel}\right|$ is equivalent to using the full distance since

$$C\left(z_1, z_2, |\vec{r}_1 - \vec{r}_2|\right) = C\left(z_1, z_2, \sqrt{\left|\vec{r}_{1\parallel} - \vec{r}_{2\parallel}\right|^2 + (z_1 - z_2)^2}\right) \tag{4.5}$$

which is of the form (4.4). This method has a setup cost of $\mathcal{O}(N^4 \log N)$ with an additional cost of $\mathcal{O}(N^4)$ operations per realization. The memory usage is $\mathcal{O}(N^3)$.

4.1.1 Realization of 3D refractivity fields

The procedure to generate fields of type (4.1) is based on Section 3.4.1. A complete description follows. In order to be consistent with Section 3.4.1, $\vec{r} = (x, y, z)$ and $\vec{k} = (k_x, k_y, k_z)$ here denote integer coordinates.

1. Choose resolution $N_x \times N_y \times N_z$ and domain size $L_x \times L_y \times L_z$. Let $\Delta x = L_x/N_x$, $\Delta y = L_y/N_y$ and $\Delta z = L_z/N_z$.
2. If a spectrum $\Phi(\vec{\kappa})$ was given as input instead of C_h, go to 4. Else sample C_h on a 3D grid with coordinates in the order

$$0,\ \Delta x,\ 2\Delta x, \ldots \pm (N_x/2)\Delta x,\ -(N_x/2-1)\Delta x, \ldots -\Delta x \tag{4.6}$$

   (assuming N_x is even) and similarly in the y and z directions.
3. Compute $\Phi_{\vec{k}}$ using an FFT from the array above with normalization factor $1/(N_x N_y N_z)$. Then proceed to step 5.
4. Sample $\Phi(\vec{k})$ on a 3D grid with coordinates ordered as

$$0,\ \Delta k_x,\ 2\Delta k_x, \ldots \pm (N_x/2)\Delta k_x,\ -(N_x/2-1)\Delta k_x, \ldots -\Delta k_x, \tag{4.7}$$

   where $\Delta k_{x/y/z} = 2\pi/L_{x/y/z}$. Again this is for even N_x and the y and z axis have the same ordering. The array $\Phi_{\vec{k}}$ is obtained by dividing the sampled array by $\Delta k_x \Delta k_y \Delta k_z$.
5. Generate an array $\tilde{n}_{\vec{k}}$ of $N_x N_y N_z$ independent normal, zero mean, pseudorandom variables where the real and imaginary parts both have variance $\Phi_{\vec{k}}$.
6. By FFT compute the sum

$$\tilde{n}_{\vec{r}} = \mathrm{Re}\left(\sum_{\vec{k}} \hat{n}_{\vec{k}} \exp(2\pi \mathrm{i}\vec{k}\cdot\vec{r})\right). \tag{4.8}$$

7. Multiply the $\tilde{n}$ array by the inhomogeneity to obtain the output array

$$n_{\vec{r}} = \sqrt{C_v(z\Delta z)}\tilde{n}_{\vec{r}}. \tag{4.9}$$

4.1.2 Time evolution

When evolving a field in time, the field is stored in the frequency domain in the form of the $\hat{n}_{\vec{k}}$-array. Whenever the field is needed in real space it can be generated using steps 6 and 7 above.

Three forms of time evolution have been implemented: frozen field update (3.31), intrinsic Markovian evolution (3.33) and convection by an arbitrary horizontal wind

$$n(\vec{r}, t + \Delta t) = n(\vec{r} - \Delta t \vec{v}_{\parallel}(z, t), t). \tag{4.10}$$

The last one must be combined with an intrinsic update for n to reach a stationary distribution. Note that this will make the distribution non-isotropic when there are vertical gradients in $\vec{v}_{\parallel}$.

Frozen field updates are done with a simple phase shift in the frequency domain

$$\hat{n}_{\vec{k}}(t + \Delta t) = \hat{n}_{\vec{k}}(t) \mathrm{e}^{-\mathrm{i}\vec{k}\cdot\vec{v}\Delta t}. \tag{4.11}$$

To see that this recreates (3.31) simply insert it into (4.8). Intrinsic updates are performed using

$$\hat{n}_{\vec{k}}(t + \Delta t) = \mathrm{e}^{-\gamma(\vec{k})\Delta t}\hat{n}_{\vec{k}}(t) + (1 - \mathrm{e}^{-2\Delta t \mathrm{Re}\,\gamma(\vec{k})})^{1/2}\xi_k(t), \tag{4.12}$$

where $\xi_k(t)$ has independent normal distributed real and imaginary parts; each havin mean zero and variance $\Phi_{\vec{k}}$. The "dispersion relation" $\gamma(\vec{k})$ is arbitrary as long as it has non-negative real part. The exact form is the subject of turbulence modelling, see Section 3.4.2.

For height dependent wind, the frequency field $\hat{n}_{\vec{k}}$ is first transformed with an FFT along the z-axis into $\tilde{n}_{\vec{k}_{\parallel},z}$. Then a phase shift

$$\tilde{n}_{\vec{k}_{\parallel},z}(t + \Delta t) = \tilde{n}_{\vec{k}_{\parallel},z}(t) \mathrm{e}^{-\mathrm{i}\vec{k}_{\parallel}\cdot\vec{v}(z,t)\Delta t} \tag{4.13}$$

is applied. Finally $\tilde{n}_{\vec{k}_{\parallel},z}(t + \Delta t)$ is transformed back into $\hat{n}_{\vec{k}}(t + \Delta t)$ using an inverse FFT along the z-axis.

4.1.3 Validation

There are many factors of 2 and 2π in the derivations of Chapter 3 where the implementation can go wrong. Hence it is vital to verify that it produces fields with the right structure.

A good test case is a Gaussian covariance function

$$C(\vec{r}) = \exp(-\left|\vec{r}\right|^2 / a^2) \tag{4.14}$$

since its spectrum has a simple analytical expression

$$\Phi(\vec{k}) = \frac{1}{(2\pi)^3} \int \exp\left(-\left|\vec{r}\right|^2 / a^2 - \mathrm{i}\vec{k}\cdot\vec{r}\right) \mathrm{d}^3\vec{r} = \frac{a^3}{8\pi^{3/2}} \exp\left(-a^2 \left|\vec{k}\right|^2 / 4\right). \tag{4.15}$$

Furthermore, both the covariance and spectrum are highly smooth and decay quickly to zero. For later purposes, note that in the limit $a = 0$, C remains finite but $\Phi = 0$.

This corresponds to white noise of power 0; for each $\vec{r} \in \mathbb{R}^3$, $n(\vec{r}) \sim N(0,1)$ and is independent from other part of the field.

The implementation is considered correct if it generates fields with the correct structure function $\langle (n(\vec{r}_1) - n(\vec{r}_2))^2 \rangle =: S(|\vec{r}_1 - \vec{r}_2|)$. Using (4.14) the expected structure is.

$$S(r) = 2(C(0) - C(r))2\left(1 - \exp(-r^2/a^2\right). \tag{4.16}$$

This is enough since $n_{\vec{r}}$ by constryction follows a zero mean multivariate normal distribution and is therefore entirely determined by its covariance/structure function.

To make the test as unbiased as possible, it was performed by first choosing 1000 index differences $(\Delta i, \Delta j, \Delta k)$ at random (with a bias towards small differences). Then 1000 realizations of the random field were computed. For each realization, 100 origin indices (i, j, k) were selected at random (uniformly). For each combination of origin and index difference the quantity $(n_{i+\Delta i, j+\Delta j, k+\Delta k} - n_{i,j,k})^2$ was added to a position corresponding to $(\Delta i, \Delta j, \Delta k)$ in an accumulator array. Afterwards each element of the accumulator was divided by $100 \cdot 1000$ to obtain an estimate of $S_{\Delta i, \Delta j, \Delta k}$. Or equivalently $S(r)$, where

$$r = \sqrt{(L_x \Delta i / N_x)^2 + (L_y \Delta j / N_y)^2 + (L_z \Delta k / N_z)^2}. \tag{4.17}$$

The resolution and size parameters were chosen as $N_x = 100$, $N_y = 128$, $N_z = 64$ and $L_x = 20$, $L_y = 30$, $L_z = 25$ and $a = 2$ respecively. This produces an isotropic covariance in real space, but an asymmetric one in index space which is good for testing correctness. Note that the maximum value of r in this geometry is

$$r_{\max} = \frac{1}{2}\sqrt{L_x^2 + L_y^2 + L_z^2} \approx 22. \tag{4.18}$$

The test was performed both with a covariance (4.14) and with a spectrum (4.15) as input to the simulation. The results are plotted in Figure 4.1 together with the theoretically expected structure function (4.16). The agreement is perfect within statistical uncertainty.

To ensure that realizations are reasonable it is a good idea to visualize them. This is done in Figure 4.2 for two choices of the length scale a with resolution 128^3 in a volume of size 10^3. Note that these fields are not fractal; there is a single relevant length scale a. For shorter distances than a the field is highly smooth; for longer distances it looks like uncorrelated noise.

An example of the utility of visualization is demonstrated in Figure 4.3. In the left panel, clear cubical artefacts are visible; despite the fact that the input covariance is isotropic. This is immediately clear from the visualization, but would be less obvious for one dimensional statistical measures. The reason for the artefacts is that the covariance function does not reach zero quickly enough within the cube which introduces discontinuities when it is approximated as a periodic function. The effect is absent when using the exact spectrum (4.15) as input (see right panel). This corresponds to using the periodized covariance function

$$C_{\text{periodic}}(\vec{r}) \coloneqq \sum_{\vec{R}} C(\vec{R} + \vec{r}) \tag{4.19}$$

where $\vec{R}$ is a Bravais lattice vector (i.e. $\vec{R} = (kL_x, lL_y, mL_z)$ for any and all $k, l, m \in \mathbb{Z}$).

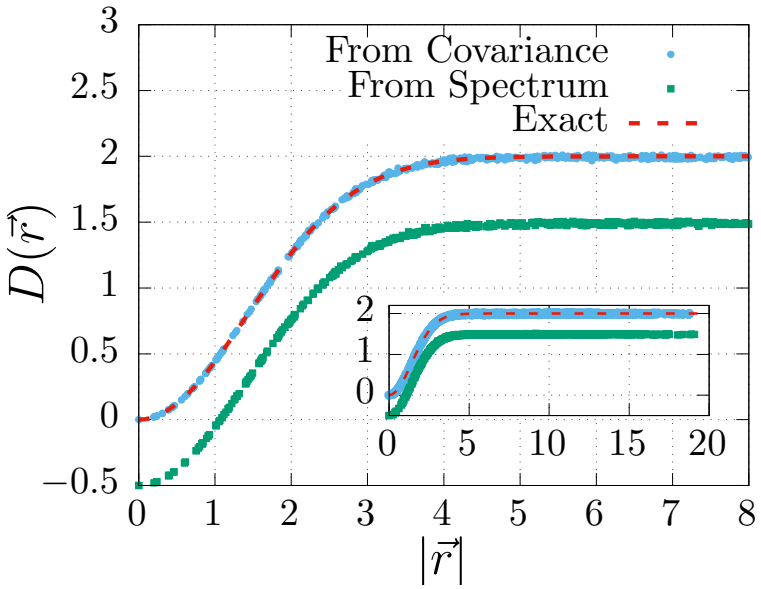

Figure 4.1: Numerically computed structure functions compared with the theoretically expected one. "From Covariance" is computed when the covariance (4.14) is used as input while "From Spectrum" is computed with the spectrum (4.15) as input. The points "From Spectrum" have been shifted downwards 0.5 units to improve readability. The inset shows the same data for the full range of distances.

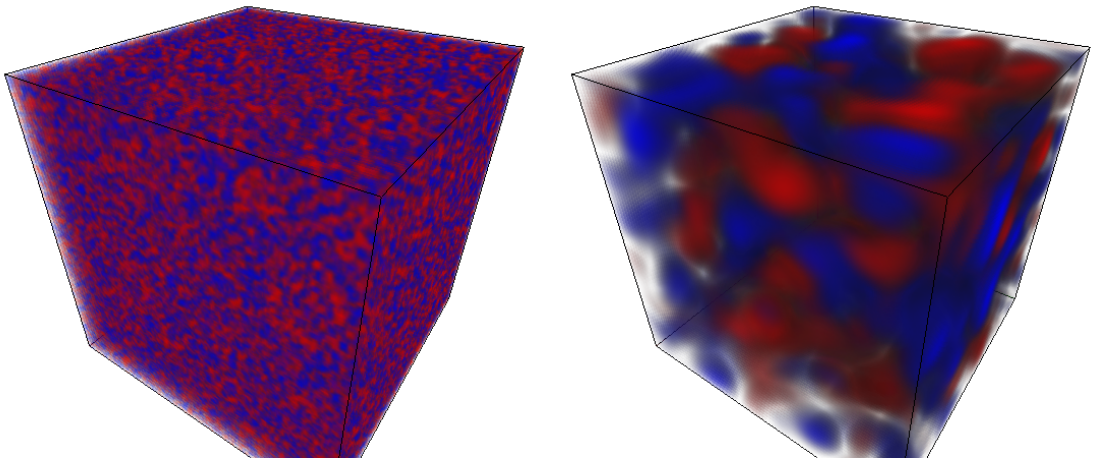

(a) $a = 0.1$ **(b)** $a = 1$

Figure 4.2: Gaussian random fields in 3D with input covariances given by (4.14) with resolution 128^3 and box size $L_x = L_y = L_z = 10$. The field is positive in red regions and negative in blue regions while values close to zero are transparent. In between the colour is interpolated smoothly.

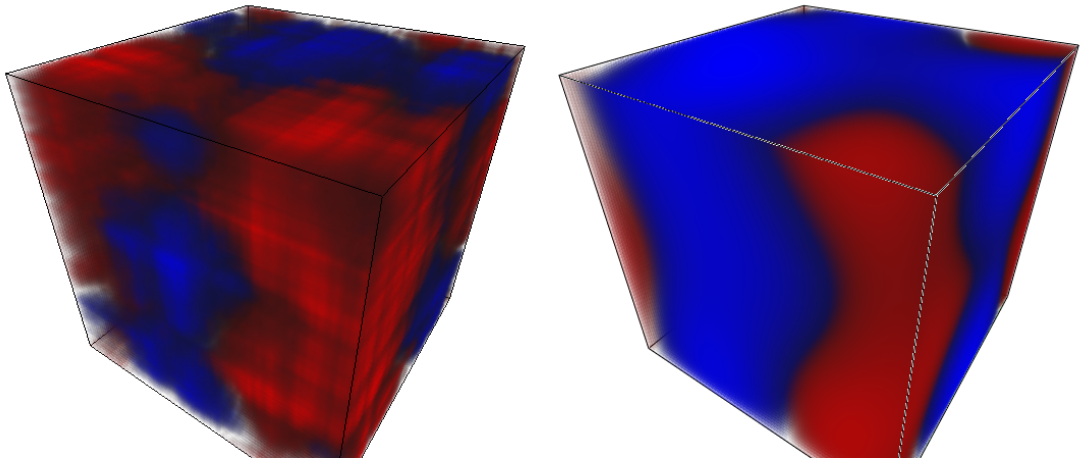

(a) From covariance **(b)** From spectrum

Figure 4.3: Illustration of aliasing when the covariance function is discontinuous because the box is too small. In both panels $a = 5$, but (a) was generated by first computing an FFT of (4.14) while (b) was generated directly from (4.15).

4.2 Integration

In the following, the excess refractivity field will be denoted $n(\vec{r})$. There are two kinds of signal delays to be generated from the field. The first is delays along straight lines, as appropriate to simulate GNSS and VLBI signals. The other is integrals over cones, as appropriate to simulate radiometer signals. In the first case the excess delay along a straight path starting at $\vec{r}_0$ in the direction $\hat{n}$ is

$$\Delta L = \int_0^\infty n(\vec{r}_0 + s\hat{n})\mathrm{d}s. \tag{4.20}$$

In the second case the average delay over radiometer beam is

$$\Delta L = \int_0^\infty \int_{S^2} n(\vec{r}_0 + \hat{n}'r)g(\theta)\mathrm{d}\Omega(\hat{n}')\mathrm{d}r, \tag{4.21}$$

where the inner integral with respect to the variable $\hat{n}'$ is taken over the unit sphere and r is a radial coordinate. Furthermore, θ is the angle between $\hat{n}$ and $\hat{n}'$ and g is the antenna gain with normalization

$$\int g(\theta)\mathrm{d}\Omega = 1. \tag{4.22}$$

These integrals must be approximated when n is sampled at a regular grid.

One way to do so is to interpret the full field as being constant in (parallelepipedal) cells centred about grid-points. Line integrals can then be approximated with

$$\Delta L = \sum_i \ell_i n(\vec{r}_i), \tag{4.23}$$

where the sum goes over cells intersected by the path and ℓ_i is the length that the path traverses through cell i. See Figure 4.4 for a 2D illustration. However,

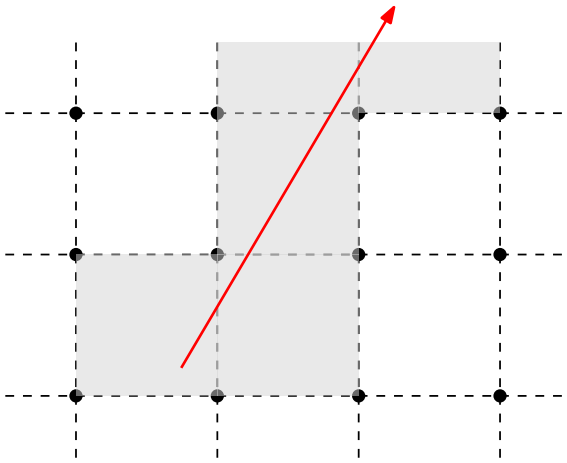

Figure 4.4: Illustration of line/cell intersection in 2D. The red arrow shows the integration path and the shaded squares show where the integration is performed.

it was observed that artefacts appeared for paths aligned with the grid, unless the resolution was very high. The artefacts were substantially reduced by interpolation of the field.

Suppose n is sampled at regularly spaced points $n(i\Delta x, j\Delta y, k\Delta z)$ with $1 \leq i \leq N_x$ (and so forth). A continuous extension of n can be obtained by setting

$$n(\vec{r}) = \sum_i w_i(\vec{r}) n(\vec{r}_i), \tag{4.24}$$

where $w_i(\vec{r}) \geq 0$, $\sum_i w_i(\vec{r}) = 1$ and the the sum goes over some set the corners closest to the point $\vec{r}$. Typically these are the 8 closest gridpoints that form the cornest of a "cube" containing $\vec{r}$. For a 2D illustration, see Figure 4.4.

There are many ways to choose the weight functions. Perhaps the simplest is to use trilinear weights. These are cubic polynomials that are 1 at some corner and decrease linearly along each axis to zero. It is convenient to represent these in normalized coordinates where (0,0,0) and (1,1,1) are diagonally opposed cell corners. The weight functions of trilinear interpolation are of the general form

$$w_{ijk}(x, y, z) = (i + (1 - 2i)x)\,(j + (1 - 2j)y)\,(k + (1 - 2k)z)\,, \tag{4.25}$$

where i, j and k are all 0 or 1. Two examples are

$$w_{000}(x, y, z) = xyz \text{ and } w_{011} = x(1 - y)(1 - z). \tag{4.26}$$

Finding all cells intersected by a line is done with a standard ray tracing algorithm. The integral over the interpolated field can be written

$$\Delta L = \sum_{i \in \text{cells}} \sum_{j \in \text{corners}(i)} n(\vec{r}_{i,j}) \int w_{i,j}(\vec{r}_0 + \hat{n}s)\mathrm{d}s \tag{4.27}$$

where the weight integrals can be evaluated analytically. They are polynomials of coordinates of line-cell intersections. This enables the method to be used for line parameters that change each simulation step. This could be useful to generate

signals in VLBI/GNSS geometry where the sources move continuously across the sky.

The procedure is a little different for integration over radiometer cones. There the interpolated delay takes the form

$$\Delta L = \sum_{i \in \text{cells}} \sum_{j \in \text{corners}(i)} n(\vec{r}_{i,j}) \int_i w_{i,j}(\vec{r}) \frac{g(\theta)}{r^2} \mathrm{d}V \tag{4.28}$$

so a large number of volume integrals

$$\int_i w_{i,j}(\vec{r}) \frac{g(\theta)}{r^2} \mathrm{d}V = \int_{x_{i,j}}^{x_{i,j}+\Delta x} \int_{y_{i,j}}^{y_{i,j}+\Delta y} \int_{z_{i,j}}^{z_{i,j}+\Delta z} \frac{w_{i,j}(x,y,z) g(\theta(x,y,z))}{x^2+y^2+z^2} \mathrm{d}x\mathrm{d}y\mathrm{d}z \tag{4.29}$$

have to be evaluated. Note that Cartesian coordinates $\vec{r} = (x, y, z)$ are used since that is easier to work with for parallelepipedal cells. In order to suppress alignment artefacts and other systematic errors, the integration must be done accurately. Typically this takes significantly more time than generating a refractivity field, so it is best to precompute the integrals. Then the delay for a specific field can be efficiently computed using the stored weights as $\Delta L = \sum_i w_i n(\vec{r}_i)$. This means that only a fixed number of directions/positions can be used. Howerver, that is not unrealistic. Radiometers usually have an observation schedule that cycles through a relatively small number of orientations.

Usually the antenna gain $g(\theta)$ is approximated as a Gaussian with a given half power beam width

$$g(\theta) = g(0) \exp\left(-\ln(2)(2\tilde{\theta}/\text{hpbw})^2\right), \tag{4.30}$$

where $\tilde{\theta}$ is the angular deviation from the central direction. Since the integral over all angles should be one and hpbw is usually small, we have

$$g(0)^{-1} \approx \int_0^\infty \exp\left(-\ln(2)(2\tilde{\theta}/\text{hpbw})^2\right) 2\pi\tilde{\theta}\mathrm{d}\tilde{\theta} = \frac{\pi \text{hpbw}^2}{4\ln 2}. \tag{4.31}$$

The antenna gain is usually sharply peaked with a beamwidth of a few degrees. This means that most of the weights will be very small and do not need to be considered. It is possible to find the most important cells from geometry, but getting the indices right when programming is painstaking. A more automatic approach is the following:

1. Choose a cutoff c, all weights below this value will be discarded. The order of magnitude for weights in the farthest cells is $\frac{\Delta V}{L^2} g(0)$ so a good choice is $c = \varepsilon \frac{\Delta V}{L^2} g(0)$ where ε is some small number, say 10^{-3}.
2. Keep two sets `visited_cells` and `cells_to_visit`. Add the radiometer origin to `cells_to_visit`.
3. If `cells_to_visit` is empty we are done. Otherwise pop the next cell in `cells_to_visit` and add it to `visited_cells`. Evaluate the weight integrals for the current cell and store them. If any weight was larger than c, add all neighbours not already in `visited_cells` to `cells_to_visit`.
4. Repeat step 3 until `cells_to_visit` is empty.

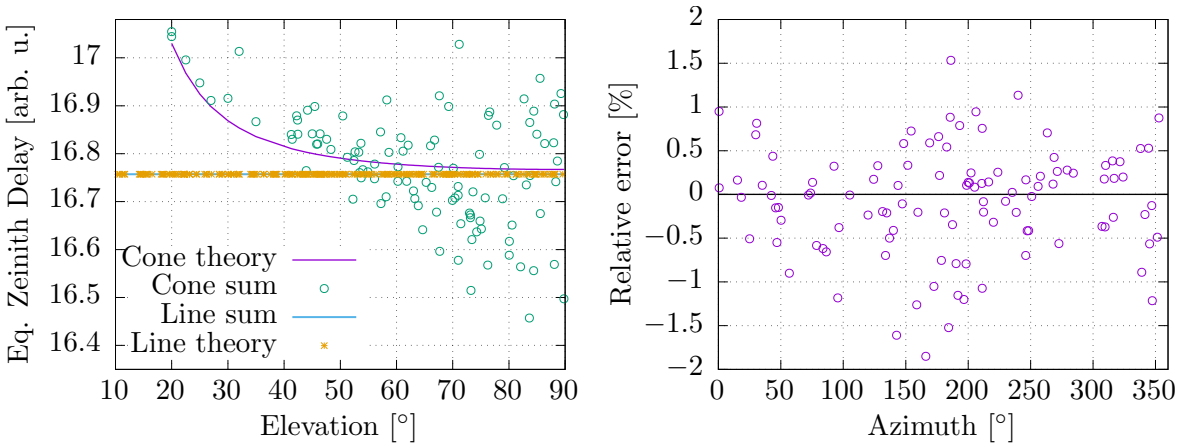

Figure 4.5: (Left) Elevation dependence of the equivalent zenith delay for a spatiallly constant refractivity (sum weights) together with expected values. (Right) Relative error of sums of cone weights as a function of azimuth.

4.2.1 Optimal interpolation – Kriging

For well resolved fields, higher accuracy might be achieved by using a technique called Kriging [32, pp. 2–3]. This is a statistical treatment of interpolation. Assuming the covariance function is known, the unknown value $n(\vec{r})$ can be estimated optimally in the least squares sense from a given set of values, such as at cell corners or some larger computational molecule. This was tested, but it was found to underestimate variances. This is natural since the method finds an unbiased minimal variance estimator.

4.2.2 Validation

The first check is if the sum of weights has the correct value. This corresponds to the delay through a constant refractivity field. For lines the sum of weights should equal the exact value $L_z/\sin\theta$ to within numerical precision (θ is elevation). For cone weights it is the average of $L_z/\sin\theta$ taken over varying θ in the radiometer cone. For easier comparison it is a good idea to multiply with the inverse mapping function $\sin(\theta)$ (see Section 2.4) to remove the dominant source of variation. This has been done with the sum of radiometer and line weights in Figure 4.5.

The test protocol was to randomly choose an elevation, an azimuth and an origin in the bottom plane. If the method works correctly, then neighbouring points in a sum of weights vs elevation plot should be very close in value. If there is an azimuth or origin dependence, some spread is to be expected. The line weights are exact to within numerical precision. On the other hand, the sums of cone weights show errors as large as 2 %. It is likely that the errors for individual cells are even larger. Notably the errors show no azimuth dependence. This can be considered an unresolved bug in the implementation.

It is also important to test whether or not the integrals have the right statistics. In Figure 4.7 the observed variance for line integrals over fields with Gaussian covariance (see (4.14) and Figure 4.2) is plotted. When the length scale is small

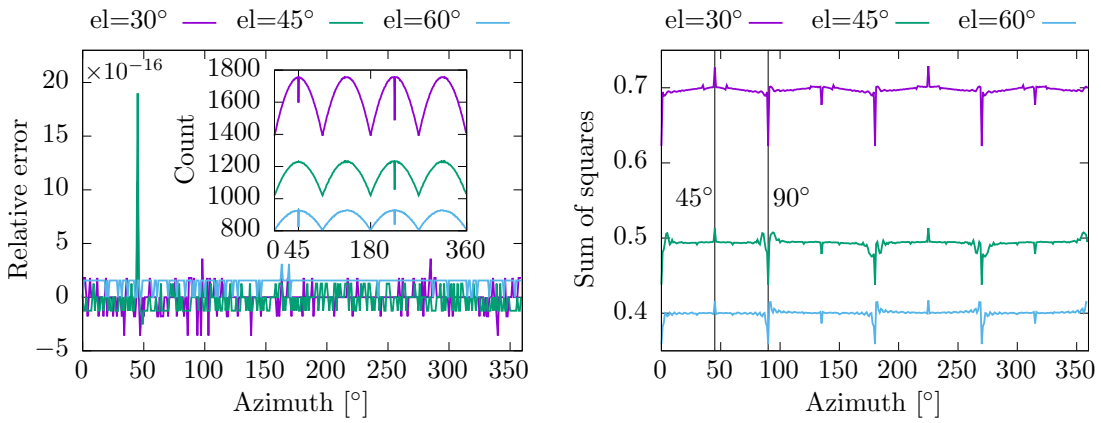

(a) $\sum_i w_i$ and $\sum_i 1$ **(b)** $\sum_i w_i^2$

Figure 4.6: (a) The sum of line weights $\sum_i w_i$ is azimuth independent to numerical precision. The inset shows the number of weights as a function of azimuth. The squared weight sum varies with the number of intersected cells, but more so with respect to how close to an axis the line passes.

compared to the grid spacing ($10/128 \approx 0.08$) the field is essentially white noise and all points are essentially independent. Hence the variance is expected to be

$$\mathrm{Var}\left[\sum_i w_i n_i\right] = \sum_i \mathrm{Var}\left[w_i n_i\right] = \sum_i w_i^2 \tag{4.32}$$

since the variance of each sample 1. Here the geometry should be the only factor influencing the result. This can be seen in the $a = 0.01$ case where downward spikes appear at multiples of 90° when the lines are aligned with the grid. This corresponds precisely to the sum of squares plotted in Figure 4.6b. Note that the absolute values are very close and not just the shape. When the field is smooth enough for interpolation on the grid to make sense (even $a = 0.1$ seems to be enough), the effect is reduced to a large degree.

Another problem appears when the correlation length a is comparable to the dimensions of the volume (a cube of side length 10). This is illustrated in the $a = 2$ case in Figure 4.7 where low enough elevations have maxima at grid alignments ($n \cdot 90°$ and $45° + n \cdot 90°$). The reason is periodicity of the field. The volume is not flat enough to support elevations lower than 45° without wrap-around. For some azimuths the line comes closer to its tail than others in which case there is a large correlation and hence a high signal. The cure for this problem is to use a volume wide enough accommodate all elevations of interest.

4.3 Example applications

The fractal structure of fields generated from Kolmogorov inspired structure functions such as

$$S(r) \propto \frac{r^{2/3}}{L^{2/3} + r^{2/3}} \tag{4.33}$$

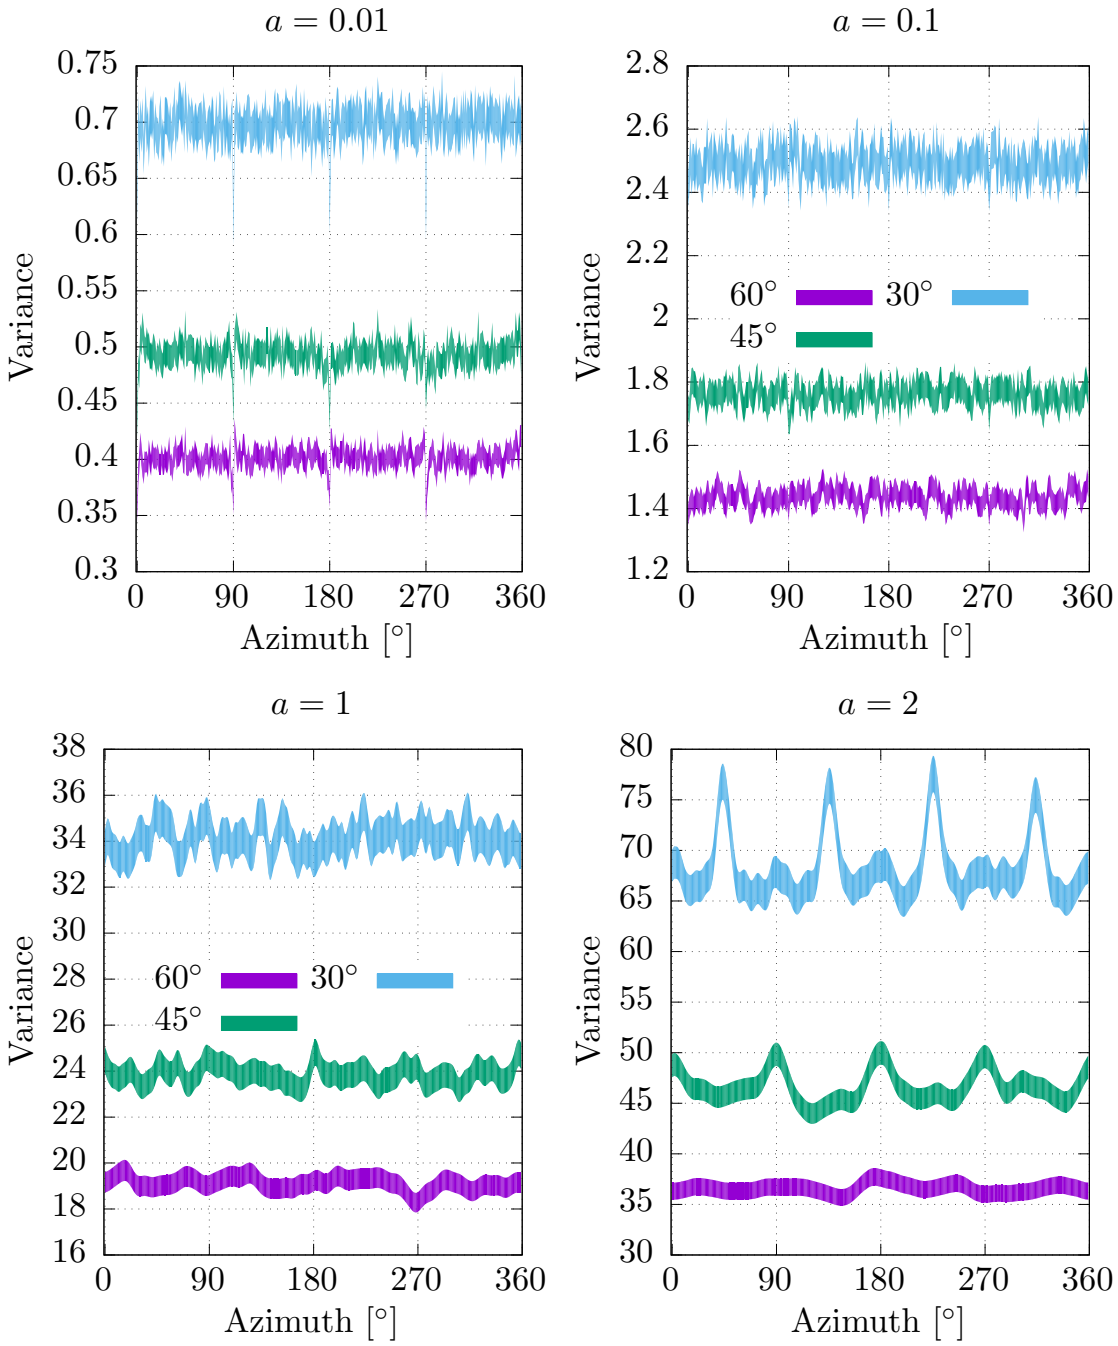

Figure 4.7: Variance of line integral through Gaussian fields (4.14) (resolution 128^3 size 10^3) from 10^4 realizations with 90 % confidence intervals. "Under-resolution" is evident from the spikes of $a = 0.01$. At $a = 2$ periodicity causes abnormally large variance in certain directions.

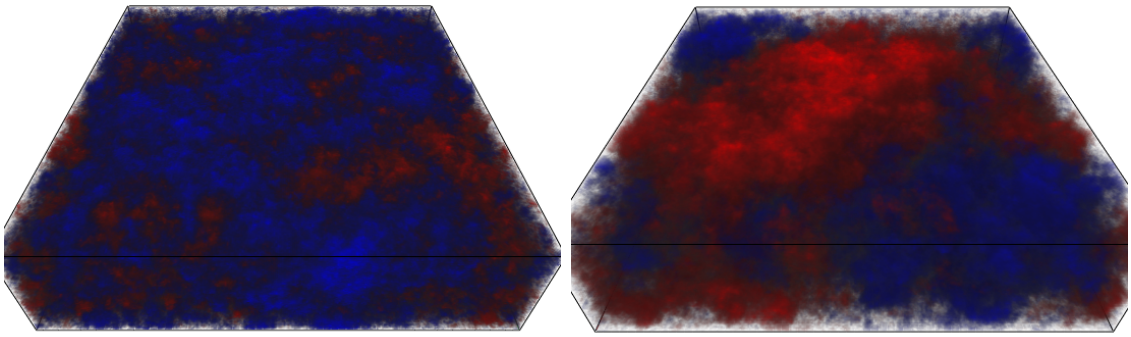

(a) $L = 1$ **(b)** $L = 5$

Figure 4.8: Visualization of a refractivity field with covariance given by (4.2) with $H = 0$. The grid resolution is $256 \times 256 \times 128$ and the physical size is $40 \times 40 \times 10$ for two different regularization lengths L.

can be glanced in Figure 4.8. Visually, very little difference is seen when the normalizing length scale is increased from $L = 1$ to $L = 5$. In both cases there are coherent structures of all available length scales. Note that L has no effect on the shortest length scales since $S(r) \propto r^{2/3}$ for small enough r, no matter what L is. Its only effect is to set a soft limit for the longest scales before decorrelation. However, since the correlation only decreases as $r^{-2/3}$ it takes a lot for the fields to look fully noisy; even for L comparable with the grid spacing, some correlations can be visually perceived.

The intended use of the software is to simulate delay time series. In Figure 4.9 such a timeseries is analyzed in the same way as two GPS timesseries were in Figure 2.8. The simulation used covariance of the form (4.2) with $L = 1\,\text{km}$, $H = 2\,\text{km}$ and $C_n^2 = 10^{-14}\,\text{m}^{-2/3}$. Furthermore, a log-normal field was generated at the same time using similar parameters and according to the method described in Section 3.5. The grid had resolution $512 \times 256 \times 64$ and the physical volume was $40\,\text{km} \times 20\,\text{km} \times 5\,\text{km}$. Time updates were done by periodically translating the field one grid point in the x-direction per timestep in addition an intrinsic update with $\vec{k}$-independent decorrelation parameter $\gamma = \frac{1}{512}$. An equivalent ZWD timeseries was generated by averaging the equivalent zenith delays from lines of azimuth angles 15° apart at elevations 30°, 45° and 60° at each timestep. The simulation was run for 5120 timesteps, corresponding to 10 periods.

As expected, the Gaussian field produces normal distributed ZWD increments. Anything else would have required a fundamental error in the implementation. The log-normal field produces highly non-Gaussian fields even for the smallest Δt. A smaller relative oscillation could have been used to more closely match Figure 2.8.

The structure functions scale as $\Delta t^{4/3}$ for both the log-normal and normal fields. This provides some validity to the reasoning in Section 3.5. Also note that other scalings can be achieved by changing the ratio $L : H$ and by using a $\vec{k}$-dependent γ. Finally, as can be seen by a minimum of the structure function at $\Delta t = 512$, the chosen value of γ was not enough to fully decorrelate the field in one period.

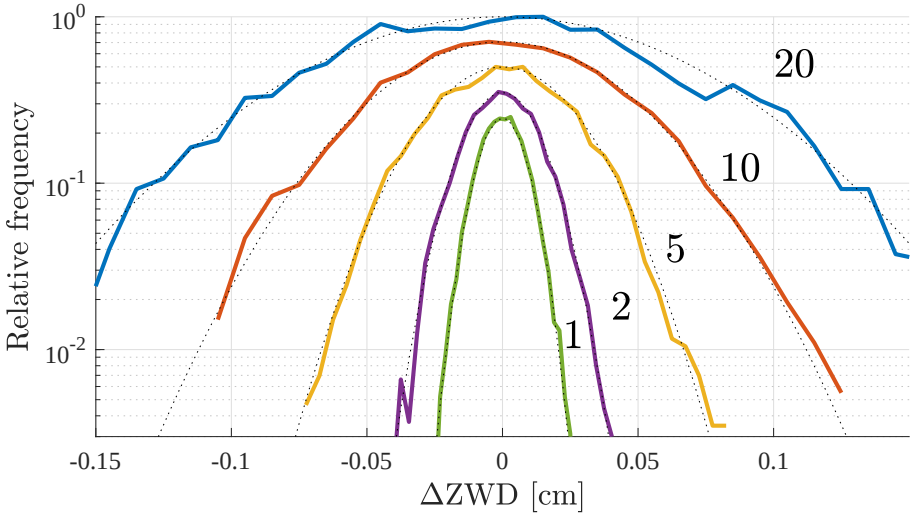

(a) Simulated PDF of ΔZWD (Normal)

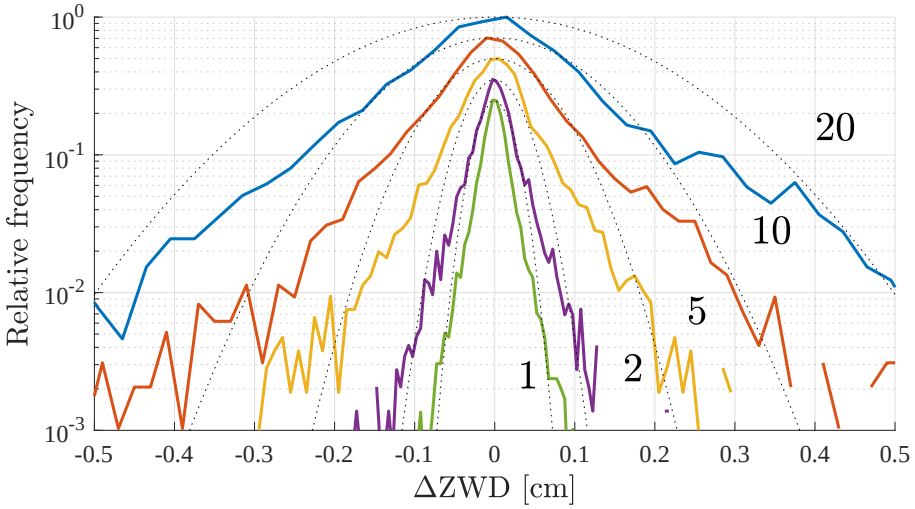

(b) Simulated PDF of ΔZWD (Log-Normal)

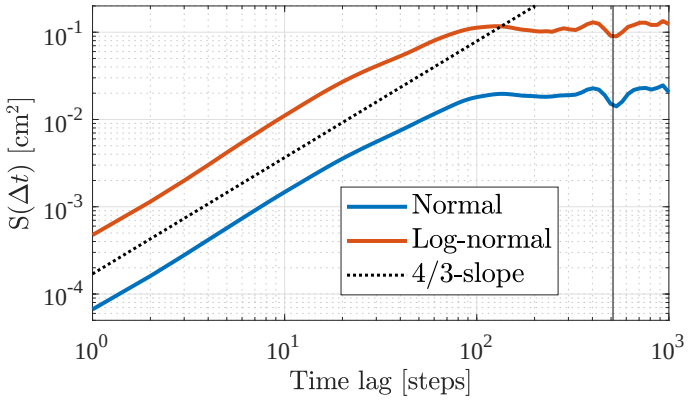

(c) Structure functions.

Figure 4.9: Replication of Figure 2.8 using two kinds of simulated data. See text for simulation parameters. (a) & (b) show estimated probability distributions of ZWD increments for various temporal separations. The numbers indicate Δt in units of simulation steps. (c) Shows the structure functions from both simulations compared to $\Delta t^{4/3}$-scaling. The vertical line at $\Delta t = 512$ indicates where artificial periodicity is expected.

Chapter 5

Discussion

As was mentioned in Section 3.3, the reason for using Gaussian fields is that it is the most tractable way to compute a random field with a given second order structure function. But there are problems with doing so. The usual motivation for using normal distributions to model measurements is the central limit theorem. In very simplified terms, it states that the sum of many independent disturbances is asymptotically normal distributed. However, this asymptotic limit really only applies to the bulk of the distribution. The more standard deviations away from the mean you look, the slower the convergence is. Furthermore, the Navier-Stokes equation (which describes fluid motion) contains non-linear terms. Therefore an effective stochastic description should contain both additive and multiplicative noise. With just additive noise, normal distributions are expected while with just multiplicative noise, log-normal distributions are expected. In reality we should get something in between.

Two modes of time evolution beyond the frozen-flow hypothesis were introduced in Section 3.4.2. The most questionable of the two is height dependent wind. The whole idea behind using a stochastic model for a deterministic phenomenon is based on excluding enough details that the remaining observables seem random. Some effects of convection by varying wind can be considered to already be included in the Markov update step. Adding additional shearing motion might not be meaningful.

Simulating on a grid is not the only way to obtain time-dependence outside the frozen flow hypothesis. The spectrum (3.35) with Kolmogorov scaling inserted can be written as

$$\Phi(k,\omega) \propto \varepsilon k^{-3}\omega^{-2}\left(C + \varepsilon^{2/3}k^{4/3}\omega^{-2}\right)^{-1}, \tag{5.1}$$

where C is some constant. This spectrum could be used directly to compute signal covariances. These can then be used in methods such as those in [8, 9] to generate realisations of random delay signals.

In some special cases, it is possible to simulate with higher resolution (or faster) by setting the resolution to 1 in some direction. If all observations take place in a single vertical plane then that can be done in one of the horizontal directions. Another possibility is to make the simulation horizontal, and use a mapping function to obtain elevation dependence.

Finally, what are the prospects of validating the simulations against reality? There are many parameters and the phenomenon is inherently stochastic, which makes

comparison difficult. Especially since measurements themselves have noise and the parameters can change in less time than it takes to gather enough statistics to test the model. Validating turbulence models is a research field all on its own. Still, one possibility is to compare stochastic refractivity fields against Large Eddy Simulations (LES), which solves the Navier-Stokes equations approximately for turbulent flows. Then all parameters can in principle be controlled and full information about fields are available.

5.1 Conclusions

There are many reasons to doubt the validity of models of atmospheric refractivity fields based on Kolmogorov theory; at least when applied to the full depth of the troposphere.

Realizations using Gaussian random fields cannot capture all observable effects as the non-Gaussian tails of probability distributions of measured data shows. That particular flaw may be possible to overcome by using log-normal random fields computed from Gaussian random fields.

In this thesis, efficient methods and software for the generation of 4D Gaussian and log-normal random fields were implemented. They are largely agnostic to the turbulence model and are powerful enough to handle all structure functions proposed in existing literature known to the author. The method works, but is somewhat inconvenient due to spurious periodicity and integration through noisy fields.

Finally, two novel models of time evolution of refractivity fields complementing the frozen field hypothesis were proposed and implemented. The initial motivation was as a way to deal with the finite geometry, but with a specific modelling choice it was found to be more compatible with Kolmogorov scaling.

Bibliography

[1] M. C. Roggermann et al. “Improving the resolution of ground-based telescopes”. In: *Reviews of Modern Physics* 69.2 (1997), pp. 437–505. DOI: 10.1103/RevModPhys.69.437.

[2] O. J. Sovers et al. “Astrometry and geodesy with radio interferometry: experiments, models, results”. In: *Reviews of Modern Physics* 70.4 (1998), pp. 1393–1454. DOI: 10.1103/RevModPhys.70.1393.

[3] A. L. Fey et al. “The Second Realization of the International Celestial Reference Frame by Very Long Baseline Interferometry”. In: *The Astronomical Journal* 150.2 (2015). DOI: 10.1088/0004-6256/150/2/58.

[4] Z. Altamimi et al. “ITRF2014: A new release of the International Terrestrial Reference Frame modeling nonlinear station motions”. In: *Journal of Geophysical Resarch: Solid Earth* 121.8 (2016), pp. 6109–6131. DOI: 10.1002/2016JB013098.

[5] A. Ricarte and J. Dexter. “The Event Horizon Telescope: exploring strong gravity and accretion physics”. In: *Monthly Notices of the Royal Astronomical Society* 446.2 (2014), pp. 1973–1987. DOI: 10.1093/mnras/stu2128.

[6] M. M. Alizadeh et al. “Ionospheric Effects on Microwave Signals”. In: *Atmospheric Effects in Space Geodesy*. Ed. by J. Böhm and H. Schuh. Springer, 2013, pp. 35–71. DOI: 10.1007/978-3-642-36932-2_2.

[7] T. Nilsson and R. Haas. “Impact of atmospheric turbulence on geodetic very long baseilne interferometry”. In: *Journal of Geophysical Research* 115.B03407 (2010). DOI: 10.1029/2009JB006579.

[8] T. Nilsson et al. “Simulations of atmospheric path delays using turbulence models”. In: *Proceedings of the 18th European VLBI for Geodesy and Astrometry Working Meeting*. Ed. by J. Böhm et al. Vienna, Austria, 2007-04, pp. 175–180.

[9] M. Vennebusch and S. Schön. “Generation of Slant Tropospheric Delay Time Series Based on Turbulence Theory”. In: *Geodesy for Planet Earth*. Ed. by S. Kenyon et al. Vol. 136. International Association of Geodesy Symposia. Springer, Berlin, Heidelberg, 2012, pp. 801–807. DOI: 10.1007/978-3-642-20338-1_101.

[10] T. Hobiger et al. “Fast and accurate ray-tracing algorithms for real-time space geodetic applications using numerical weather models”. In: *Journal of Geophysical Research* 113.D20302 (2008). DOI: 10.1029/2008JD010503.

[11] J. V. Iribarne and H.-R. Cho. *Atmospheric Physics*. D. Reidel Publishing Company, 1980. ISBN: 978-94-009-8954-2. DOI: 10.1007/978-94-009-8952-8.

[12] G. Elgered. "Tropospheric radio-path delay from ground-based microwave radiometry". In: *Atmospheric Remote Sensing by Microwave Radiometry*. Ed. by M. A. Janssen. Wiley, 1993. Chap. 5, pp. 215–258.

[13] T. Nilsson. "Measuring and modelling variations in the distribution of atmospheric water vapour using GPS". PhD thesis. Göteborg, Sweden: Chalmers University of Technology, 2008.

[14] K. Lagler et al. "GPT2: Empirical slant delay model for radio space geodetic techniques". In: *Geophysical Research Letters* 40.6 (2013), pp. 1069–1073. DOI: 10.1002/grl.50288.

[15] T. Ning et al. "Multi-technique comparisons of 10 years of wet delay estimates on the west coast of Sweden". In: *Journal of Geodesy* 86.7 (2012), pp. 565–575. DOI: 10.1007/s00190-011-0527-2.

[16] National Oceanic and Atmospheric Administration (NOAA). *Integrated Global Radiosonde Archive*. URL: www.ncdc.noaa.gov/data-access/weather-balloon/integrated-global-radiosonde-archive (visited on 2018-05-13).

[17] C. Nordling and J. Österman. *Physics Handbook. for Science and Engineering*. 8:10. Lund, Sweden: Studentlitteratur, 2006. ISBN: 978-91-44-04453-8.

[18] L. Davidson. *An Introduction to Turbulence Models*. Department of Thermo and Fluid Dynamics Chalmers University of Technology, Gothenburg, 1997. URL: www.tfd.chalmers.se/~lada (visited on 2018-06-05).

[19] A. N. Kolmogorov. "A refinement of previous hypotheses concerning the local structure of turbulence in a viscous incompressible fluid at high Reynolds number". In: *Journal of Geophysical Research* 13.1 (1962), pp. 82–85. DOI: 10.1017/S0022112062000518.

[20] A. M. Obukhov. "Some Specific Features of Atmospheric Turbulence". In: *Journal of Geophysical Research* 67.8 (1962), pp. 3011–3014. DOI: 10.1029/JZ067i008p03011.

[21] A. Ishimaru. "Temporal frequency spectra of multifrequency waves in turbulent atmosphere". In: *IEEE Transactions on Antennas and Propagation* 20.1 (1972), pp. 10–19. DOI: 10.1109/TAP.1972.1140133.

[22] R. N. Treuhaft and G. E. Lanyi. "The effect of the dynamic wet troposphere on radio interferometric measurements". In: *Radio Science* 22.2 (1987), pp. 251–265. DOI: 10.1029/RS022i002p00251.

[23] H. B. Janes et al. "Comparison of simultaneous line-of-sight signals at 9.6 and 34.5 GHz". In: *IEEE Transactions on Antennas and Propagation* AP-80.4 (1970), pp. 447–451. DOI: 10.1109/TAP.1970.1139715.

[24] A. S. Gurvich and V. I. Tatarskii. "Coherence and intensity fluctuations of light in the turbulent atmosphere". In: *Radio Science* 10.1 (1975), pp. 3–14. DOI: 10.1029/RS010i001p00003.

[25] S. Schön and F. K. Brunner. "A proposal for modelling physical correlations of GPS phase observations". In: *Journal of Geodesy* 82.10 (2008), pp. 601–612. DOI: 10.1007/s00190-008-0211-3.

[26] M. Vennebusch et al. "Temporal and spatial stochastic behaviour of high-frequency slant tropospheric delays from simulations and real GPS data". In: *Advances in Space Research* 47.10 (2011), pp. 1681–1690. DOI: 10.1016/j.asr.2010.09.008.

[27] T. Nilsson et al. “Correlations Between Slant Wet Delays Measured by Microwave Radiometry”. In: *IEEE Transactions on Geoscience and Remote Sensing* 43.5 (2008), pp. 1028–1035. DOI: 10.1109/TGRS.2004.840659.

[28] R. H. Kraichnan. “Diffusion by a Random Velocity Field”. In: *The Physics of Fluids* 13.1 (1970), pp. 22–31. DOI: 10.1063/1.1692799.

[29] A. T. A. Wood and G. Chan. “Simulation of Stationary Gaussian Processes in $[0, 1]^d$”. In: *Journal of Computational and Graphical Statistics* 3.4 (1994), pp. 409–432. DOI: 10.2307/1390903.

[30] C. R. Dietrich and G. N. Newsam. “Fast and Exact Simulation of Stationary Gaussian Processes through Circulant Embedding of the Covariance Matrix”. In: *SIAM Journal on Scientific Computing* 18.4 (1997), pp. 1088–1107. DOI: 10.1137/S1064827592240555.

[31] A. Mantoglou and J. L. Wilson. “The Turning Bands Method for Simulation of Random Fields Using Line Generation by a SpectralMethod”. In: *Water Resources Research* 18.5 (1982), pp. 1379–1394. DOI: 10.1029/WR018i005p01379.

[32] M. L. Stein. *Interpolation of spatial data. Some Theory for Kriging*. Springer Series in Statistics. Springer, 1999. ISBN: 978-1-4612-7166-6. DOI: 10.1007/978-1-4612-1494-6.

Appendix A

General height dependent algorithm

In the following the grid size is taken to be $N_x \times N_y \times N_z$. The simulation can be divided into a setup part and a realization part. The setup consists of the following:

1. For each pair (z_m, z_n), $1 \leq m \leq n \leq N_z$:
   (a) Compute the $N_x \times N_y$ array $C(z_1, z_2, \vec{r}_\parallel^{(i,j)})$ and its 2D DFT with respect to $\vec{r}_\parallel$ denoted $\phi(z_1, z_2, \vec{k}_\parallel^{(i,j)})$.
   (b) To reduce the amount of data from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$, approximate the array $\phi(z_1, z_2, \vec{k}_\parallel)$ with

$$\bar{\phi}(z_1, z_2, k_\parallel^i) = \underset{\left|\vec{k}_\parallel^{(j,k)}\right| \in B_i}{\text{mean}} \phi(z_1, z_2, \vec{k}_\parallel^{(i,j)}). \tag{A.1}$$

   where each *bin* B_i is an interval containing $k_\parallel^i$. Typically $k_\parallel^i = i\Delta k_\parallel$ and $B_i = [k_\parallel^i - \frac{\Delta k}{2}, k_\parallel^i + \frac{\Delta k}{2})$, but many other partitions are possible. The total number of bins N_k is chosen $\propto \sqrt{N_x N_y}$.
2. For each $k_\parallel^i$, $1 \leq i \leq N_k$: Compute the Cholesky decomposition $D_{kl}^{(i)}$ of $\bar{\phi}(z_m, z_n, k_\parallel^i)$ interpreted as a matrix in the indices m and n. I.e. $\sum_k D_{mk}^{(i)} D_{nk}^{(i)} = \bar{\phi}(z_m, z_n, k_\parallel^i)$. The cost of this step is $\mathcal{O}(N_k N_z^3) = \mathcal{O}(N^4)$.

The realization consists of the steps:

1. For each $\vec{k}_\parallel^{(i,j)}$, $1 \leq i \leq N_x$, $1 \leq j \leq N_y$: Find which bin B_l that $\left|\vec{k}_\parallel^{(i,j)}\right|$ is in. Then generate a realization $n_{\vec{k}_\parallel^{(i,j)}}(z)$ of a gaussian stochastic process with zero mean and covariance $\left\langle n_{\vec{k}_\parallel^{(i,j)}}(z_1) n_{\vec{k}_\parallel^{(i,j)}}(z_2) \right\rangle = \phi(z_1, z_2, k_\parallel^l)$ using the stored Cholesky decomposition $D^{(l)}$. The cost is $\mathcal{O}(N_x N_y N_z^2) = \mathcal{O}(N^4)$ in addition to $N_x N_y N_z$ calls to the PRNG.
2. For each z_i, $1 \leq i \leq N_z$: Compute the 2D DFT of $n_{\vec{k}_\parallel^{(j,k)}}(z)$ to obtain a horizontal slice of the refractivity field $n(x_j, y_k, z_i)$. The cost is $\mathcal{O}(N_z N_x N_y \log(N_x N_y)) \sim \mathcal{O}(N^3 \log N)$.

The total amount of memory required for both steps is roughly $N_k N_z^2 + 4 \times N_x N_y N_z$ numbers (single or double precision).

Preprocessing is in practice much faster for covariance functions of the form

$$\langle n(\vec{r}_1) n(\vec{r}_2) \rangle = C_\mathrm{v}(z_1, z_2) C_\mathrm{i}(\vec{r}_1 - \vec{r}_2) \tag{A.2}$$

Table A.1: The times below are from execution on my desktop. Windows 7 64-bit running on a 3 GHz Intel Core i3 CPU with 4 GB of RAM available.

Time of computation in seconds		**Size** (N)			
Step	Description	64	128	200	256
Gen.1	Compute covariance matrices	0.14	2.40	18.7	68.4
Gen.2	Cholesky factorization	0.01	0.17	0.86	2.25
Real.1	Generate random numbers	0.12	1.69	7.73	26.6
Real.2	FFT to real space	0.00	0.04	0.21	0.64
-	Write 3D field to disk	0.07	0.33	2.00	4.50

or sums thereof. The complexity is the same, but this form allows C_v and C_i to be precomputed with $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$ operations respectively instead of $\mathcal{O}(N^4)$ operations when placed within the big loop. Since the C:s often contain multiple expensive operations like `exp`, `pow` and `sqrt` this can increase the speed by a large factor. The N^2 2D FFTs must still be evaluated which means $\mathcal{O}(N^4 \log N)$ operations are still needed overall, but they are relatively cheap.

Some example runtimes of the different steps of the C++ implementation is given in Table A.1. In all cases $N_x = N_y = N_z = N$.

Appendix B

Data sources

The Onsala Space Observatory regularly participates in VLBI experiments. It also has a variety of other astronomical, geodetic and meteorological experiments. For instance a ground based weather station, GNSS receivers and two water vapour radiometers Astrid and Konrad.

The radiometer based estimates of zenith wet delay used in this thesis come from Konrad and were computed by Peter Forkman. GPS-based estimates come from the ONSA and RECF recievers and were computed by Grzegorz Klopotek.

The weather and radiometer data is publicly accessible from `wx.oso.chalmers.se/data/` in subfolders `weather/` and `radiometer/` respectively. The format of the weather data is documented in a public text file, but the radiometer data format is not. Some other datasets are also available, such as processed GPS-data.

For weather forecasting and climate studies, radiosondes (weather balloons) are released simultaneously two times per day (00UTC and 12UTC) on various locations around the globe. The closest one to OSO is at Landvetter. The radiosondes are carried by weather balloons up to the lower stratosphere where the balloons break and the radiosondes deploy a parachute and fall gracefully back to the ground. The data is transmitted by radio to the launching station. Data from radiosonde launches back to 1905 are publicly available through the Integrated Global Radiosonde Archive [16]. The data format is coherent over the entire timespan.

There are a large number of fixed GNSS receiver stations around the globe used for various purposes. For some of them, arrival time data is publicly available. There are several archives of such data. See for instance `http://sopac.ucsd.edu` and `http://www.sonel.org`. It is possible, but far from trivial, to compute wet delays from this data.

There are open databases with DNS simulations of turbulence on the internet. One example is the Johns Hopkins Turbulence Database `http://turbulence.pha.jhu.edu`. This kind of data can be compared with statistical models of turbulence.