UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

DEVELOPMENT AND VALIDATION OF A SODAR SIMULATOR FOR ATMOSPHERIC BOUNDARY LAYER CHARACTERIZATION

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

CHARLOTTE ELIZABETH WAINWRIGHT
Norman, Oklahoma
2014

DEVELOPMENT AND VALIDATION OF A SODAR SIMULATOR FOR ATMOSPHERIC BOUNDARY LAYER CHARACTERIZATION

A DISSERTATION APPROVED FOR THE
SCHOOL OF METEOROLOGY

BY

Dr. Phillip Chilson, Co-Chair

Dr. Robert Palmer, Co-Chair

Dr. Evgeni Fedorovich

Dr. Petra Klein

Dr. Jeffrey Kelly

Dedication

To my parents, Pat and Geof Wainwright, without whom none of this would have been possible. This is for you.

Acknowledgements

The National Science Foundation is acknowledged for the support of the work presented in this dissertation through the grant ATM-1016153.

I would like to thank all of my committee members for selflessly giving of their time and energy to help me complete my dissertation. Special thanks go to my committee co-chairs, Drs. Phil Chilson and Bob Palmer. They took a chance by bringing me onto this project, for which I shall always be grateful. I have learnt so much from both of them, not just on the intricacies of sodar signal processing, but also about how to work to the best of my abilities. The guidance they have provided will stay with me as I move forward through life. I would also like to thank Dr. Evgeni Fedorovich for his tireless efforts to help me understand the details of the structure function parameter. His comments and suggestions have greatly improved my work.

I am deeply grateful to the School of Meteorology and Advanced Radar Research Center staff for all their help over the years. Thank you, Krysta, Jo Ann, Lauren, Celia, Marcia, Nancy, and Becky, for always providing a friendly face in the office.

My parents have always encouraged me to follow my dreams, even when those dreams have taken me far from home. I will be forever grateful for their tireless support in cheering me on, from primary school to graduate school. They have enabled me to embrace every opportunity, and I owe much of my success to the sacrifices they have made for me.

My friends have made graduate school a pleasant place to be, and kept me sane throughout. I deeply appreciate their support, and will always fondly remember the many fun times that we've had. Special thanks to my officemates Tim and Jeremy, who have provided, at various times and in copious amounts, LES and UAS data, coffee, donuts, and moral support. Special thanks also to Phil Stepanian, who has brought me such happiness over the last three years, and without whose love and support I would certainly not have been able to complete this work.

Contents

List Of Tables

List Of Figures

Abstract

A sodar simulator capable of producing time-series data emulating sodar signals has been developed and tested. The atmospheric fields used to populate the sodar simulator are taken from output of a large eddy simulation code. The characteristics of the sodar (number of beams, azimuth and zenith angles of the beams, beamwidth, transmit frequency, range resolution, etc.) are defined by the user to allow evaluation of and comparison with existing systems. The range of the reflected acoustic signal is calculated based on a temperature-dependent speed of sound. Realistic acoustic background noise is simulated using filtered white noise. The raw acoustic time-series data are processed using a Fourier transform to yield acoustic Doppler spectra, from which the radial velocities are calculated. The design of the simulator allows for the testing of and comparison between various signal-processing techniques and averaging periods.

Several different methods for validating the results of the time-series sodar simulator are presented, including validation of the derived wind speeds against a simulated instrumented tower and a moment sodar simulator, which evaluates the wind components using the radial velocities measured by the sodar beams, and as such includes the effects of the DBS technique. The results of the sodar simulator are also used to evaluate vertical profiles of the structure function parameter for temperature, and these profiles are compared against those derived from a simulated unmanned aerial system, and evaluated from the large eddy simulation output directly.

Two example cases are presented of populating the sodar simulator with large eddy simulation data representative of one developing and one fully developed convective boundary layer, and a third example case using large-eddy simulation output of a stable boundary layer.

Chapter 1

Introduction

1.1 Motivation

Sodars have been widely used to study the atmospheric boundary layer (ABL) since their development in the 1970s (Kallistratova and Coulter 2004). Sodars are used in the field of air pollution monitoring and forecasting (Pekour and Kallistratova 1993; Singal et al. 1994; Walczewski 1997) as they can provide profiles of the low-level three dimensional wind-field, and estimates of the mixing height, which is an important parameter in air pollution modeling.

Sodars are able to capture in detail boundary layer phenomena, such as the nocturnal low-level jet (e.g., Beyrich 1994; Banta et al. 2002; Kallistratova et al. 2009; Kallistratova et al. 2013), convective structures in the daytime ABL (e.g., Eymard and Weill 1982; Taconet and Weill 1983; Seibert and Langer 1996; Petenko and Bezverkhnii 1999), and elevated inversions (e.g., Weill et al. 1978; Keder 1999). Sodars are also well suited to provide information about the structure of the ABL turbulence for the turbulence monitoring community (e.g., Greenhut and Mastrantonio 1989).

As the wind energy industry continues to grow, there has been a renewed interest in acoustic remote sensing. There are now several sodar models marketed specifically for this application. Bradley et al. (2005) reported on the calibration of sodar systems for wind energy applications and several recommendations were made for improving the use of sodars for such applications, including the development of a means of self-calibration, improving the signal to noise ratio by transmitting in several beams directions simultaneously, and the development of more sophisticated data filtering

and digital signal processing techniques. As such, there is a clear need to develop improved scanning strategies and signal processing techniques to allow the maximum extraction of useful information from sodar returns.

Despite their extensive employment in ABL studies, sodars do have some limitations. For instance, they typically cannot be used in highly populated areas due to the noise pollution they cause. The performance of sodars in complex terrain has been discussed (e.g., Neff 1988; Soler et al. 2003; Bradley 2008b) in connection with the level of surface heterogeneity that could affect sodar measurements. Since sodars derive the three-dimensional wind components using the DBS technique Balsley and Gage (1982), it is assumed that the wind at each of the locations sampled by the distinct beams is identical. Bradley (2008b) illustrated that the invalidation of this assumption caused by complex terrain creates errors in the derived wind speeds, via simulation of flow over a hill and a ridge. Errors in the wind speed resulting from the use of the DBS technique on heterogeneous flow fields were calculated by Bradley (2008b) to be in the range of 5 - 20%. The effect of applying the Doppler Beam Swinging (DBS) technique across an area with horizontal wind shear or nonuniform vertical velocities was investigated for wind profiling radars by Koscielny et al. (1984), who found that the horizontal wind shear or nonuniform vertical velocity introduced a bias in the wind estimates, and that this bias could not be reduced by averaging more measurements. Crescenti (1996) summarized many comparison studies between winds derived from sodar measurements and those derived from cup/vane anemometers, sonic anemometers, and lidar. The mean correlation between sodar-derived winds and those from the comparison suite of instruments across twenty different comparison studies was found to be 0.92 (Crescenti 1996), with better agreement between sodar and other measurement platforms when the boundary layer wind field

was horizontally homogeneous. Moreover, Crescenti (1996) also claimed that an inhomogeneous boundary layer can complicate comparisons and make interpretation of results extremely difficult.

Since commercial sodar generally make use of the DBS technique to derive the three-dimensional wind field from three or more independent beams, these derived winds will also be biased if the wind field is not horizontal homogeneous. Investigations into the minimum averaging time required for comparison of horizontal wind components under convective conditions were made by Kaimal et al. (1980), who suggested a minimum averaging time of 20 minutes, with the stipulation that this applies for horizontally homogeneous flows. An appropriate averaging period for comparison purposes in the stable boundary layer was similarly investigated by Gaynor et al. (1991), who concluded that 20-minute averaging periods were appropriate for comparisons in the stable boundary layer. Due to the use of the DBS technique on heterogeneous wind fields, there exist questions as to the range of spatial and temporal scales over which sodar observations of the wind can be considered representative. This cannot be investigated by closely collocating several sodars in the field, as close spacing of several sodars would result in acoustic interference among the sodars.

The questions that exist as to the accuracy of sodar observations of wind in the horizontally inhomogeneous boundary layer provide partial motivation for the research presented in this dissertation. Additional motivation lies in the stated need for improvements in sodar digital signal processing techniques for the wind energy industry. The goals of the presented research are outlined in the following section.

1.2 Goals

The goal of the presented research is the development of a sodar simulator capable of ingesting large eddy simulation (LES) output fields and producing data streams which

mimic output that would be produced by a sodar operating under the atmospheric conditions simulated by the LES.

It is hoped that the development of a sodar simulator will allow for an examination into the errors in sodar-derived winds resulting from the use of the DBS technique on horizontally inhomogeneous data. The use of the sodar simulator on simulated atmospheric fields from the output of a large eddy simulation (LES) allows the LES to act as a 'truth' against which to compare results of the sodar simulator.

The sodar simulator will also allow for the repeated testing of the same dataset with different simulated sodar scanning strategies and signal processing techniques, which is not possible with observed sodar data. This should facilitate a fair comparison between the different scanning strategies and signal processing techniques.

By formulating the sodar simulator to follow the process used by an operational sodar as closely as possible, it is hoped that tests of the simulator performed on LES output will produce realistic simulated sodar output. Testing the sodar simulator on LES output representing different atmospheric conditions should allow for recommendations to be made on the optimal sodar operating parameters for use under various atmospheric conditions.

1.3 Structure of the dissertation

The dissertation is comprised of seven chapters. Chapter 2 provides a review of boundary layer characterization techniques, introducing the parameters and variables that are used to characterize the boundary layer and a description of the *in-situ* and remote sensing instrumentation available to observe these variables. The second chapter also contains a section detailing the use of sodar as a remote sensing tool for the study of boundary layer structure, and a brief history of some of the specific boundary layer topics which have been addressed through the use of sodar.

The theory behind acoustic remote sensing is presented in Chapter 3. Details are provided regarding how the choice of sodar design and the selection of operating parameters such as frequency, transmitted signal power, and sampling time will affect the sodar output. Aspects of the signal processing techniques employed by commercially available sodars are also discussed, with a focus on the effects of the signal processing technique upon sodar output parameters.

Chapter 4 prescribes the specific setup of the sodar simulator which was developed as the basis of this study. An overview of the algorithmic structure of the sodar is provided, and details are provided on how each algorithm is implemented within the sodar simulator. A brief overview of the large eddy simulation code used in the later case studies is given at the beginning of this chapter.

The different validation techniques used to examine the results from the sodar simulator are discussed in Chapter 5. Several methods for calculating the structure function parameter for temperature are presented, including calculation from observed and simulated sodar output, calculation using data from an operational unmanned aerial system, and calculation from both a simulated unmanned aerial system and directly using the LES output. The format of the unmanned aerial system simulator, which was also designed as part of the presented research, is also presented in detail within Chapter 5.

Chapter 6 presents three case studies for which the sodar simulator is driven with output from a large eddy simulation. Two convective boundary layer case studies are presented, using simulations representing 31 May 2009 and 24 May 2013. The 24 May 2013 simulation represented a fully developed convective boundary layer and was nudged with WRF model data to more closely represent the atmospheric conditions measured at a site in Purcell, Oklahoma, during a field experiment. The field experiment provided data from an operational sodar and an unmanned aerial system, against which the results of the simulated sodar can be compared. A third

case study, in which the large eddy simulation represents a stable boundary layer is also described in this chapter and results from the case study are presented and discussed.

Finally, Chapter 7 offers further discussion of the results presented in the preceding chapter, and conclusions following from the work are presented. Suggestions for future work to enhance the sodar simulator are also given in Chapter 7.

Chapter 2

Review of boundary-layer characterization techniques

2.1 Atmospheric boundary layer overview

The Atmospheric Boundary Layer (ABL) is the lowest portion of the atmosphere, and it can be defined as the part of the troposphere that is directly influenced by the presence of the Earth's surface and responds to surface forcings on a timescale of an hour or less (Stull 1988). These surface forcings include effects from evapotranspiration, heat transfer, frictional drag, emission of pollutants, and flow modification caused by variations in the terrain. In the following sections, some aspects of studying the boundary layer will be presented, with a focus on those parameters which are important throughout the remainder of the dissertation.

2.1.1 Boundary-layer depth

The depth of the boundary layer over land is highly variable in time and space, and is dependent upon a number of factors. These factors include surface characteristics, atmospheric stability, time of day, incoming solar radiation, and atmospheric pressure. The boundary layer over land has a well defined-structure that evolves with the diurnal cycle (Stull 1988).

The diurnal cycle of the boundary layer is shown in Fig. 2.1. Three components of the boundary layer, the mixing layer, the residual layer, and the stable boundary layer can be seen in Fig. 2.1. The surface layer is located at the bottom of the ABL, and is defined as the region where turbulent fluxes vary by less than 10% of their magnitude. As illustrated in Fig. 2.1, the surface layer can be a part of either the mixing layer (typically during the day) or the stable boundary layer (typically during the night).

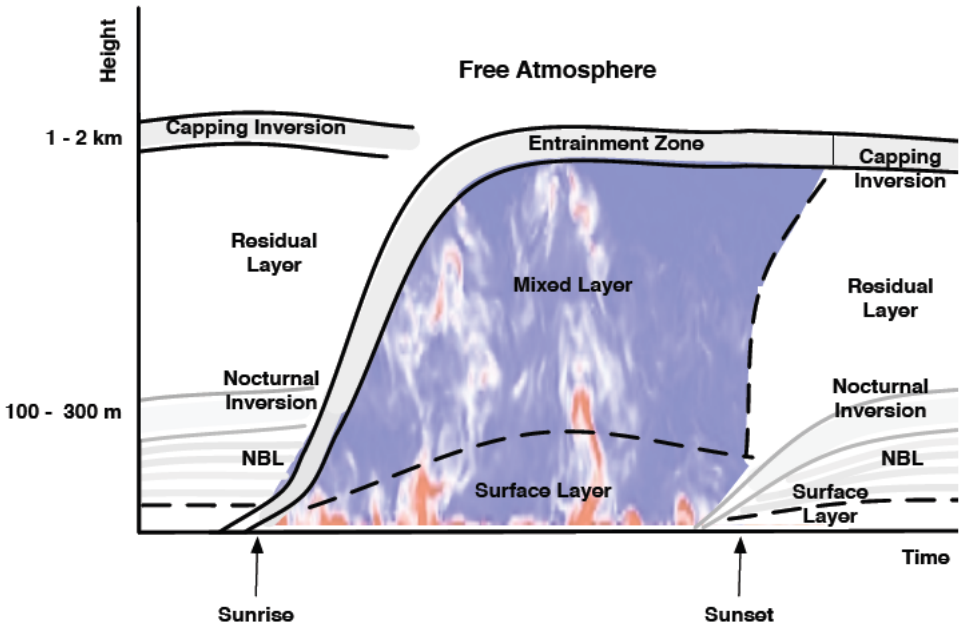

Figure 2.1: Depiction of the diurnal variation of the boundary layer (Scipión 2011, adapted from Stull 1988).

The diurnal cycle shown in Fig. 2.1 can be summarized as follows: after sunrise, the surface is heated by solar radiation. As solar heating increases, turbulent eddies in the mixing layer begin to grow in size, increasing the height of the mixing layer. The mixing layer also grows by entrainment of less-turbulent air from the free atmosphere into the mixing layer. This process occurs in the entrainment zone, which can be seen capping the mixing layer in Fig. 2.1. The mixing layer reaches its maximum depth in the late afternoon. Shortly before sunset, as the solar incidence angle decreases, thermals (illustrated in red in Fig. 2.1) stop forming and the turbulence in the mixing layer begins to decay. The resulting layer is called the residual layer. Since the turbulence decays following sunset, particles and pollutants which were lofted above the surface layer in the mixing layer will remain lofted in the nocturnal boundary layer. During the night, the lowest portion is transformed into a stable boundary layer due to contact with the ground, which is colder after the removal of solar heating. The stable nocturnal boundary layer is characterized by weaker and sporadic turbulence (Stull 1988). Static stability in the nocturnal boundary layer can be enhanced by strong radiational cooling on clear nights. Turbulence in the nocturnal boundary layer can be generated by the wind shear caused by the development of a region of accelerated wind speeds aloft known as the low-level jet. The development of the low-level jet is a fairly common nocturnal occurrence in many location, particularly in the Southern Great Plains region of the US. A climatology of the low-level jet over the Southern Great Plains region can be found in Song et al. (2005).

2.1.2 Wind in the boundary layer

Stull (1988) divides wind in the boundary layer into three categories: the mean wind, turbulence, and waves. All three categories can co-exist or exist separately in the boundary layer. The mean wind is responsible for the horizontal transport of pollutants, moisture, heat, and momentum in the boundary layer, and the vertical

transport of these quantities is dominated by turbulence. Mean horizontal wind speeds between 2 m s^{-1} and 10 m s^{-1} are typical within the boundary layer, while the vertical wind speeds are much smaller, generally below 1 m s^{-1} outside of updrafts and downdrafts.

The mean wind can be split into contributions from large-scale variations and turbulence by averaging the wind speed over a period of 30 minutes to one hour, to give mean values $\overline{u}$, $\overline{v}$, and $\overline{w}$. This averaging has the effect of removing fluctuations due to turbulence. The instantaneous velocity perturbations, u', v' and w' can then be evaluated by subtracting the mean velocity value from the instantaneous measurements,

$$u' = u - \overline{u}, \tag{2.1}$$

$$v' = v - \overline{v}, \tag{2.2}$$

and

$$w' = w - \overline{w}. \tag{2.3}$$

The mean flow ($\overline{u}$, $\overline{v}$, $\overline{w}$) can be regarded as the part of the wind flow that varies at timescales longer than the averaging period used in their evaluation. The perturbation velocity components, u', v' and w' represent velocity variations at timescales shorter than the averaging period used to evaluate the mean flow components (Stull 1988).

2.1.3 Turbulence

Turbulence can be thought of as a random, three-dimensional state of motion which is characterized by a high degree of chaotic velocity. The structure of turbulence is made up of a collection of eddies of variable sizes which interact both with each other and with the mean flow. Two major types of turbulence exist in the ABL: mechanical turbulence caused by the instability of the vertical wind shear, and thermal turbulence caused by buoyancy forces. A thorough review of atmospheric turbulence can be found in Fedorovich et al. (2004c).

2.1.3.1 Turbulence kinetic energy

Kinetic energy associated with flow of the wind can be divided into contributions from the mean wind and contributions from turbulence. The kinetic energy resulting from the mean flow can be calculated as

$$KE = \frac{1}{2}(u^2 + v^2 + w^2). \tag{2.4}$$

The turbulence kinetic energy resulting from the mechanical production of turbulence can be written as

$$TKE = \frac{1}{2}(\sigma_u^2 + \sigma_v^2 + \sigma_w^2) \tag{2.5}$$

where σ_u^2, σ_v^2, and σ_w^2 are the variances of the velocity in the zonal, meridional, and vertical directions, respectively. Turbulence kinetic energy can be considered as a measure of the intensity of turbulence, and it is related to the transport of momentum, heat, and moisture through the atmosphere. As such, it is considered an important variable in the field of micrometeorology (Stull 1988).

2.1.3.2 Structure function

One way to examine the strength of turbulence is to look at the spatial variation of quantities which are assumed to be advocated by turbulent flows in the surface layer. One such example is temperature. The structure function parameter for temperature is a way to quantify the effects of small-scale turbulence on the variability of atmospheric temperature. The structure function parameter is explained as such: for any scalar, k, whose value is known at at least two locations separated in space, the structure function parameter is defined as

$$C_k^2 = \overline{(\delta k)^2} r^{2/3} \tag{2.6}$$

where $\overline{(\delta k)^2}$ is the structure function for k (explained below), and r is the separation distance between the points at which the variable is evaluated (Tatarskii 1961).

The structure function, $\overline{(\delta k)^2}$, is most simply described as a quantification of the difference between the value of a variable at two known locations. It is evaluated as

$$\overline{(\delta k)^2}(\mathbf{r}, t) = \overline{[\mathbf{k}(\mathbf{x}, t) - k(\mathbf{x} + \mathbf{r}, t)]^{\mathbf{2}}} \tag{2.7}$$

where the overbar represents ensemble averaging, $\mathbf{x}$ is the position vector, $\mathbf{r}$ is the separation vector between the two points, and t is time.

The most commonly used structure functions in boundary layer layer studies are those for temperature (C_T^2), for the magnitude of the horizontal velocity (C_V^2), for moisture (C_q^2), and for the refractive index (C_n^2). The structure function formulation is applicable only when the turbulent fluctuations responsible for the spatial variation of the quantity being measured are assumed to be locally isotropic and homogeneous, and within the inertial subrange of turbulent scales. The inertial subrange can be thought of as the region of the turbulence spectrum where energy is neither being created nor lost, i.e., the region in which shear and buoyant production of turbulence is balanced by the dissipation of turbulent eddies at small sizes (Kolmogorov 1941).

2.2 Instrumentation for boundary-layer studies

Instrumentation used for boundary layer studies can be split into two main categories: *in-situ* instruments and remote sensing instruments. During an *in-situ* measurement, the sensor is in direct contact with the object whose property is being measured, e.g., a thermometer in direct contact with the surrounding air. For measurements made with remote sensors, on the other hand, the remote sensing instrument does not come into direct contact with the object at the location the measurement represents, but rather the remote sensor analyzes radiation which is emitted or scattered by the measurement object. Generally most remote sensors involve theory in order to estimate the desired variables, rather than measuring directly observable quantities.

A comprehensive review and comparison of each of the instruments described in Sections 2.2.1 and 2.2.2 can be found in Emeis (2010), which details the development

and applications of all the instruments in the following section, as well as several others.

2.2.1 *In-situ* instrumentation

There are several different kinds of sensors used for making *in-situ* measurements within the boundary layer. Due to the nature of *in-situ* measurements, sensors of this kind are either ground based, or they must be attached to a structure to raise the instrument to height at which measurements are desired. This obviously creates limitations as to which variables can be measured, and the locations in which the measurements can be made.

A secondary caveat to *in-situ* instrumentation is that generally, *in-situ* sensors provide only point measurements, i.e., the measurement can be considered representative only of the exact point location at which it was made. Whether or not this is of importance obviously depends upon the variable being measured, the spatial variation of the variable, and for what purpose the measurements are being made. In order to provide coverage over a finite spatial area, an array of *in-situ* instrumentation must be used.

2.2.1.1 Ground-based instrumentation

A large portion of *in-situ* meteorological measurements fall into the category of ground-based measurements. The measurement of temperature, pressure, humidity, wind speed, and precipitation at the surface is highly important, as it is the surface values of these properties which have the largest effect on the population, as well as effects on agriculture.

The ground-based data that are used for this research are provided by the Oklahoma Mesonet (Brock et al. 1995). The Oklahoma Mesonet is a distributed network

of meteorological measurement stations throughout the state of Oklahoma. Each station measures the air temperature, humidity, barometric pressure, wind speed and direction, precipitation, solar radiation and soil temperature, and updates a public website with this information at 5-minute intervals. A short description of the instruments used to measure some of these variables is given below, focusing on those variables which are used (either from observations, or via simulation) within this work.

Each Mesonet station measures air temperature at 1.5 m and 9 m above the ground. The temperature is measured using interchangeable chip thermistors with an accuracy of $\pm 0.4^{\circ}$C (McPherson et al. 2007). The thermistor at 1.5 m above ground level is housed in an aspirated radiation shield, and the sensor at 9 m is housed in an unaspirated radiation shield, which can bias readings in calm conditions with high solar radiation. The humidity is measured using a capacitive relative humidity sensor, which Mesonet laboratory tests designated to have an accuracy of 3% at relative humidity values between 10% and 98%. Precipitation is measured at a height of 0.6 m above the ground using an unheated tipping-bucket rain gauge. The rain gauge has a 30-cm opening and is surrounded by a shield to decrease the effects of wind on the falling precipitation. The accuracy of the rain gauge when tested in the Mesonet laboratory was $\pm 5\%$ over a range of 0 - 5 cm hr^{-1} (McPherson et al. 2007). Since the rain gauges are unheated, measurements of the water equivalent from frozen precipitation are delayed until the air temperature rises above freezing. The wind speed and directed at the Mesonet sites are measured at 10 m height using a combination propellor and vane anemometer. The accuracy of the measured wind speed is ± 0.3 m s^{-1} for wind speeds between 1.0 m s^{-1} and 60.0 m s^{-1}. The starting wind speed threshold of the propellor is 1.0 m s^{-1} and that for the vane is 1.1 m s^{-1}, so wind speed below these values cannot be accurately measured. These variables are those which are used within this study. Information about the Mesonet measurements

of pressure, soil moisture, and solar radiation, including instrument accuracies and details on the particular instrument models used at the Mesonet stations can be found in McPherson et al. (2007).

2.2.1.2 Radiosondes

Radiosondes provide a method to extend *in-situ* instrumentation to measure vertical profiles of atmospheric variables. A radiosonde consists of a small instrument package that is suspended below a large balloon. The balloon is inflated with helium and released into the atmosphere. The instrument package records measurements of temperature, pressure, and relative humidity as the balloon ascends to heights of up to 35 km. If the package contains a GPS location marker, then the wind speed and direction can be derived by tracking the location of the balloon relative to ground as it ascends.

The US National Weather Service operates a network 92 of upper-air observation stations, which typically release radiosondes at twelve hour intervals. Guidelines for the accuracy and range of the temperature, pressure, and humidity sensors can be found in Organization (1996). The purpose of the network of stations releasing radiosondes is to provide vertical profiles of temperature, pressure, and humidity, which can be used for forecasting. More information about the history of radiosonde use and the specific sensors used in sevrel commercial radiosonde models can be found in Dabbert et al. (2002)

2.2.1.3 Instrumented aircraft observations

Instrumented aircraft provides a way to extend the field of *in-situ* remote sensing to areas above the surface without requiring instrumented towers. The premise of instrumented aircraft observations is simple: a manned aircraft is outfitted with meteorological instrumentation, and measurements are recorded while the aircraft is in

flight. There are two main benefits of instrumented aircraft, the first being that as well as making measurements at heights greater than could be reached by an instrumented tower, an aircraft can make observations in areas which are impractical to have surface-based instrumentation, such as in remote locations. The second benefit is that aircraft measurements can provide information about the variation of meteorological parameters such as temperature and humidity over a range of spatial scales, while stationary ground-based sensors can provide information only on temporal variations of these parameters.

The use of instrumented aircraft for meteorological research primarily began in the 1950s (e.g., Brewer 1954; Newton 1955) and has been since used to study a wide variety of research topics. Many instrumented aircraft are outfitted with a mixture of both typical *in-situ* instruments such as thermometers and relative humidity sensors, and instrumentation for specialized research purposes. For example, the instrumented aircraft operated through the University of Wyoming, King Air, is equipped with standard surface instruments to measure temperature, wind speed, pressure, and humidity, and is also outfitted with microphysical instrumentation capable of measuring cloud droplet spectra and droplet effective radius. The King Air also contains a suite of instruments for measuring trace gases and aerosol properties. The wide variety of instrumentation available on research aircraft makes them suitable for research on many different topics. For example, the T-28 instrumented research aircraft operated by the South Dakota School of Mines and Technology has been used to study electrical fields in thunderstorms, to measure trace gases in storms to track the storm circulation, and to make microphysical observations in convection (Detwiler et al. 2012). Instrumented aircraft observations have been used in research varying from examination of katabatic winds in the Antarctic (Parish and Bromwich 1989) to flights through hurricanes in the Atlantic to measure the pressure and radius of the maximum wind (Jorgensen 1984).

2.2.1.4 Unmanned aerial systems

The idea of using small remotely controlled aircraft to measure meteorological variable was first suggested by Konrad et al. (1970). Over the last two decades, the technology required for this idea has been developed, and remotely controlled aircraft equipped with meteorological sensors have since been used for meteorological research, among other uses. The remotely controlled aircraft are known as unmanned aerial systems (UAS). The use of UAS for meteorological research has grown in popularity since their development in the early 2000s (e.g., Holland et al. 2001, Shuqing et al. 2004). The UAS platform can be outfitted with many of same sensors as a full-size manned aircraft, but are limited by the size and weight of the UAS platform. UAS measurements of temperature, relative humidity, and wind speed and direction have been used to evaluate boundary layer schemes within numerical models (Mayer et al. 2011). More recently, UAS platforms have been used to examine turbulence, through the estimation of the structure function parameter for turbulence (van den Kroonenberg et al. 2012, Bonin et al. 2014).

The use of UAS to measure atmospheric variables provides many of the same benefits as manned instrumented aircraft flights. Profiles of atmospheric variables such as temperature and humidity can be measured up to heights which could not be reached by an instrumented tower. As with manned instrumented flights, the UAS is also able to measure the spatial variation of atmospheric parameters. At present, UAS flights made for research purposes by scientific institutions are monitored by the Federal Aviation Authority (FAA). UAS flights must adhere to the codes of practice laid out in the FAA certificate of authorization (COA), which provides safety regulations and emergency procedures. A COA must be obtained before UAS flights can be performed.

The UAS flights which provided data used within this research were conducted using the UAS operated at the University of Oklahoma (full details on this UAS

platform are provided in Section 5.2.1) under FAA COA 2012-CSA-57. The COA stipulates that the UAS flights must remain within a 1-mile radius of the stated experiment site, and the height of the flight is limited to 3000 ft. The COA regulations stipulate that the platform cannot be flown during hours of darkness, and it must remain within eye sight of the ground pilot operating the remotely controlled UAS. This limits the use of UAS in complex terrain, in which a line-of-site could be difficult to maintain. Further details on the FAA regulations pertinent to the UAS used within this research can be found in Bonin et al. (2014).

2.2.2 Remote sensing instrumentation

Remote sensing can be further divided into the categories of passive and active remote sensing. Passive remote sensing involves an instrument monitoring naturally occurring energy (e.g., thermal infrared energy). Active remote sensing, on the other hand, involves the instrument emitting a signal, and then receiving scattered energy from that signal.

There are many types of instruments for remote sensing of the boundary layer. The following sections provide details on four of the most widely-used remote sensors for boundary layer studies: wind profiling radars, lidars, scintillometers, and sodars. Details on several remote sensing techniques which are not discussed here can be found in Emeis (2011).

2.2.2.1 Wind profiling radar

Unlike weather radars, wind profile ring radars are capable of detecting atmospheric echoes from optically clear air. Most wind profiling radars and Doppler weather radars transmit a pulse-modulated electromagnetic wave into the atmosphere. However, while Doppler weather radars receive backscatter from Rayleigh scattering of

hydrometeors in the atmosphere, wind profiling radars can measure backscatter received from Bragg scatter produced by fluctuations of the refractive index, n, in clear air. Wind profiling radars are sensitive to Bragg scatter from turbulent eddies at spatial scales of size $\lambda/2$, where λ is the transmit wavelength of the radar. For a 400-MHz radar wind profiler, $\lambda/2$ equates to turbulent eddies of size 0.37 m, and for a 900-MHz wind profiling radar, the Bragg scale equates to approximately 0.16 m. These frequencies are commonly used for wind profiler operation in the US.

In the US, there is a network of operational wind profiling radars, the NOAA Profiling Network (NPN). The majority of the NPN radars are located within the central and southern Great Plains region of the country, with one additional profiler located in upstate New York and three located in Alaska. All of the profiling radars within the contiguous US operate at either 404 or 449 MHz and the three radars in Alaska operate at 449 MHz. NPN radars produce moments for each range gate every at 6-minute intervals and mean hourly profiles of the vertical and horizontal wind speed and direction from 500 m above the ground up to a height of 16.25 km (Barth et al. 1994). This is achieved through the use of two separate modes, which alternate operation at 1-minute intervals. The low mode samples the lower atmosphere from heights of 500 m to 9.25 km above the ground, and the high mode samples heights from 7.5 km to 16.25 km. The low mode provides the wind profile at a height resolution of 320 m, and the high mode produces wind profile measurements at 900 m vertical resolution. The horizontal wind speed is derived from the frequency spectra produced by the off-vertical beams, through use of the Doppler beam swinging technique (DBS; Balsley and Gage 1982), with the vertical wind component derived from the frequency spectra resulting from the vertical beam. The hourly wind profile measurements are publicly available on the NOAA profiler website.

As well as providing wind profile measurements throughout the troposphere, wind profiling radars also output the so-called echo power. The echo power is measure of

the relative backscatter from each range gate, and this can be used to identify cloud layers and the vertical extent of convection.

2.2.2.2 Lidar

Lidars emit light beams and use the measured backscatter to derive concentrations of trace substances in the atmosphere. There are four different kinds of lidar used for studying the boundary layer: backscatter lidars and ceilometers, differential absorption lidars, Raman lidars, and Doppler wind lidars (Emeis 2011). The main purposes of the four types of lidar are discussed in the following paragraphs.

Backscatter lidars show the backscatter resulting from scattering by aerosol particles and trace gases in the atmosphere. As such, the vertical profile of aerosols in the atmosphere can be determined from the returned signal power. The ceilometer is a simple form of a backscattering lidar. It records the optical backscattering intensity in the near infrared, and are typically used to monitor the height of the cloud base. For this reason, ceilometers are widely used at airports, since cloud base monitoring is highly important for aviation.

Differential absorption lidars emit two separate pulses of light at slightly different frequencies. The frequency of the two pulses are chosen such that one of the frequencies is optimally absorbed by the trace gas of interest. Absorption of the lidar pulse by trace gases in the atmosphere is highly frequency-dependent, and thus by comparing the two returned pulse intensities at the optimal absorption frequency of the trace gas of interest, the concentration of the absorbing gas is detectable. Absorption by aerosol particles occurs across a broader frequency band, so both lidar pulses experience the same absorption by aerosols. Differential absorption lidar can be used to examine the concentration of a wide range of trace gases, such as monitoring the concentration of sulphur dioxide following volcanic eruptions (e.g., Edner et al. 1994). However, differential absorption lidar is most commonly used to derive

vertical profiles of water vapor in the atmosphere (e.g., Giez et al. 1999). A thorough theoretical analysis of the use of differential absorption lidar to determine the profile of water vapor in the atmosphere conducted by Wulfmeyer and Bösenberg (1990) concluded that water vapor profiles could be retrieved with error below 5% throughout the troposphere.

Raman lidars are also used to detect the concentration of trace gases in the atmosphere. This lidar detects Raman scattering from trace gas molecules. Raman scattering occurs when a trace gas molecule absorbs a photon, enters an excited state, and then drops back to an energy state close to, but not exactly, its original energy state. The frequency of the scattered photon changes slightly resulting from this procedure, and this frequency shift is measured by the Raman lidar. Each trace gas causes a characteristic frequency shift, so the presence of trace gases, including water vapor, can be identified by examination of the frequency of the returned photon. As such, the Raman lidar can be used to derive the profiles of water vapor and other trace gases in the atmosphere (e.g., Turner et al. 2002).

The final type of lidar, the Doppler wind lidar, is used to study the boundary layer wind profile. Doppler wind lidar can be further divided into two categories: coherent lidar systems and incoherent/direct detection lidar systems. Coherent lidar systems determine radial velocity along the beam by determination of the Doppler shift through comparison of the frequency of backscattered optical radiation to that of a reference beam. Direct detection (incoherent) lidar systems determine the frequency shift by passing the light through an optical filter (Chanin et al. 1989). Both systems produce profiles of the vertical wind speed, and if off-vertical beams are employed, then the horizontal wind speed and direction can also be detected through the use of the DBS technique (Balsley and Gage 1982). Doppler wind lidars typically provide wind profiles from heights of 40 m up to a few kilometers, although they are limited

by the presence of aerosol particles. In very clean air, with few aerosol or trace gas particles, there are no scattering particles and so no signal return is received.

2.2.2.3 Scintillometer

Scintillometry measures fluctuations in the intensity of electromagnetic radiation due to turbulence. The scintillometer consists of a transmitter and receiver, which are separated by a known distance of tens of meters to a few kilometers. The scintillometer transmits electromagnetic radiation at optical or radio wavelengths (most typically in the visible or infrared portion of the electromagnetic spectrum). The receiver measures fluctuations in the intensity of the electromagnetic radiation, which are caused by turbulent fluctuations.

The refraction of light in the atmosphere is primarily controlled by temperature fluctuations, which cause fluctuations in the refractive index, n. Secondary contributions come from humidity and pressure fluctuations. This allows scintillometers to provide path-averaged measurements of the structure function parameter for temperature, C_T^2.

Scintillometers can be divided into classes based on the aperture and resulting path length. Small aperture scintillometers use a path length of approximately 50 - 200 m, as the signal is extinguished for path lengths greater than 250 m (Meijninger et al. 2002). Large aperture scintillometers use a longer path length of between 500 m and 5 km. The large aperture scintillometer signal is saturated at path lengths of approximately 5 km and greater (Meijninger et al. 2002).

Small aperture scintillometers can provide both the structure function parameter for refractive index, C_N^2, and the inner scale of turbulence, l_0. Large aperture scintillometers can be used to calculate area-averaged sensible and latent heat fluxes, by the employment of Monin-Obukhov similarity theory. Scintillometry has been shown to provide good estimates of the area-averaged heat fluxes for paths over homogeneous

terrain (e.g., Hill 1997). Experiments calculating the area-averaged heat flux over heterogeneous terrain via scintillometry have shown reasonable agreement with heat fluxes calculated an eddy covariance measurements (Meijninger et al. 2002).

2.2.2.4 Sodar

Sodar systems operate in a similar manner to radar and lidar systems, but the transmitted pulse is acoustic. The acoustic signal is scattered by turbulent fluctuations of temperature at the Bragg scale (approximately 0.05 - 0.3 m for typical sodar operating frequencies). The backscattered acoustic signal is recorded by microphones, and the power spectrum of the returned signal is used to determine the Doppler shift due to non-zero radial velocity along the sodar beams. This allows for the determination of radial velocity profiles. For a sodar operating using a single vertical beam, only the vertical velocity, w, is retrieved. For sodars operating with three or more independent beams (typically three or five beams are used), the horizontal wind components can be derived using the DBS technique, as for radar and lidar.

Simulation of a sodar system is the main focus of this dissertation, so a review how sodars are used to monitor the ABL is provided in Section 2.3. Further detail on all of the sodar operating parameters and issues pertaining to sodar hardware and software are presented in Chapter 3.

2.2.2.5 Radio acoustic sounding system

Radio Acoustic Sounding Systems (RASS) simultaneously release both an acoustic and an electromagnetic pulse (Marshall et al. 1972). RASS systems can be divided into two types, Bragg-RASS and Doppler-RASS systems. Bragg-RASS systems are operated as an addition to a wind profiling radar, while Doppler-RASS are operated as an addition to a sodar. The main utility of RASS is that the range of the returned acoustic and electromagnetic pulses can be matched according to the returned signal.

Exact knowledge of the height of each range gate (provided by the electromagnetic pulse) allows for the calculation of a speed of sound profile from the acoustic pulse return, which can in turn be used to derive a profile of the acoustic virtual temperature, T_{av}, which is proportional to the square root of the speed of sound. The deviation of the acoustic virtual temperature measured by RASS, T_{av}, from the air temperature, T, is dependent on the humidity content of the air. The relation between T and T_{av} can be written as

$$T_{av} = T(1 + 0.513q) \tag{2.8}$$

where q is specific humidity. For most purposes, especially if the air is relatively dry, it is sufficient to use the acoustic temperature as a proxy for the air temperature.

2.3 The use of sodar for studying the boundary layer

The first use of sodars to study the ABL occurred during the early 1970s (Kallistratova and Coulter 2004), and their use has continued to increase to this day. The specific applications of sodar for boundary-layer monitoring and characterization are many, of which some examples will be provided here.

2.3.1 Turbulence structure

The examination of turbulence structure in the atmosphere is one of the primary uses for sodars. The development of sodar as a remote sensing technique was based upon the knowledge that the propagation of acoustic waves is affected by small-scale turbulent structures in the atmosphere.

Acoustic scattering intensity, as a function of the angle between the transmitted beam and the backscatter, was shown by Tatarskii (1971) to depend upon the variability in acoustic refractive index (specifically, through the structure function parameter for refractive index, C_n^2) at the Bragg scale. When only direct 180° backscatter is considered, the dependence simplifies to be proportional to the structure function

parameter for temperature, C_T^2, at the Bragg scale. The acoustic returned power from a monostatic sodar (i.e., for backscattering at 180°) is known to depend directly upon the structure function parameter for temperature, C_T^2 (Little 1969). This means that an examination of the sodar echo return provides information about the thermodynamic structure of the atmosphere.

After sunrise, the boundary layer grows in height as convective heating causes turbulent mixing. As the heating increases during the day, the turbulent eddies grow in size. The maximum size of the turbulent eddies defines the height of the mixing layer. Sodars can be used to identify the mixing layer height, and many studies have sought to evaluate sodar-derived mixing layer heights against those from other instruments (see Section 2.3.2.1).

Sodar echograms (i.e., charts of acoustic returned power profiles with height) have also been used to study more fine-scale turbulence structure in both the convective and stable boundary layer. Many questions still exist in the topic of turbulence under stable and very stable conditions, and sodar is well placed to allow an examination of this topic. Anderson (2003) used sodar to examine the fine-scale turbulence structure over Halley Research Station in Antarctica, a region well-known for its strongly stable profile. The high-resolution sodar observed the propagation of internal gravity waves, and comparisons with tethersonde measurements indicated that the presence of gravity waves precludes a directly proportional relationship between the acoustic returned power and C_T^2. However, it was found that when gravity waves and other disturbances were not present, i.e., in ideal conditions, that C_T^2 profiles could be calculated directly from sodar returns.

Measuring vertical profiles of turbulence throughout the whole boundary layer is a challenge in any experimental campaign. Measurements made using *in-situ* sensors are typically limited to the lowest few meters of the surface layer, unless mounted on a tower (which is impractical for many field locations), and they typically provide only

point measurements. While these are useful, they do not provide any information about the spatial variation of variables of interest. Most of the widely used remote sensing instrumentation (e.g., wind profiling radars, lidars) have a minimum range that is on the order of a hundred meters. Sodar is well-placed to bridge the gap in vertical measurements between *in-situ* sensors and lidars or wind profiling radars.

The minimum range of a sodar is dependent upon the particular sodar model used. Most commercial systems have a minimum range of tens of meters above the ground, and a range resolution of around 10 m. The use of high-frequency mini-sodars can reduce the minimum range and allow for more detailed sampling of the surface layer.

2.3.2 Boundary layer climatology

The primary use for sodar systems is to monitor the atmospheric boundary layer dynamics and structure. This can take a number of forms, as sodars provide a number of variables that are useful for boundary layer studies. One of the most frequent uses for sodars is to derive or estimate the height of the mixing layer. This is discussed in more detail in section 2.3.2.1.

Sodars operating in three- or five-beam modes also have the capability to provide the three-dimensional wind components over the lowest few hundred meters of the atmosphere (the maximum height range is dependent upon both the specific make and model of sodar used, and the atmospheric conditions). These wind profiles are widely used for monitoring the boundary layer, and as inputs to weather forecasting models, in particular, for use in air pollution models.

Air pollution meteorology is an area of research for which sodars are highly useful. Sodars can provide several of the parameters necessary as inputs to air pollution modeling: aside from profiles of the three-dimensional wind components, sodars can also provide estimations of the height of the mixing layer (as discussed in further detail in Section 2.3.2.1). Sodar power return data also provides information on the

atmospheric stability, and the height of any present inversions, both of which are also necessary for air pollution modeling. Sodars are also capable of providing information about the variance of the horizontal and vertical velocity components, σ_u^2, σ_v^2, and σ_w^2, which are used in air pollution models.

Several papers have discussed in detail the use of sodar data to enhance air pollution and dispersion modeling. Gera et al. (2013) used a decade of sodar observations to examine the ventilation over Delhi, India, as it relates to pollution dispersion. Using the sodar power returns, the timing and duration of fumigation hours over the city was investigated. It was found that downward fumigation of pollutants trapped in the stable boundary layer was highly seasonally dependent, with average daily fumigation duration of 6 hours during the winter and only 2 hours during the summer (Gera et al. 2013). The fumigation duration is an important parameter for respiratory health hazard prediction, illustrating how sodar data has the potential for utility in a number of topics outside of meteorology.

2.3.2.1 Mixing layer height estimation

The mixing layer height is a very important parameter in air quality and pollution modeling. Within the well-mixed layer, turbulence resulting from surface heating mixes passive quantities such as gases and particles, resulting in a somewhat uniform distribution. Entrainment of air above the mixing layer is impeded by a capping inversion, limiting the height of the mixing. Thus the mixing layer height plays an important role in determining the maximum transport of air pollutants.

Sodars have been used to estimate the height of the mixing layer since their development in the 1970s (e.g., Russell and Uthe 1978; Coulter 1979). Sodars are well placed to evaluate the mixing height, since the sodar returned power is a function of the spatial variation of temperature, hence, strong temperature inversions show up strongly in the sodar return.

Aside from direct determination from sodar power returns, the mixing layer height can also be estimated from sodar-derived Doppler wind speed or vertical velocity variance, although the latter methods are appropriate for use only in the well-developed convective boundary layer. A full discussion and critique of each of these methods can be found in Emeis et al. (2008).

Mixing layer height estimation using vertical velocity variance assumes that the height of maximum vertical velocity variance occurs at a height of az_i, where z_i is the mixing layer height and a is a constant of typical value 0.35 - 0.4. This allows for mixing layer height estimates that are up to 2.5 times higher than the maximum sodar range. Due to the extrapolation involved, this method is considered unreliable (Beyrich 1997).

The most widely used method for estimating mixing height from sodar return makes use of the sodar power return, which is a function of the acoustic backscatter intensity. Within the well-mixed layer, the acoustic backscatter intensity is high, due to thermal fluctuations. Evaluation of the gradient of the sodar return power identifies the mixing layer height by a sharp drop-off in the power of the returned acoustic signal. Secondary maxima in returned power above this height are indicative of lifted inversions.

Several algorithms exist for the automatic determination of mixing layer height from sodar returns. The basic method utilized involves an analysis of the gradient of acoustic backscatter intensity, as described above. A discussion of the specific threshold values used in some of these algorithms is provided by Emeis et al. (2008).

2.3.3 Sodar usage in mountainous terrain

Sodar, like many other remote sensing technologies, uses the DBS technique. With DBS, three or more independent co-planar beams are used to derive the three-dimensional wind field. The DBS technique requires the assumption that wind is homogeneous over the sampling volume.

In situations where this assumption holds, i.e., the winds are horizontally homogeneous within the area bounded by the off-vertical sodar beams, sodar wind estimates have been shown to match very well with winds measured by mast-mounted instruments. Crescenti (1996) compiled twenty years of sodar comparison studies, and found that the mean correlation coefficient between reference and sodar-measured wind speeds across the twenty studies was 0.91. The more recent UpWind (2011) study compared sodar and mast-mounted instruments for steady flow over homogeneous terrain found a correlation with $R^2 > 0.985$.

However, when the winds are not horizontally homogeneous (e.g., in complex terrain), the radial velocities measured by each independent beam contain different wind components. It can be considered that for n independent beams, the radial velocities contain contributions from n different values of u (u_1, ..., u_n), v, and w.

Studies by Behrens et al. (2012) and Bradley (2012b) found errors in wind speed estimates of up to 8% when the DBS technique was used in complex terrain. Quantification and correction of these errors have been performed for through the use of modeling the flow over specific complex terrain (e.g., Bingöl et al. 2009; Bradley 2012b).

2.3.4 Sodar usage in remote terrain

One of the advantages of sodar technology is that once installed, it requires little routine maintenance. As such, sodars have been installed in remote locations, such as the Antarctic, where other remote sensors cannot provide long-term measurements.

Initially, sodar use in the Antarctic was limited to sodar systems located at manned stations with mains electric power (e.g., Hall and Owens 1975; Neff 1978b). During the 1980s, sodars were used in measurement campaigns at Antarctic coastal sites, although they were always located at a research station with access to the electric grid (Anderson et al. 2005). Sodars located in the Antarctic had required access to a power grid in order to run the antenna heating, which prevents a build-up of rime on the speakers or antenna, which would interfere with the measurements.

Anderson et al. (2005) report on the installation of an autonomous Doppler sodar in the remote Coats Land area of the Antarctic. The power required by the sodar was provided by two photovoltaic solar panels and two wind generators. This allowed the sodar to run autonomously for a two-year field experiment, providing the first comprehensive multi-year wind-profile measurements at an unmanned remote Antarctic station.

The autonomous Doppler sodar provided the first winter-time measurements of the Antarctic katabatic winds at a remote station. These measurements illustrated that for flows that are primarily katabatically driven (i.e., not synoptically driven), there is a strong dependence of the depth and strength of the katabatic low-level jet on wind speed. A positive correlation was found between the maximum wind speed in the low-level jet and both the height of the jet maxima and the depth of the katabatic flow (Renfrew and Anderson 2006).

2.3.5 Sodar usage in urban areas

Sodars can also be used to measure wind profiles in the complex urban environment, albeit with some caveats. In order to use sodar in an urban or suburban environment, noise pollution effects must be considered, both in the effects of the noise caused by the sodar on the local population; as well as the effects of high background noise inherent in the urban environment, and reflection of acoustic waves by the irregular

urban geometry, upon the sodar measured variables. However, knowledge of the wind profiles in urban areas is very important. Air pollution modeling above urban areas is a topic of increasing importance, and accurate urban wind profiles are needed as inputs for these models.

Yushkov et al. (2007) analyzed a full year of sodar measurements taken at two sites within Moscow, Russia, finding that the high levels of background noise, specifically from automotive traffic, adversely affected the sodar signal-to-noise ratio, which reduces the maximum useable retrieval height. Sodar measurements recorded over a period of 17 months at an urban site in Hannover, Germany were used by Emeis and Türk (2004) to examine the mixing-layer height. For the sodar measurements recorded in Hannover, the sodar was placed 550 m upstream of the street canyon which was the focus of the study. This ensured that the sodar was located in an industrial area away from residential properties, and as such was not creating noise pollution for local residents.

A major challenge in using sodar in urban areas is that it can be very difficult to separate the effect of the urban area itself on the wind profile from effects due to the surrounding orography. One way to overcome this issue is for a network of sodars to be used, encompassing the transition from rural to suburban to urban topography. This method was employed by Kallistratova and Kouznetsov (2012), who used three sodars to study low level jets over Moscow. The first sodar was located in a rural area 50 km west of Moscow, the second at a site in the south-west of the city, and the third at a site in the city center. The use of the three sodars at different locations around the city allowed for an examination of the effect of urbanization (specifically, the urban heat island) on the formation and duration of low level jets. It was found that the increased heat storage in the urban area, coupled with the increased surface roughness (compared to the rural site) caused the urban area to see low level jets that were weaker and higher, and emerged later, than the corresponding rural site.

2.3.6 Electromagnetic wave propagation

One of the first uses of sodar technology was as a tool for evaluation of the ease of optical propagation through the atmosphere. Small-scale fluctuations in the temperature, and thus refractive index, alter the propagation of both electromagnetic and acoustic waves through the atmosphere. The direct relationship between the structure function parameter for temperature, C_T^2, and the structure function parameter for refractive index, C_n^2, motivated the use of sodar for studying scintillation and the propagation of light waves through the boundary layer.

The lower the value of C_T^2, the smaller the impact of temperature fluctuations on light propagation. Sodars can be used to evaluate the vertical profiles of C_T^2 in the atmosphere. This is an important topic for astronomy, as the optimal siting of observatories with large-scale telescopes depends upon finding locations with the best optical "seeing" conditions.

The use of sodar for examining telescope test sites was investigated by Travouillon et al. (2011), who compared turbulence profiles derived from two different sodar models with profiles measured by a multi-aperture scintillation sensor coupled with a differential image moon monitor. The sodar-derived turbulence profiles compared well with those from the scintillometer, and the scatter between profiles measured by four sodars of two different models was within acceptable limits. Travouillon et al. (2011) thus concluded that sodar is a useful instrument for measuring turbulence profiles in the lower atmosphere, and sodar can be used complementary to a multi-aperture scintillation sensor to provide turbulence information over the entire height region of interest.

Chapter 3

Sodar theory

Prior to accurate simulation of a sodar, it is important to understand fully the underlying principles and techniques used in sodar technology. In this chapter, the various components of commercial sodars, both hardware- and software-based, will be discussed. Section 2.1 provided an overview of aspects of the boundary layer for which sodars provide a means to study.

The theory underlying both acoustic scattering and sodar remote sensing is well established. The following section details the development and basic principles of acoustic scattering theory. The remainder of the chapter outlines the theory behind sodars, the basics of which have remained unchanged since they were laid out by Little (1969). These principles will be expanded upon in detail, with special attention paid to how the choice of sodar hardware, and the choice of operating parameters, affects the sodar retrievals.

3.1 Acoustic scattering theory

The scattering of acoustic waves by turbulence in the atmosphere has long been observed: it was initially noted, albeit not well-understood, by Reynolds (1876). In this paper, experimental data is reported in which rockets powered by gunpowder were launched along The Wash, UK, and whether or not ships located at varying distances and directions from the rocket could hear the sound was noted. Reynolds remarks upon the "stoppage of sound by the heterogeneity of the atmosphere", but states that an atmosphere of the required heterogeneity to produce such an effect would surely be observed to be flickering. For this reason, the difference in reported

sound of the explosion at the ship locations due to scattering of the sound waves was instead ascribed to refraction in the atmosphere.

A theoretical approach to the scattering of sound by turbulence in the atmosphere was developed during the 1940s. An outline and a general solution to the theoretical problem of acoustic scattering was first laid out by Obukhov (1941). It was asserted by Blokhintzev (1946) that sound at audible frequencies is primarily scattered within the internal subrange of turbulence. As acoustic scattering theory continued to develop in the 1950s Russia, theoretical aspects of acoustic scattering due to turbulence were being developed in parallel in the USA by Lighthill (1953), Kraichnan (1953), and Batchelor (1957).

The calculation of the acoustic scattering cross section in terms of temperature and velocity spectra followed the work of Monin (1962). A comprehensive theory of the effects of atmospheric turbulence upon acoustic wave propagation was laid out by Tatarskii (1961, 1971), which provides the basis of the underlying theory behind acoustic remote sensing that is still in use at the current time. The theoretical background following the work of Tatarskii is outlined in the following section.

3.1.1 Underlying principle of acoustic remote sensing

The underlying principle of acoustic remote sensing relies on the transmission of an acoustic pulse from a speaker or set of speakers. The transmitted pulse is of a known frequency and has directionality. The pulse propagates through the atmosphere, and as it propagates, the energy in the pulse is scattered by small-scale temperature inhomogeneities. Some of this energy is scattered, in the form of an acoustic echo, back towards the speaker (in the case of a monostatic sodar system) or toward the receiving speaker (in the case of a bistatic sodar, discussed briefly in section 3.2), which records the received echo. The temperature inhomogeneities are caused by turbulence, and these small-scale inhomogeneities are assumed to move with the mean wind.

The scale of the turbulence (and of the temperature inhomogeneities), l_{t1}, is related to the transmitted acoustic pulse through

$$l_{t1} = \frac{\lambda_a}{2 \sin \theta_B} \tag{3.1}$$

where l_{t1} is the scale of the turbulence expressed in m, λ_a is the wavelength of the transmitted acoustic frequency expressed in m, and θ_B is the Bragg angle, i.e., the incident angle of the acoustic pulse. In other words, temperature inhomogeneities of a size l_{t1} maximize their contribution to the scattering of acoustic pulses of wavelength λ_a at angles of θ_B. This simplifies to $l_{t1} = \lambda_a/2$ for direct backscattering , where θ_B=90°.

The propagation of the acoustic wave is also affected by Fresnel diffraction, which accounts for the diffraction of the wave caused by its passage through temperature inhomogeneities. This effect of the Fresnel diffraction is to contribute to fluctuations in the amplitude of the acoustic wave. For the same acoustic pulse with wavelength λ_a, the scale of the inhomogeneities responsible for Fresnel diffraction is

$$l_{t2} \approx (\lambda_a L)^{1/2} \tag{3.2}$$

where L is the distance over which the acoustic wave propagates.

For acoustic pulses of typical sodar frequencies (i.e., ∼2000 Hz) and for L representative of propagation within the lower boundary layer, both l_{t1} and l_{t2} are on the order of centimeters to meters (Kallistratova 2002).

The classical theory of acoustic scattering in a turbulent medium proposed by Tatarskii (1971) uses the formulation of Monin (1962), which defines the acoustic scattering cross section as

$$\sigma_s(\theta) = 2\pi k_a^4 \cos^2\theta \left(\frac{\Phi_T(k_t)}{2T^2} + \cos^2(\theta/2) \frac{\Phi_v(k_t)}{c^2} \right) \tag{3.3}$$

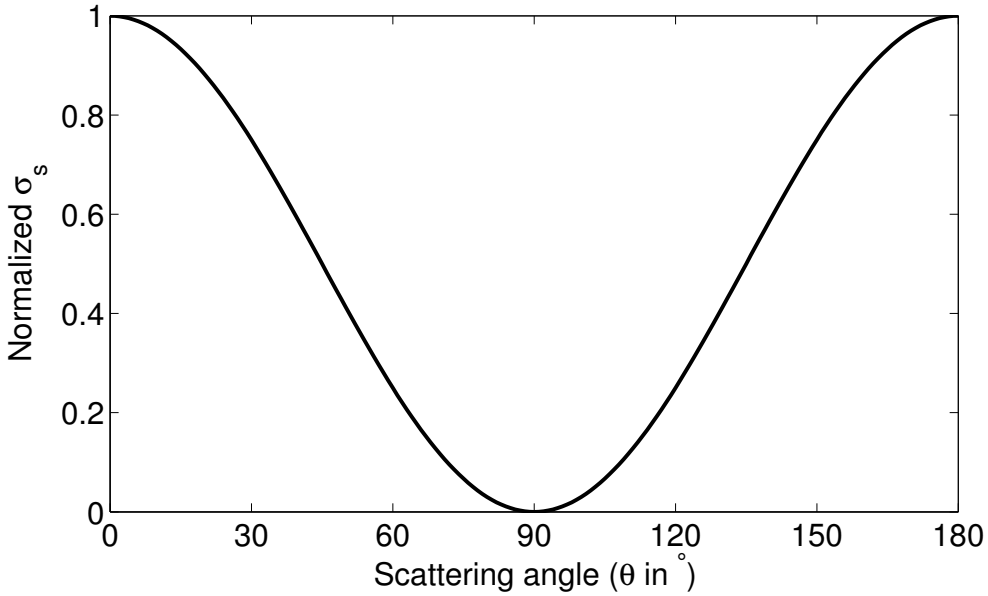

Figure 3.1: The variation of the acoustic scattering cross section σ_s with scattering angle, calculated from equation 3.3. The scattering cross section is normalized.

where $\Phi_T(k_t)$ and $\Phi_v(k_t)$ are the Kolmogorov spectra for temperature fluctuations and wind speed fluctuations for locally homogeneous and isotropic turbulence, respectively. These are calculated as

$$\Phi_T(k_t) = 0.033 C_T^2 k_t^{-11/3} \tag{3.4}$$

and

$$\Phi_v(k_t) = 0.061 C_v^2 k_t^{-11/3} \tag{3.5}$$

where C_T^2 and C_v^2 are the structure function parameters for temperature and velocity. In equations (3.3) - (3.5), $k_a = 2\pi/\lambda_a$ represents the wavenumber of the acoustic pulse, and $k_t = 2\pi/l_{t1}$ represents the wavenumber of turbulence of length scale l_{t1}. The turbulent wavenumber and the acoustic wavenumber are related as $k_t = 2k_a$, through

$$k_t = \frac{2\pi}{l_{t1}} = \frac{k_a \lambda_a}{l_{t1}} = 2k_a. \tag{3.6}$$

As with equation 3.1, the acoustic scattering cross section is greatly simplified when considered for direct backscattering (θ=180°) only. In this case, σ_s is reduced to

$$\sigma_s(180^\circ) = 2\pi k_a^4 \frac{\Phi_T(2k_a)}{2T^2}, \tag{3.7}$$

i.e., all contributions due to velocity inhomogeneities are removed and the backscattering cross section depends on turbulent temperature inhomogeneities only. Also, as can be seen from Fig. 3.1, for given values of C_T^2 and C_v^2, the scattering cross section is maximized for a scattering angle of 180°, i.e., for direct backscattering.

For turbulence within the inertial subrange, where the turbulence can be considered locally isotropic and homogeneous at the characteristic scale of scattering inhomogeneities, the structure function C_T^2 can be related directly to the acoustic backscattering cross-section by combining equations 3.4 and 3.7 to give

$$C_T^2 = 140\sigma_s(180^\circ)T^2\lambda^{-1/3}. \tag{3.8}$$

This is useful, as the acoustic backscattering cross-section can be calculated from the sodar return power (through equation 3.10), as discussed in section 3.3.2. Since σ_s can be calculated from a sodar output parameter, this means that sodar returns can be used to estimate C_T^2. Full details of the method for estimating C_T^2 from sodar output parameters are given in section 5.3.5.

Kallistratova (2002) states that the condition of l_{t1} being within the inertial subrange of turbulent scales is always met for sodars of operating frequencies 1500 - 6000 Hz, which corresponds to acoustic pulse wavelengths of λ_a=0.055 - 0.22 m for a speed of sound of 330 m s^{-1}. All commercially available sodars fall within this frequency range, which ensures that the estimation of C_T^2 from commercially available sodars is appropriate.

3.2 Types of sodar

There are two main types of sodar: those using dish antennae and those using phased array antennae. Sodars using dish antennae are designed such that a downward-facing speaker is mounted above a parabolic dish. Sodars of this type are less common, particularly amongst commercially available sodars, and are less widely-used than the phased array type antennae. As such, parabolic dish antennae-based sodars will not be considered further here.

Sodars using phased array antennae have two conventional speaker setups. The speakers are arranged either in an upward-facing horizontal array, or they are in an array which is placed at an angle from a reflector panel. In the second case, the speakers themselves may be aligned in a vertical array, and placed inside a weather-proof housing. The use of housing minimizes weathering of the speakers, and protects them from debris. Speakers arranged in the former setup, the horizontal array, are typically recessed to prevent precipitation interacting with the speakers directly. This is done in order to mitigate the acoustic noise caused by precipitation. However, both of the mentioned array structures are susceptible to precipitation noise, whether the precipitation is interacting with the outside of the speakers (as with the horizontal array), or the reflector panel (as with the reflector array). In most commercially available sodars, the acoustic signal is overwhelmed by noise during precipitation, meaning that atmospheric properties cannot be accurately retrieved during periods of precipitation.

Both types of phased array antennae discussed above have similar beam geometries, and the differences between the two types are mostly of importance for ease of equipment maintenance, and protection from the elements.

Sodars can also be monostatic or bistatic: this refers to whether or not the same speakers are used for signal transmission and receiving. A monostatic sodar contains only one speaker array, used for both transmitting and receiving. A bistatic sodar,

on the other hand, possess two sets of speaker arrays, one of which is used exclusively for signal transmission, whilst the other is used only for receiving the signal. The two speaker arrays on a bistatic sodar are typically separated by a baseline distance of 100-300 m.

There are advantages and disadvantages of both monostatic and bistatic systems. Monostatic sodar systems do not require much space to deploy, while bistatic systems require at least as much space as the baseline distance. Monostatic systems use the DBS, which involves the assumption of spatial continuity in the measured wind field across the area encompassed by the off-vertical beams. On the other hand, bistatic systems determine the winds from a common volume of air, so the spatial continuity assumption is not necessary. Since monostatic systems use the same speakers for sound transmission and receiving, the effective angle of scattering is restricted to 180° (i.e., direct backscattering only), such that equation 3.3 reduces to equation 3.7. This ensures that the structure function parameter for temperature can be estimated from the power return of monostatic sodars. On the other hand, bistatic sodars receive signals scattered by both thermal and mechanical turbulence. Crescenti (1996) states that monostatic sodars are more susceptible to contamination from environmental background noise, whilst bistatic sodars are more susceptible to ground clutter effects from refraction and reflection of the signal by objects on the ground along the baseline distance. In order to estimate C_T^2 from bistatic sodars, a method for estimating the velocity spectra and removing the contributions of Φ_v to σ_s is required. Currently, all commercially manufactured sodars are of the monostatic type, while bistatic sodars are used primarily for specific research purposes (e.g., Bradley et al. 2012). As such, only monostatic sodar systems will be considered further in this study.

3.3 Beam transmission

Parameters which affect the transmission of a sodar beam can be split into two categories. The first category is parameters which are controlled by the physical sodar design, including parameters which are limited by the sodar hardware. The second category are parameters which the end user is free to set, albeit perhaps from a limited range. Naturally, there is overlap between the two categories, for parameters such as transmission frequency, which is chosen by the user from a pre-specified range, dependent upon the particular make and model of sodar.

Parameters of either type which directly affect the beam transmission are the transmitted frequency, the transmitted power, the acoustic pulse length, the wait time between pulses, the signal rise time, the pulse beamwidth, and the off-vertical beam elevation angle. Of these parameters, those which are controlled by the sodar design explicitly are the maximum transmit power, the acoustic pulse length, the signal rise time, and the pulse beamwidth.

3.3.1 Frequency

All sodars have a range of transmit frequencies within which they can operate. The user can select their desired transmit frequency from within this range. As an example, the Metek PCS.2000 sodar operated by the University of Oklahoma has an operating frequency range of between 1700 Hz and 4000 Hz.

The chosen transmit frequency has an affect upon the maximum range of the beam, as acoustic absorption is highly frequency-dependent (discussed further in section 4.2.5). The choice of transmit frequency also affects the sodar signal-to-noise ratio (SNR), as the background acoustic noise is frequency-dependent. More details on acoustic background noise are given in section 3.5.2, but it is noted here that the background noise amplitude tends to increase at decreasing frequencies and vice versa. It is clear that there are competing effects between maximizing range and

maximizing SNR, and as such, the sodar transmit frequency should be chosen based upon the desired maximum range, with an allowance for the typical background noise at a particular site.

3.3.1.1 High frequency mini-sodar

During the last few decades, as sodar technology has continued to develop, sodar has expanded outside of the original frequency range of 1000-2000 Hz. One of the shortcomings of sodar operating within this range of frequencies is the minimum height of the retrieved signal is often several tens of meters above the ground. The reason for this is that if the sodar speakers are of the voice-coil-based electo-acoustic type, and the sodar is monostatic (i.e., the same speakers are used for transmission and receiving), then a transmit-to-receive delay is required. The speaker membrane continues to oscillate for a short time after the acoustic pulse, so a transmit-to-receive delay is used to delineate between the residual speaker oscillations, and real atmospheric acoustic echoes (Argentini et al. 2012).

Attempts to circumvent the problem of the high minimum range resulted in the development of the high-frequency mini-sodar. The idea for a high-frequency mini-sodar was tested by the NOAA Wave Propulsion Laboratory in the 1970s. Their system used the same type of electro-acoustic transducers as convectional sodars with tiny parabolic dishes, and was known as the Suitcase Sounder (Owens 1974). Further developments by Moulsley and Cole (1979) used an array of piezo-electric transducers rather than voice-coil-based transducers, significantly reducing the ringing time of the speakers, thus shortening the necessary transmit-to-redeive delay. For the piezo-electric transducers, an increase in the transmit frequency from 1730 Hz to 4700 Hz allowed a reduction in the minimum range from 28 m to 7 m (Moulsley and Cole 1979). The main limitation of piezo-electric transducers is that they can only handle a certain

acoustic power load, which limits the maximum retrieval range due to attenuation of the signal.

Improvements in the handling of the electronic system for recording the sodar echoes lead to the development of a high resolution sodar with a very low minimum range, designed for studying the surface layer (the surface-layer mini-sodar; SLM-sodar, Argentini et al. 2012). In this system, three separate electro-acoustic transducer speakers operating at 5000 Hz are placed around the receiving speaker. The acoustic echoes recorded by the receiving speaker are handled separately from the transmitted signals within the electronics, which removes the "ringing" effect following the tone emission in the received signal (Argentini et al. 2012). The receiving speaker can thus begin recording immediately following the tone emission, yielding a minimum range of 2 m. The tone is emitted every 1 s for a duration of 10 ms, which results in a vertical resolution of approximately 1.7 m. The high-resolution mini-sodar was tested in Rome comparisons were made between the resulting wind profiles and those measured by a collocated conventionally-operating sodar and three sonic anemometers. It was found that the nocturnal boundary layer contained fine-scale structures that were of too small a size for the conventional sodar to resolve, but the high range resolution of the SLM sodar enabled these structures to be clearly seen. Hence the SLM-sodar illustrated its suitability for nocturnal surface layer studies, with the ability to accurately portray high-resolution temporal and vertical inhomogeneities within the surface layer structure.

3.3.1.2 Multi-frequency sodar

Conventional sodars use a single transmit frequency for all of the independent beams, and the pulse is usually transmitted as a sine wave of the desired frequency. The use of multiple frequencies in acoustic sounders was initially developed with the goal

of improving sodar performance. Several different methods of incorporating multiple frequencies into the pulses have been tested, and these are outlined below.

One of the earlier methods used different transmit frequencies for each independent beam direction (or, in the case of bistatic sodars, for each separate transmit antenna). For monostatic sodars, using a different transmit frequency for each beam can allow for a reduced time between pulses, since return signals from the separate beams are easily distinguishable if the transmitted frequencies are sufficiently separated such that there will be no crossover between the returned acoustic signals from each beam. For bistatic sodars, using different transmit frequencies for each antennae allows for easy separation of the received signals from each antennae for the same reason.

More recently, the transmission of multiple frequencies within a single acoustic pulse has been achieved by the use of pulse compression. The most common pulse compression methods used with sodars are the release of a chirped signal, or the use of phase encoding. The chirped signal method uses an acoustic pulse whose transmit frequency changes linearly throughout the duration of the pulse (typically, the frequency is increased throughout the pulse). Some sodars also use approximations to the chirp method, e.g., a step-chirp, where the pulse frequency increases in a series of abrupt steps throughout the pulse. The phase encoding method uses a sine wave whose phase is altered as a step-change by multiples of π during the pulse (Bradley 1999).

An advantage to both of these methods is that they increase the bandwidth of the signal. Each of the pulse coding methods requires the addition of a matched filter in the processing of the received signal. The matched filter compresses the pulse, decreasing the pulse duration and enhancing the range resolution. Bradley (1999) investigated the applicability of several different pulse coding systems to sodar systems by testing coded waveforms on simulated atmospheric targets, both with and without a Doppler-adaptive filter. It was found that without a Doppler-adaptive filter, the

step-chirp pulse performed the best with regards to a combination of being immune to Doppler influence and increasing the sodar range resolution. Three Doppler-adaptive filters were also tested, with the optimized frequency shift method performing the best, producing the highest accuracy in wind speed estimates (Bradley 1999). This study proved theoretically that using pulse coding techniques in sodars can provide accurate wind speed estimates, and provide both an increase in signal-to-noise ratio and a decrease in range resolution, when compared to a single frequency transmitted pulse.

The multi-frequency pulse coding was tested in the field on a parabolic dish antenna by Kouznetsov (2009), with the finding that the theoretical improvement in sensitivity (and corresponding increase in SNR) predicted by Bradley (1999) did occur, with maximum increases in SNR of 3-5 dB at a height of 100 m.

3.3.2 Transmitted power

The power transmitted by the sodar, in simple terms, relates to the intensity of the acoustic pulse. It is desirable for the sodar to transmit acoustic pulses at the highest possible power without damaging the speakers, while mitigating a noise nuisance to the surroundings. The reason that a high transmitted power is preferable is that the power of the received signal is proportional to the power of the transmitted signal. The returned sodar signal power was first defined by Little (1969) as

$$P_R = \frac{P_T G A_e \sigma_s L c \tau}{2r^2} \tag{3.9}$$

where P_R is the returned signal power measured in W, P_T is the transmitted signal power measured in W, G is a dimensionless factor representing the gain of the antenna, A_e is the antenna effective aperture area expressed in m^2, σ_s is the acoustic backscattering cross-section expressed in m^{-1}, c is the speed of sound in m s^{-1}, τ is the pulse duration expressed in s, r is the signal range in m, and L is an attenuation factor which accounts for sound transmission and receiving efficiency, as well

as the effects of classical, molecular, and excess attenuation. So, the more power in the transmitted signal, the greater the power in the returned signal. For every field site, there is a threshold of background noise level above which the signal return becomes undistinguishable from the background noise. Maximizing the transmitted signal power ensures that the return acoustic signal is detectable across a wide range of field sites.

The attenuation term L in equation 3.9 was examined by Neff (1978a) and separated into two terms, $\exp(-2\alpha r)$ and $L_E(f)$. The $\exp(-2\alpha r)$ represents the roundtrip loss of power due to viscous and molecular relaxation processes, and the $L_E(f)$ term represents excess attenuation which was interpreted as a ratio of the signal amplitude received from two different transmit frequencies.

Combining 3.9 with the exponential term from losses due to attenuation suggested by Neff (1978a), we arrive at the sodar power equation in the form which is widely used today,

$$P_R = P_T G A_e \sigma_s \frac{c\tau}{2} \frac{e^{(-2\alpha r)}}{r^2}, \tag{3.10}$$

where P_T is the transmitted power in W, G is the antenna gain, A_e is the antenna effective area in m^2, σ_s is the backscattering cross section as described in equation 3.7, c is measured in m s^{-1}, τ is measured in s, α is measured in dB m^{-1}, and r is in m. A detailed description of the calculation of acoustic attenuation α can be found in section 4.2.5.

A comparison of sodars from many manufacturers (including both commercial and specialized research sodar models) performed by Crescenti (1996) found that the transmitted power across the different sodar models varied from 0.5 - 300 W, with the bulk of the commercial sodars having transmitted powers within the 50 - 200 W range.

3.3.3 Acoustic pulse length

The choice of the acoustic pulse duration, τ, affects the sodar output in a number of ways. The returned signal power is directly proportional to τ (as seen in equation 3.10), so a longer pulse duration will help to increase the returned signal power.

The choice of τ also affects the frequency resolution (and hence radial velocity resolution) of the sodar. The frequency resolution is calculated as

$$\Delta f = \frac{1}{\tau}. \tag{3.11}$$

A smaller value for Δf results in a higher radial velocity resolution, which allows for distinction between increasingly smaller changes in radial velocity. Both of these effects would suggest that a longer pulse duration is desirable. However, the pulse duration is also proportional to the sodar range resolution, which is calculated as

$$\Delta r = \frac{c\tau}{2}. \tag{3.12}$$

The competing effects of τ upon both the radial velocity resolution and the range resolution must be taken into account. This inability to satisfy both the need for a short pulse duration to maximize the sodar range resolution and the need for a longer pulse duration to maximize the velocity resolution is known as the Doppler dilemma. Some sodar studies may require a greater emphasis on high resolution measurements, while others may prefer to sacrifice high range resolution in favor of more accurate wind speed values. Since the pulse duration is typically a feature of the sodar preset by the manufacturer, the limitations caused by the chosen pulse duration are predetermined.

The 1996 sodar comparison study by Crescenti mentioned in the previous section also reported the pulse length of twenty one different sodar models. The reported pulse lengths varied from 7.8 - 400 ms, with most of the sodars in the 50 - 150 ms range.

3.3.4 Signal rise time

A rise time, $\beta\tau$, is introduced into the acoustic signal, where β is a dimensionless scaling factor. The value of β is equivalent to the fraction of the acoustic pulse length τ which is taken to ramp up the amplitude of the signal to its maximum value. The rise time is included by multiplying the pulse voltage by a filter, typically of the Hanning or Gaussian shape. The effect is to introduce a ramp up and ramp down to each pulse, rather than a quick jump from zero to maximum outage amplitude. The reason for this is twofold: the ramp up and down protect the speakers from damage which might occur from a sudden onset of maximum amplitude voltage, and the filtered (or windowed) pulse limits spreading in the frequency spectrum. Figure 3.2 shows the pulse voltage of an unfiltered pulse, and pulses filtered using Hanning and Gaussian shape windows, both using a rise time of $\beta = 0.25\tau$. A longer rise time will result in a smoother pulse envelope, which produces a smoother and wider spectrum. This has both advantages and disadvantages, as a smoother frequency spectra reduces the possibility as a spurious noise maxima being identified as the spectral peak, but a wider spectrum can reduce the sharpness of the spectral peak, making it harder to identify.

3.3.5 Pulse repetition frequency

The pulse repetition frequency is inversely proportional to the time length between pulses. The time duration between pulses is determined by the maximum desired range. A longer duration between pulses allows for a higher maximum range. It is necessary that the delay between pulses is sufficiently long that the acoustic pulse propagate to the maximum range of interest, and return to the speaker, before the next pulse is released. This ensures that the returns from each pulse can be identified and handled separately. One way to mitigate a long pulse delay time is the use of different frequency signals for each separate beam, as described in section 3.3.1.2.

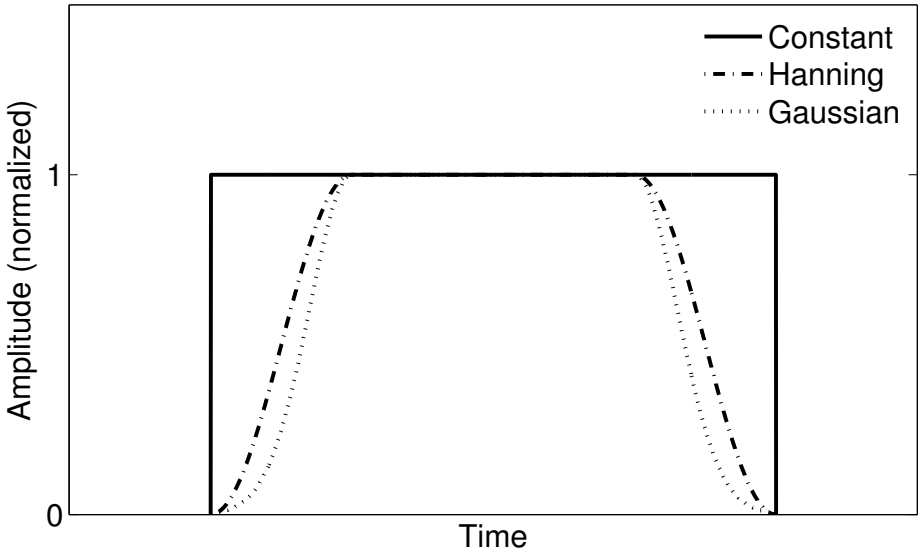

Figure 3.2: Filtering of pulse voltages. The solid line represents an unfiltered pulse, the dashed line is Hanning filtered, and the dotted line is Gaussian filtered ($\mu = \bar{t}$,$\sigma = 0.15$). Each filter has a rise time of 0.25τ.

3.3.6 Beam zenith angle

The zenith angle of the beam is a function of the sodar design, the transmitted acoustic frequency, and the atmospheric virtual temperature. The sodar design parameters upon which the zenith angle of the off-vertical beams depends are the spacing of the individual speakers within the speaker array, d, the number of speakers contained within the array, N.

For phased array sodars, generation of the off-vertical beams does not require the movement of any of the sodar hardware components, thus reducing wear on the system. The off-vertical beams are steered by using a phase lag between the speaker array elements. For a phase lag between array elements of φ, the zenith angle of the resulting beam is calculated as

$$\theta = \sin^{-1}\left(\frac{\varphi}{2\pi}\frac{\lambda}{d}\right) \tag{3.13}$$

where d is the speaker spacing, and λ is the transmitted acoustic wavelength (Ito et al. 1989). The acoustic wavelength, λ, is dependent upon the transmitted acoustic

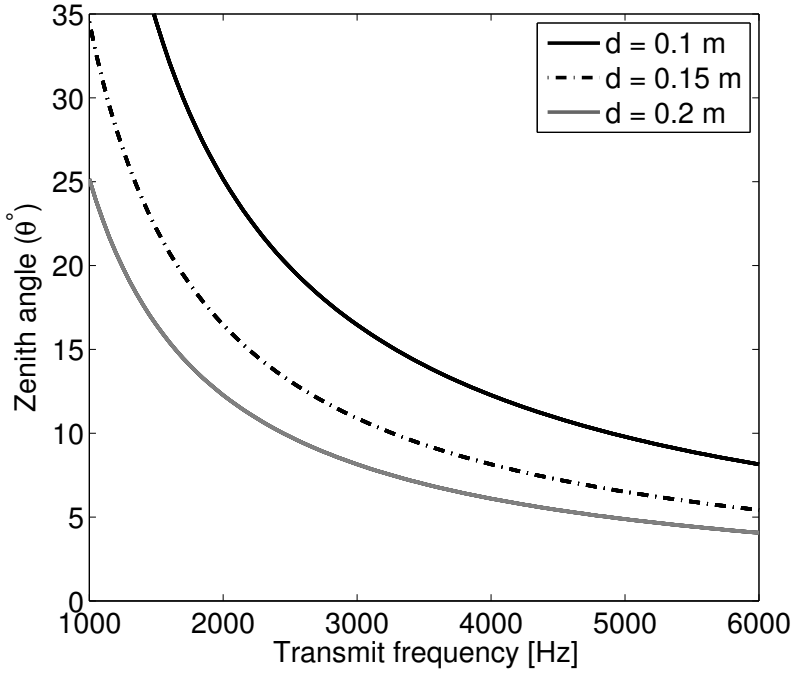

Figure 3.3: The variation of the off-vertical beam zenith angle with transmit frequency for three different speaker spacings d. The speed of sound was assumed constant at 340 m s^{-1} and the phase lag was kept at $\pi/2$.

frequency, and the speed of sound, which is a function of virtual temperature, $c = 20.05\sqrt{T_v}$. The dependency of θ on the transmit frequency is illustrated in Fig. 3.3. The zenith angle was calculated for three different values of the speaker spacing, 0.1 m, 0.15 m and 0.2 m, and for transmit frequencies ranging from 1000 Hz to 6000 Hz. The phase lag, φ, was held constant at $\pi/2$ and the speed of sound was assumed to be 340 m s^{-1}. As is clearly shown in Fig. 3.3, the zenith angle decreases significantly at higher transmit frequencies. For the same transmit frequency, a smaller spacing between the individual weaker array elements will produce a larger off-vertical beam zenith angle.

Sodars typically apply the phase lag between speakers in increments of $\pi/2$. For an acoustic pulse fashioned as a sine wave of the form

$$\sin \omega t, \tag{3.14}$$

where ω is equivalent to $2\pi f$ and f is the frequency of the transmitted pulse, this gives four resultant pulses of the forms $\sin \omega t$, $\sin(\omega t + \pi/2) = \cos \omega t$, $\sin(\omega t + 2\pi/2) = -\sin \omega t$, and $\sin(\omega t + 3\pi/2) = -\cos \omega t$.

For commercially available sodars, the user typically defines the transmit frequency, and not the zenith angle directly. The zenith angle is then calculated for the given frequency and measured temperature. It must also be noted from equation 3.13 that since the acoustic wavelength is dependent upon temperature, for a given input transmit frequency, the resulting zenith angle will vary slightly due to the virtual temperature of the atmosphere.

Of the eighteen sodar models for which the off-vertical beam zenith angle is reported in Crescenti (1996), the angles varied from 10° to 45°, with fifteen of the eighteen sodars using zenith angles of between 18° and 30°.

3.3.7 Beamwidth

The beamwidth of each of the individual beams is related to aspects of the sodar hardware, including the number of individual speakers contained in the array. The beamwidth is also closely related to the beam zenith angle, as both can be examined in terms of the antenna intensity pattern. For a square array comprising $M \times M$ speaker elements, the resulting antenna intensity can be written as

$$I \propto \frac{2J_1((M/2)kd\sin\theta)}{\sin((1/2)kd\sin\theta)} \tag{3.15}$$

where J_1 is a Bessel function of the first kind, d is again the spacing between array elements, and k is the acoustic transmitted wavenumber (equivalent to the inverse of the transmitted frequency). The resulting intensity pattern for a 8×8 array of speakers is shown in Fig. 3.4. It is clear from Fig. 3.4 that each beam receives contributions from a range of angles centered upon the given zenith angle. The finite beamwidth effectively results in spectral broadening.

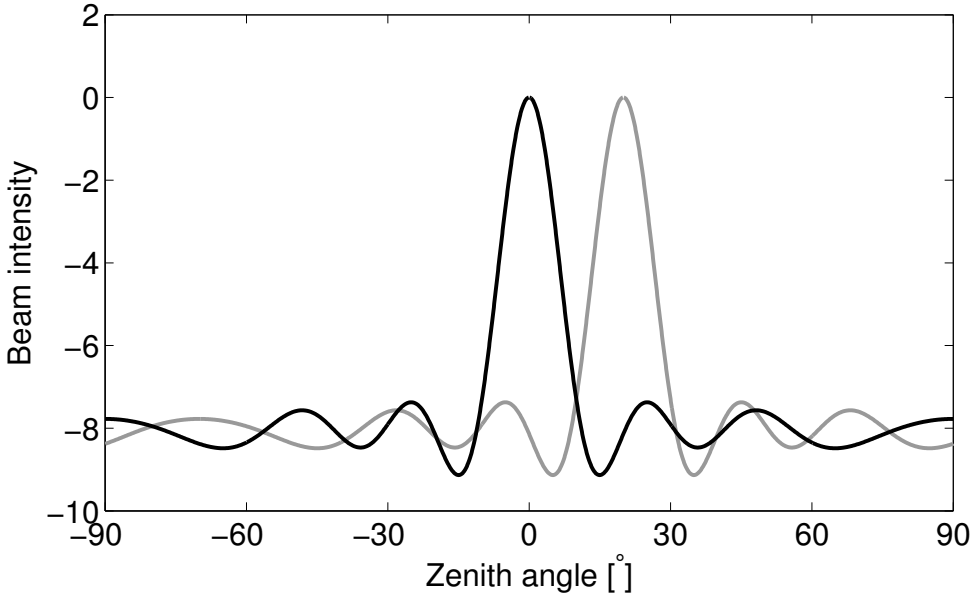

Figure 3.4: The beam intensity pattern from the vertical beam (black line) and an off-vertical beam with zenith angle of 20° (grey line). The represented array has 8×8 elements, zero phase gradient, and d/λ=5.

The intensity pattern resulting from equation 3.15 also does not include the effects of the sodar baffles, which add a level of complexity. The sodar baffles affect the beam propagation in a number of ways: the baffles block the beam at high zenith angles, which acts to minimize the side lobes. The baffles can also cause diffraction of the beam, with the diffraction pattern being particular to the specific baffle design.

In practical terms, equation 3.15 also limits the maximum zenith angle of the off-vertical beams. Figure 3.4 illustrates that the antenna intensity pattern produces side lobes at angles removed from the main beam direction. In order to avoid ambiguity between the main and secondary lobes, the maximum zenith angle is limited to a certain range, typically 15-30°.

Crescenti (1996) compiled sodar statistics from previously published sodar comparison studies, and reported the beamwidth used by seven different sodar models. The beamwidth angle in degrees varied from 7.5° to 11° in the reported sodars.

3.4 Beam receiving

After the acoustic pulse is transmitted, it interacts with and is scattered by the turbulent atmosphere. Some of this scattering is directed back toward and subsequently received by the speaker array (for a monostatic sodar).

Once the pulse is transmitted, there is a short delay before the speakers switch to receiving mode. The reason for the delay is that the speakers continue oscillating for a short time after pulse transmission. The use of a transmit-to-receive delay allows for delineation between residual speaker oscillations and atmospheric echoes. The length of the transmit-to-receive delay is the driver of the sodar minimum range. A shorter transmit-to-receive delay time clearly allows for a reduced minimum range, which is desirable for studying the lower ABL. Two ways to reduce the transmit-to-receive delay, through the use of piezo-transducers, or the incorporation of the use of different transmit frequencies for each beam, are discussed in sections 3.3.1.1 and 3.3.1.2.

3.4.1 Sampling time

The sampling time is determined from the maximum desired range at which sodar echoes are received. The sampling time must be shorter than the delay between pulses, with the transmit-to-receive delay removed, to ensure that there is no overlapping received signal from consecutive pulses. In practical terms, the maximum desired range is chosen by the user, and the sodar software translates this into a sampling time which encompasses all the range gates, based upon the speed of sound c calculated from the measured atmospheric temperature.

3.4.2 Range gate determination

Once the returned signal from one pulse is received, the signal is divided into batches which represent the return from different height layers.

Many commercial sodars allows the user to set their desired range resolution, Δr. The range resolution depends on both the pulse duration (as in equation 3.12), the FFT length, N_{FFT}, and the sampling frequency, f_x, as

$$\Delta r = \frac{cN_{FFT}}{2f_x}. \tag{3.16}$$

In many commercial sodars, the pulse duration to be used is calculated from the desired range resolution by rearranging equation 3.12 as $\tau = 2\Delta r/c$. However, the speed of sound, c is a function of the virtual temperature. For a sodar system operating without RASS, the speed of sound is calculated based upon the temperature measured by a sensor within the sodar housing. This speed of sound evaluated at the instrument height is used in determination of the range gates throughout the entire range encompassed by the sodar.

The calculated speed of sound c is used to determine the time duration assigned to a particular height range. This will only assign correct range values under isothermal conditions. If the temperature change with height is essential, such as under conditions of strong radiative cooling/heating or within a strong inversion, the reported measurement heights may be off by several meters (Fig. 3.5). For general atmospheric monitoring, this small difference between the reported and actual measurement height may not be significant, but it could be important for cloud base monitoring, for which sodars are sometimes used for in aviation, for example.

3.5 Signal processing

The sodar returned power, defined in equation 3.10, relates only to the amplitude of the received signal and does not provide any phase information. However, the full signal received by the sodar is of the form

$$V(t) = I(t) + jQ(t) \tag{3.17}$$

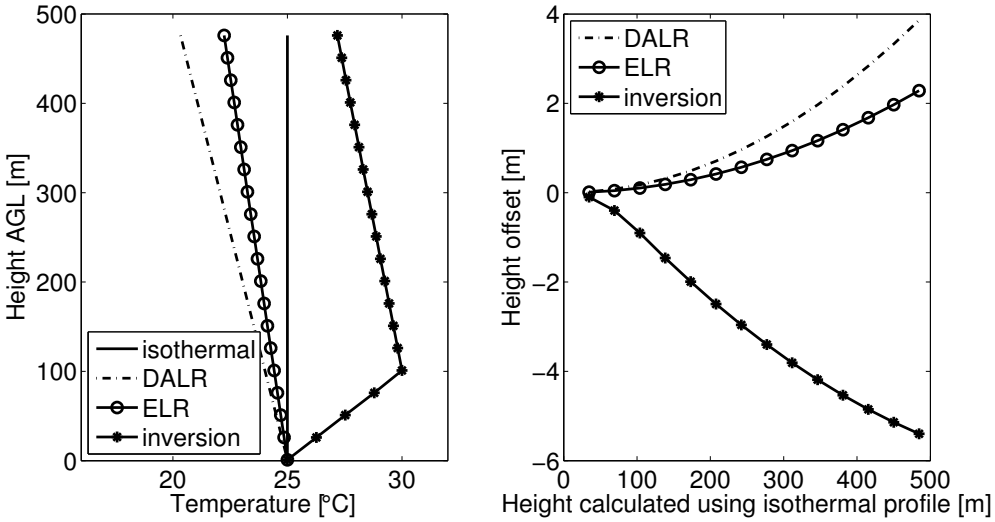

Figure 3.5: Left panel: four example vertical temperature profiles: isothermal, dry adiabatic, standard-atmosphere and a low level inversion. Right panel: offset in actual height of range gates for the different temperature profiles if the isothermal profile is assumed.

where $I(t)$ and $Q(t)$ are the in-phase and quadrature-phase portions of the signal and j denotes $\sqrt{-1}$. These are related to the amplitude, phase and Doppler shift of the received signal as

$$I(t) = \sqrt{P_R}\cos(2\pi\Delta f t + \varphi) \tag{3.18}$$

and

$$Q(t) = \sqrt{P_R}\sin(2\pi\Delta f t + \varphi). \tag{3.19}$$

This allows the full received signal to be written in the form

$$I(t) + jQ(t) = \sqrt{P_R}e^{(j2\pi\Delta f t + \varphi)}. \tag{3.20}$$

The signal is received in the form of a complex time series, defined by equations 3.17 and 3.20. The complex voltages are processed in data chunks corresponding to each range gate. A Fourier transform of the data is taken, and the resulting frequency spectra are processed to determine the frequency shift, and thus radial velocity. More details on each of these processes are included in the following sections.

3.5.1 Fourier transform

The received complex time series is split into chunks of a finite length for the purposes of performing a Fourier transform, which converts the signal from the time domain into the frequency domain. Sodars typically implement the Fourier transform as a Fast Fourier Transform (FFT, developed by Cooley and Tukey 1965), since the FFT algorithm provides a computationally efficient method of performing a discrete Fourier transform. The FFT is performed on a group of samples which defines a single range gate. Typically, the FFT is performed on samples of size 2^n, where n is a finite positive integer. The sample size of 2^n is chosen because the FFT is most computationally efficient when performed upon sample sizes that are a power of 2.

The use of the FFT algorithm to perform spectral estimation is widespread; it is used in all commercially available sodars. However, the FFT does have some limitations, which are discussed in detail Kay and Marple (1981). The main limitations of the FFT as it relates to sodar signal analysis are the finite and limited frequency resolution, and spectral smearing or leakage due to the windowing of the signal that occurs during the FFT. The limited spectral resolution hampers precise estimation of the spectral peak (see equation 3.22). Increasing the length of the FFT for the same sampling frequency increases the spectral resolution, although no new information is added: rather, the spectrum appears more smooth. The spectral leakage due to windowing in the FFT manifests itself as leakage of the signal from the main spectral lobe into the sidelobes, which can distort secondary spectral peaks. Both of these problems are most marked for very short data samples, i.e., when the FFT length greatly exceeds the number of samples contained within the portion of the signal to be analyzed.

Due to the mentioned limitations of the FFT, several other signal processing techniques have been suggested and tested. The main alternative technique is wavelet analysis (Morlet 1983), which has been used to identify fixed echoes in sodar returns

(Kalogiros and Helmis 1999), to identify coherent structure properties in the boundary layer (e.g., Coulter and Li 1995; Petenko and Bezverkhnii 1999) and to examine intermittent turbulence (e.g., Hadelberg and Gamage 1994). A hybrid between Fourier spectrum analysis and Wavelet analysis methods, known as the spavelet analysis was suggested by Petenko (2001). Despite improvements seen over the Fourier spectral analysis using these techniques, the FFT and subsequent spectral analysis remains the industry standard for sodars, and so the FFT is the only method considered further in this present study.

The choice of the FFT length has an effect upon the maximum vertical resolution, which can be calculated as

$$\Delta z_v = \frac{c}{2}\frac{N_{FFT}}{f_s} \tag{3.21}$$

where Δz_v is measured in meters, N_{FFT} is the length of the FFT, and f_s is the sampling frequency in Hz. The effect of the maximum vertical resolution due to FFT length is that for a particular height of interest, that range gate assigned to that height will contain contributions from $\pm\Delta z_v/2$ m above and below that height.

A typical sodar sampling frequency is around 960 Hz, and the FFT performed is usually of 64 or 128 points. For a 64-point FFT, this results in a spectral resolution of 15 Hz. The radial velocity resolution can be calculated from the spectral resolution as

$$\Delta v_r = \frac{c\Delta f}{2f_T}. \tag{3.22}$$

The aforementioned spectral resolution of 15 Hz will result in a radial velocity resolution of 1.42 m s^{-1} for a transmit frequency of 1800 Hz, and a radial velocity resolution of 0.85 m s^{-1} for a transmit frequency of 3000 Hz (assuming a speed of sound of 340 m s^{-1}, which corresponds to an air temperature of 287.5 K). For the 128-point FFT, the corresponding spectral resolution is 7.5 m s^{-1}. This gives a radial velocity resolution of 0.71 m s^{-1} for an 1800 Hz transmit frequency, and 0.43 m s^{-1} for a 3000 Hz transmit frequency.

Transmit frequency f_T (Hz)	Sampling frequency f_x (Hz)	Maximum frequency shift f_s (Hz)	Maximum radial velocity v_r (m s^{-1})	Zenith angle θ (°)	Maximum horizontal wind u_h (m s^{-1})
1800	340	±170	16.0	17	54.7
3000	340	±170	9.6	17	32.8
4500	960	±480	18.1	17	61.9

Table 3.1: Typical sodar sampling parameters, and the resulting maximum retrievable horizontal wind speeds. All calculations were performed assuming c=340 m s^{-1}.

Alternatively, some sodar models, use a very high sampling frequency, and a much higher-point FFT. The Metek sodar uses a sampling frequency of 44100 Hz, and a 4096-point FFT (as listed in Bradley 2008a). This results in a spectral resolution of 10.8 Hz. For sodars using very high sampling frequencies like these, only the portion of the resulting spectrum which surrounds the transmit frequency is used to determine the spectral peak. This is done for two reasons: firstly, to ensure that high noise levels at frequencies far away from the transmit frequency are not identified as acoustic echoes; and secondly, to reduce computational expense.

As can be seen from the previous two examples, both the high sampling frequency and high point FFT and the lower sampling frequency with lower point FFT provide a similar spectral resolution, and thus radial velocity resolution. For sodars operating with a high sampling frequency, typically only a smaller portion of the spectra resulting from the FFT (the portion surrounding the transmit frequency) is presented to the user. From a user standpoint, there is little difference between the two methods.

The maximum retrievable frequency shift (and thus the maximum retrievable radial velocity) also depends upon the sampling frequency. For a sampling frequency of f_x, this results in a maximum retrievable frequency shift of $\pm f_x/2$. Some examples of typical values for these sodar parameters are illustrated in Table 3.1.

3.5.2 Acoustic background noise

In order to retrieve the radial velocities from the frequency spectra, the spectral peak must be determined. As part of the signal processing technique, efforts are made to remove acoustic background noise from the spectra. This ensures that when the peak is detected, the spectral information is representative of the atmospheric acoustic echoes, with as little contamination by background noise as possible.

One common way that noise removal is performed involves the sodar increasing the time delay between pulse transmission for two of the beams. This allows the sodar to gather signal returns from a pre-specified noise height. The noise height is chosen to be well above the maximum retrievable range of the sodar, such that there should be no actual atmospheric echo from that height. A noise signal is thus recorded once during every new set of measurements (for typical conditions, this equates to a noise measurement every 10-30 s). The noise signal is processed in the same manner as the atmospheric signals, such that a FFT is performed, providing a resultant noise spectrum. This noise spectrum can then be subtracted from the recorded signal spectrum, resulting in an effectively de-noised signal spectrum.

Crescenti (1998) conducted a review of the impact of background noise on sodar performance, including a discussion of the effectiveness of sodar baffles at masking ambient noise. The sodar noise can be classified into broad-band and narrow-band, where narrow-band noise is fixed-frequency and generally results from fixed echoes or site-specific noise sources. The noise sources can also be classified as active or passive, where active noise sources include traffic, birds, and insects, while passive noise sources are fixed echoes such as buildings or trees, which reflect the sodar pulse. Acoustic background noise was revisited by Bradley (2012a) over a frequency range of 1000 - 6000 Hz, with the finding that the broad-band noise could be accurately represented by 'pink' noise, which follows a $1/f$ amplitude fall-off with frequency. The narrow-band noise measured by Bradley (2012a) was found again to be highly

site-specific, as well as intermittent. For this reason, the ambient noise at sodar sites should be examined prior to deployment to identify potential noise sources.

An example of a full day of noise spectra recorded by a sodar at a rural site in Oklahoma during the winter is shown in Fig. 3.6. It should be noted that the background noise is highly dependent upon the specific sodar location (e.g., rural or urban site, presence of nearby traffic noise, noise due to birds and insects) and also highly temporally variable at a single site. Figure 3.6 shows that at this rural site, the noise level increased at all frequencies during the period 0030-0130 UTC, which corresponds to the hour preceding and during sunset. As this is a rural site with a mix of prairie grass and forested vegetation types, it is likely that the increase in the noise level around sunset is due to an increase in the noise from birds and insects. It can also be seen in Fig. 3.6 that there is a noise source at frequencies close to both the upper and lower frequency limits. Both of these noise sources are persistent throughout the day, which is suggestive of electronic or instrument-based noise. Given the time of year, and the temperature during this day (a high of -1.7°C and a low of -9.4°C) a possible cause for this noise could be the electronic heating system in the sodar, or the heating system in a nearby trailer (distance of approximately 15 m from the sodar). There is a slight increase in the background noise at frequencies between 1870-1890 Hz, which is close to the transmitted frequency of 1895 Hz. This would seem to indicate that the noise spectra in Fig. 3.6 include contributions from actual atmospheric echoes, suggesting that the height from which the noise measurements are taken should be increased. The additional increase in the background noise level from approximately 2100-2230 UTC is caused by falling snow, which was measured by a Mesonet station located at the field site. Just by examining the noise spectra recorded during one day, it is clear that there are many variables which can affect the acoustic background noise.

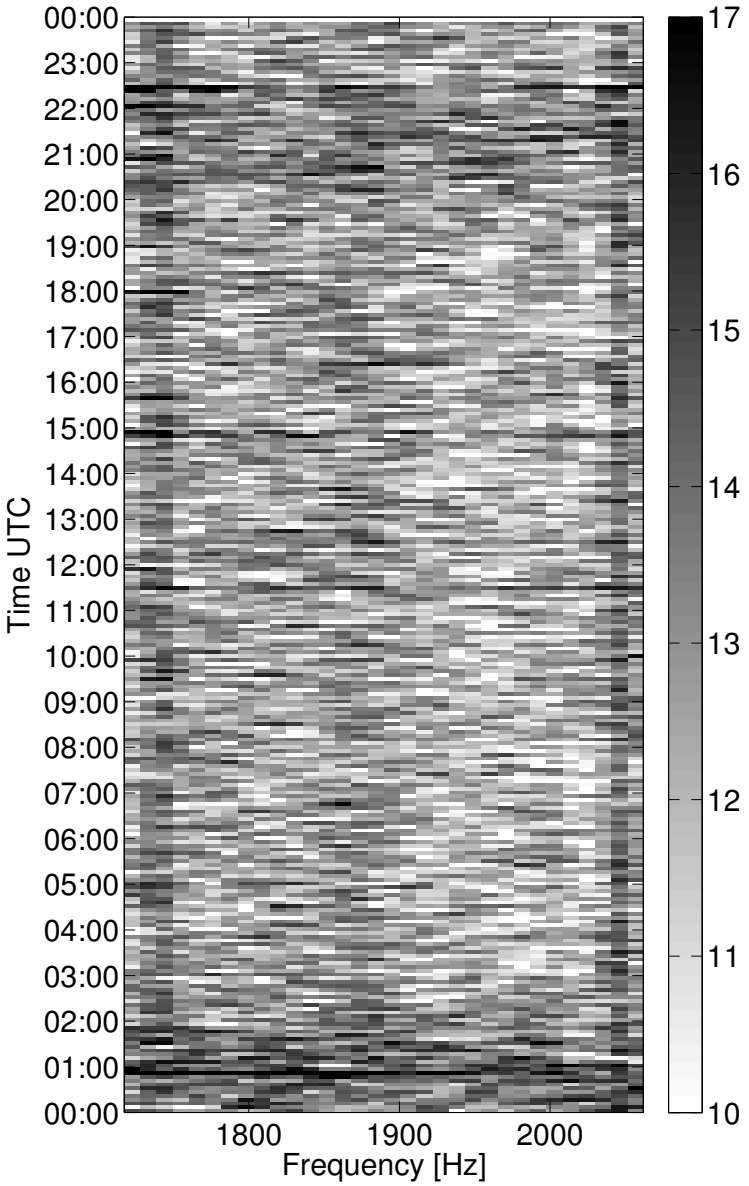

Figure 3.6: Example of the recorded sodar noise spectra over a 24-hour period during 7th-8th February 2014 at a rural field site near Purcell, Oklahoma, USA. The power is shaded in gray. The sodar transmit frequency was 1895 Hz. The spectra are recorded at approximately 15 s intervals are averaged over 5-minute periods

An alternative method of noise estimation to recording a distinct noise spectrum is to operate the sodar using a slightly higher sampling frequency, and then use the tails of the frequency spectrum to estimate the background noise. This relies on the assumption that the tails of the spectrum do not contain any acoustic echoes, which may be invalid in strong winds.

3.5.3 Peak detection

The next step in the sodar signal processing is calculation of the radial velocity from the frequency spectra produced by the FFT. In order to do this, the peak of the frequency spectrum must be determined.

Most commercially available sodars allow the user to define an averaging period. This is the length of time over which individual spectra are summed prior to peak detection and radial velocity estimation. A longer averaging period helps to increase the signal-to-noise ratio, as acoustic noise is often highly transient. If the noise is assumed to be random, then the noise amplitude is assumed to decrease as a factor of the square root of the number of additional spectra within the averaging period. An averaging period of at least five minutes is typically recommended for accurate wind retrievals.

As discussed in section 3.5.1, the spectral resolution of typical commercially available sodars is usually in the range of 5-15 Hz, with a resulting radial velocity resolution of 0.5-1 m s^{-1}. Different sodar models will use different specific algorithms to detect the spectral peak (and resulting radial velocity) to a higher accuracy than the spectral resolution. The method described in the following paragraph is used in the Metek PCS.2000 sodar, which is operated at the University of Oklahoma.

The Metek PCS.2000 uses a sampling frequency of 44100 Hz, and a 4096-point FFT, providing a spectral resolution of 10.8 Hz. Once the FFT is performed, only

the 32 spectral points surrounding the transmit frequency are kept (f_T = point 17 of 32), and the rest of the spectral information is discarded.

All the 32-point frequency spectra contained within the averaging period are summed. The corresponding noise spectra from within the averaging period are also summed, and the summed noise spectrum is subtracted from the summed signal spectrum, essentially producing a de-noised frequency spectrum. Peak detection is then performed on this average de-noised spectrum. An initial peak detection simply identifies the point on the spectra with the highest power. This peak identification results in a frequency shift which is a multiple of the spectral resolution. A more refined technique is then implemented to ensure that frequency shifts are not limited to intervals of multiples of the spectral resolution. The Metek sodar performs this by choosing a set of spectral points centered around the identified peak. The set of points is chosen to include consecutive points around the peak which have a power value of at least 25% of the power at the peak, in linear units. This equates to discarding all the spectral points with a power decrease of more than 6 dB from the power at the peak. The use of only the points surrounding the peak in the fitting procedure is done to ensure that the reported wind speed represents the true spectral peak due to the returned acoustic signal and not due to noise. A Gaussian fit is performed upon this set. The fitting of the Gaussian allows the spectral peak to shift away from multiples of the spectral resolution, providing a more accurate estimation of the spectral peak. This fitting technique is illustrated in Fig. 3.7.

3.5.4 Wind retrievals

Once the spectral peak has been identified, this is translated to a frequency shift from the transmit frequency, labeled f_s. The radial velocity, v_r, for each range gate can be retrieved from this frequency shift as

$$v_r = \frac{f_s c}{2 f_T},$$

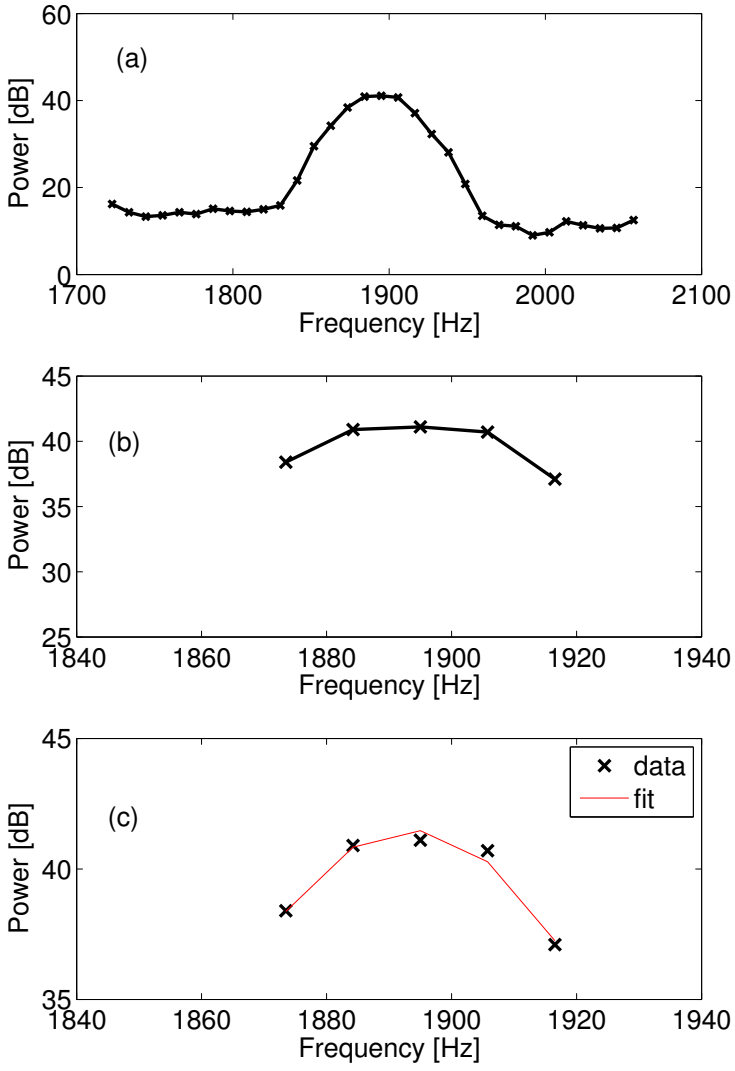

Figure 3.7: Illustration of the spectral processing and peak fitting technique used in a Metek sodar Panel (a) shows a recorded frequency spectrum from a transmit frequency of 1895 Hz. Panel (b) shows only the points with at least 25% of the peak power. Panel (c) illustrates a fit to the points in panel (b).

where the sodar estimates c using the air temperature measured at the instrument height.

If radial velocities are measured by three independent and non-coplanar sodar beams, then the three wind components, u, v, and w, can be retrieved. For typical sodars, this is usually achieved through the use of a vertical beam, and two or four mutually orthogonal beams (i.e., the off-vertical beams are spaced at 90° azimuth intervals).

The radial velocities are related to the wind components as

$$v_r(\theta, \phi) = u \sin\theta \sin\phi + v \sin\theta \cos\phi + w \cos\theta, \tag{3.23}$$

which for a five-beam retrieval can be written as the system $\mathbf{Cu} = \mathbf{v_r}$ where

$$\mathbf{C} = \begin{bmatrix} \sin\theta_1 \sin\phi_1 & \sin\theta_1 \cos\phi_1 & \cos\theta_1 \\ \sin\theta_2 \sin\phi_2 & \sin\theta_2 \cos\phi_2 & \cos\theta_2 \\ \sin\theta_3 \sin\phi_3 & \sin\theta_3 \cos\phi_3 & \cos\theta_3 \\ \sin\theta_4 \sin\phi_4 & \sin\theta_4 \cos\phi_4 & \cos\theta_4 \\ \sin\theta_5 \sin\phi_5 & \sin\theta_5 \cos\phi_5 & \cos\theta_5 \end{bmatrix}, \mathbf{u} = \begin{bmatrix} u \\ v \\ w \end{bmatrix} \text{ and } \mathbf{v_r} = \begin{bmatrix} v_{r1} \\ v_{r2} \\ v_{r3} \\ v_{r4} \\ v_{r5} \end{bmatrix}$$

with ϕ_i and θ_i being the azimuth and zenith angle for beam i, respectively, (Palmer et al. 1993). Once the calculated radial velocities are known, the wind components can be derived using the least squares solution of the matrix equation

$$\mathbf{u} = (\mathbf{C^T C})^{-1} \mathbf{C^T v_r} \tag{3.24}$$

for the five-beam system.

The equivalent retrieval for a three-beam system is also written as $\mathbf{Cu} = \mathbf{v_r}$ where

$$\mathbf{C} = \begin{bmatrix} \sin\theta_1 \sin\phi_1 & \sin\theta_1 \cos\phi_1 & \cos\theta_1 \\ \sin\theta_2 \sin\phi_2 & \sin\theta_2 \cos\phi_2 & \cos\theta_2 \\ \sin\theta_3 \sin\phi_3 & \sin\theta_3 \cos\phi_3 & \cos\theta_3 \end{bmatrix}, \mathbf{u} = \begin{bmatrix} u \\ v \\ w \end{bmatrix} \text{ and } \mathbf{v_r} = \begin{bmatrix} v_{r1} \\ v_{r2} \\ v_{r3} \end{bmatrix}.$$

This system can be simply solved as

$$\mathbf{u} = \mathbf{C}^{-1} \mathbf{v_r}. \tag{3.25}$$

Using this wind retrieval method, other useful parameters can be evaluated. By performing the wind retrieval method on individual spectra from the vertical beam, rather than the spectrum resulting from the summing across the averaging period, the variance of the vertical velocity can be evaluated. This parameter, written as σ_w, is useful for comparing changes in temporal or vertical variation of w, and is often used to quantify turbulence. The variance of the zonal and meridional winds, σ_u and σ_v can be evaluated in the same manner.

3.5.5 Power return

Aside from the three-dimensional wind components, the returned acoustic power (defined in equation 3.10) is the other most widely used sodar output. There are three specific sodar parameters parameters which are typically used to examine the turbulence profile and the thermodynamic structure of the atmosphere. These are the direct power return, the signal-to-noise ratio (SNR), and the reflectivity.

For the Metek PCS.2000 sodar, the output power return is evaluated as the maximum power value at the peak in the frequency spectra, over all spectra within the averaging period. The SNR, when evaluated for instantaneous data, is a ratio of the power at the peak of the frequency spectra resulting from the signal return to the maximum power in the noise spectra. A single value for the peak of the noise spectra is used to determine the SNR at all range gates, and for each of the n beams within a single scan. When evaluated for data with a finite averaging period, the SNR is calculated as a weighted ratio of the maximum power in the average de-noised spectra and the maximum variability in the measured noise spectra. The reported reflectivity value is the maximum value of the Gaussian curve which was fitted to the subsection of the average frequency spectra surrounding the peak.

A great deal of information about the state of the atmosphere can be inferred from the sodar power return. The power return can be used to identify the height of the

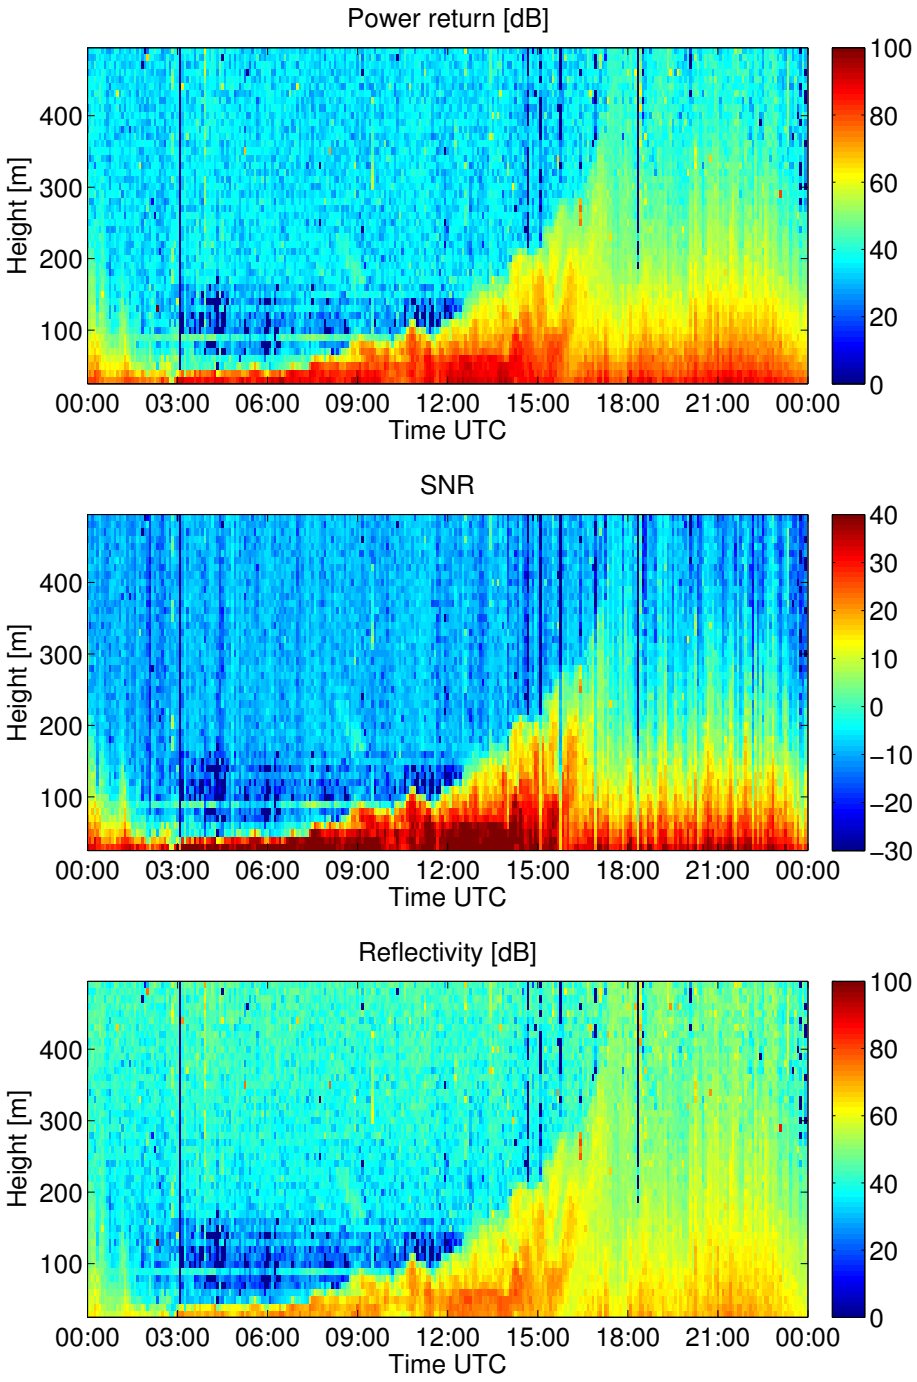

Figure 3.8: Example of the acoustic return from a Metek PCS.2000 sodar located at the Kessler atmospheric and ecological field station (KAEFS) in rural Oklahoma, on February 18th 2014. The upper panel shows the power return in dB, the central panel shows the signal-to-noise ratio, and the lower panel shows the reflectivity in dB. A five minute averaging period was used.

mixing-layer top, as discussed in section 2.3.2.1. It can also be used to give an idea idea about the atmospheric stability. For example, a sharp and sustained decrease of power or reflectivity with height is indicative of a stably stratified boundary layer.

The variation in power return, reflectivity, and SNR over a 24-hour period is shown in Fig. 3.8. The sharp decrease in power and reflectivity at around 40 m height during 0300-1300 UTC (2100-0700 LST) indicates stable stratification during the overnight period, as is expected from the traditional boundary layer diurnal cycle as illustrated in Fig. 2.1. By 1600 UTC (1000 LST), the decrease of reflectivity with height is much smoother, and the transition to the well-mixed daytime convective boundary layer is complete. It is also noticeable in Fig. 3.8 that the returned power in the lower portion of the boundary layer is roughly 20 dB higher than the reflectivity at that height. This is due to the calculation method for each of these variables: as mentioned previously, the power return is the maximum value in any of the instantaneous de-noised frequency spectra contained within the averaging period, whilst the reflectivity is the peak of a Gaussian curve fitted to the *mean* de-noised spectrum for the averaging period. This means that the reflectivity field is typically smoother then the power field, since the fitting of the Gaussian curve to the averaged de-noised spectra acts to mitigate the effect of noise on temporal scales shorter than the averaging period. The return power field, on the other hand, may be contaminated by noise sources that occur on shorter timescales than the averaging period (such as bird calls). Due to the smoothing effect of the application of the Gaussian fit to the reflectivity data, the examination of atmospheric phenomena which occur on short time-scales is best performed using the return power field. For example, an examination of the average wind speed can be performed using the reflectivity data, but this could not be used to examine the maximum wind gust speed recorded by the sodar (on the $\tilde{1}5$ s timescale) due to the Gaussian fitting procedure.

The power return field can also be used to identify elevated inversion layers. An example of a power return showing an elevated inversion can be seen in Fig. 3.9. The power return was recorded by a Metek PCS.2000 sodar operating at the Kessler atmospheric and ecological field station (KAEFS) near Purcell, Oklahoma, USA. An elevated inversion, which manifests itself as a layer of increased acoustic return power above a layer of lower power return, can clearly be seen above a height of 300 m. The elevated inversion persists for many hours, as it can be seen in Fig. 3.9 that the elevated inversion is already in place by 0700 UTC (0100 LST), and it is seen to persist until at least 1800 UTC (1200 LST), at which time the height of the elevated inversion rises above the maximum retrievable range of the sodar, as the depth of the well-mixed layer increases through the morning.

The power return can also be used to examine the onset time of the evening transition regime. The sodar returned power changes from a convective structure to a stratified and less time-varying structure as the heat flux decreases. This is marked by a decrease in the returned power, initiating from higher levels and then descending as the boundary layer collapses with the removal of solar heating. A relative minima in the sodar returned power was one of the parameters used alongside various surface meteorological measurements in Busse and Knupp (2012) to evaluate the timing of the evening transition for several low-wind, clear-sky case studies in Alabama.

In this chapter, the principles underlying acoustic remote sensing have been presented. Issues pertaining to different aspects of sodar hardware and software have been considered, with a focus on how each of the sodar operating parameters is affected by the hardware and software choices. The sodar operating parameters which can typically be set the user are reviewed, with a focus on how the choice of operating parameters can affect the sodar output parameters. In the following chapter, the algorithms used to power the sodar simulator are presented.

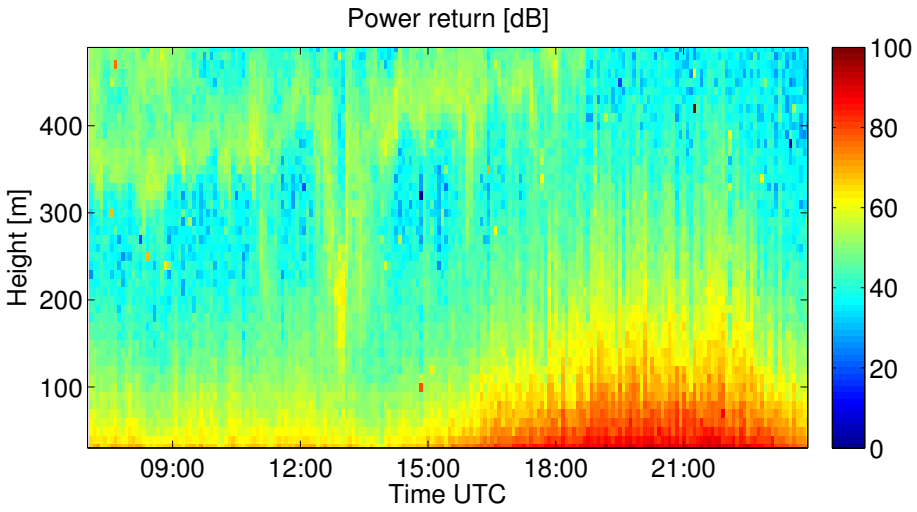

Figure 3.9: Example of an elevated inversion signature from a Metek PCS.2000 sodar located at the KAEFS site in rural Oklahoma, on February 1st 2014. A five minute averaging period was used.

Chapter 4

Simulation theory

4.1 University of Oklahoma Large Eddy Simulation

The University of Oklahoma Large Eddy Simulation (OU-LES) model is used in this study to provide the three-dimensional turbulent flow fields ingested by the sodar simulator. The governing equations used in the OU-LES model are derived from the general form of the Nävier-Stokes equations, and are outlined in Conzemius (2004).

The OU-LES model has been tested extensively under conditions representative of the convective boundary layer (CBL). Comparisons between experimental and numerical CBL data have shown that the OU-LES can realistically produce mean flow and turbulence structure under clear CBL conditions (Fedorovich et al. 2001, Fedorovich et al. 2004a, Botnick and Fedorovich 2008). The OU-LES has previously been tested within the framework of a boundary layer radar simulator (Scipión et al. 2008), which allowed for the testing of different scanning strategies and signal processing techniques in much the same way as the present study does for acoustic remote sensors.

Details about the model setup for each of the case studies are included in Chapter 6.

4.2 Sodar simulator

As detailed in Section 2.2, a wide variety of meteorological instrumentation is available for studying the atmospheric boundary layer. Clearly, different observing and remote sensing platforms are better suited for certain observational needs. The choice of a remote sensing or an *in-situ* observational instrument for a specific project depends upon a number of factors, including but not limited to: the variables of interest (e.g.,

wind speed and direction, flux measurements, mixing layer height, etc.), the desired spatial coverage and height range of interest, the temporal resolution of the retrieved parameters, and practical considerations such as instrument availability, maintenance, and operational costs. The retrieval properties of each instrument can easily be examined: the recorded parameters, the maximum retrievable range, the areal coverage, the averaging period, etc., are all known and thus available for comparison.

However, it is difficult to make a reasonable comparison of retrievals from different instruments for the same atmospheric state, as there are logistic difficulties in collocating several different remote sensing instruments. This difficulty is primarily due to the interference caused by one instrument upon another. For example, if a comparison is to be made between wind-components measured by anemometers mounted on a mast, and a radar or sodar, the mast and radar or sodar must have a separation distance short enough that the winds they are sampling can be considered homogeneous. However, an instrumented mast would act as a source of ground clutter for a closely co-located radar, or it would act as an echo source causing acoustic reflections of sodar beams.

The difficulty in overlapping different remote sensors in the field, and the impossibility of scanning the same atmospheric state repeatedly using different scanning techniques on the same instrument has led to the development of the technique of simulating remote sensing instruments. This technique involves the use of a modeled atmospheric state. The modeled flow fields are ingested into an instrument simulator, which mimics the remote sensing processes and signal processing techniques of the instrument in question, producing output fields that mimic as closely as possible what the remote sensing instrument would produce under the ingested modeled atmospheric conditions.

Simulating the scanning of atmospheric fields by remote sensing instruments is a technique that has been used with much success in the radar meteorology community.

Radar simulators, which can scan the output from high-resolution atmospheric models and large-eddy simulations, have been developed for weather radars (e.g., May et al. 2007; Cheong et al. 2008), millimeter-wavelength cloud radars (Clothiaux et al. 1996), boundary layer wind profiling radars (e.g., Muschinski et al. 1999; Scipión et al. 2009), three-dimensional imaging radars (e.g., Yu and Palmer 2001; Cheong et al. 2004), and radars with dual-polarimetric scanning capabilities (e.g., Capsoni et al. 2001; Jung et al. 2008). Scattering of sodar signals has previously been modeled using a one-dimensional slab-based approach by Bradley (1999). The technique of simulating the scanning of full three-dimensional atmospheric model output has only recently been extended to acoustic remote sensing as part of the presented work (see also Wainwright et al. 2014b).

The benefits of simulating the sodar signal based upon three-dimensional model output are manifold. The same model output fields can be repeatedly scanned using different combinations of scanning parameter values, and differences between the results can be fairly compared. The model output also acts as a 'ground truth' for the simulated sodar output, meaning that inaccuracies in the sodar output due to scanning techniques can be easily quantified. Novel scanning parameter combinations or signal processing techniques can be tested and evaluated on realistic atmospheric flow fields. Using model output as the basis for simulating the sodar signal also enables the simulation of sodar signals for locations in which the use of a sodar may be impractical, such as in remote locations where powering a sodar may be an issue, or in suburban areas where a nocturnal sodar pulse may be considered a noise nuisance. The final benefit of sodar simulation is that it allows for physically unrealistic close collocation of multiple sodars, or a sodar and an additional simulated meteorological instrument, with no adverse effects due to interference between instruments.

The following sections within this chapter describe in detail the algorithms used within the sodar simulator. A general overview of the simulator algorithms is provided

in Section 4.2.1, and then each algorithm is described in further detail in the sections that follow.

4.2.1 Overview of the simulator algorithms

The sodar simulator ingests a succession of three-dimensional flow fields containing three resolved (in the LES sense) velocity components: u, v, and w, the resolved potential temperature of air Θ, and the resolved water vapor mixing ratio q from an LES. The simulator also ingests user-defined values for certain variables, detailed below. The location of each sodar beam is then calculated based upon the temperature-varying speed of sound (Sections 4.2.2, 4.2.3). Following this, the radial velocities and a proxy for the acoustic power are calculated at the beam locations. Weighting functions are applied to the radial velocity and power (Section 4.2.4). The atmospheric absorption and acoustic backscattering cross-section are next calculated at the beam locations (Sections 4.2.5, 4.2.6). At this stage all the components are in place to calculate the phase and amplitude of the received complex acoustic signal.

For each beam, the complex acoustic signal is coherently summed across a number of beam points b within the beam to provide a time-series in range. The beams are divided into range gates of the appropriate length and time-series acoustic background noise is generated for each gate and added to the complex acoustic signal (Section 4.2.7). A fast Fourier transform (FFT) is performed on the complex noisy signal contained in that range gate. To process the synthesized signals, an average noise spectrum is generated for each beam and used to de-noise the spectra at each range gate. The de-noised Doppler spectra are finally processed to derive the radial velocity at each range gate and beam (Section 4.2.8). This process is repeated for the number of independent beams (n_b) required, and then the beam sequence is repeated until the desired runtime of the simulator has elapsed. Estimation of the three-dimensional wind field is performed on each beam sequence. Temporal averaging of

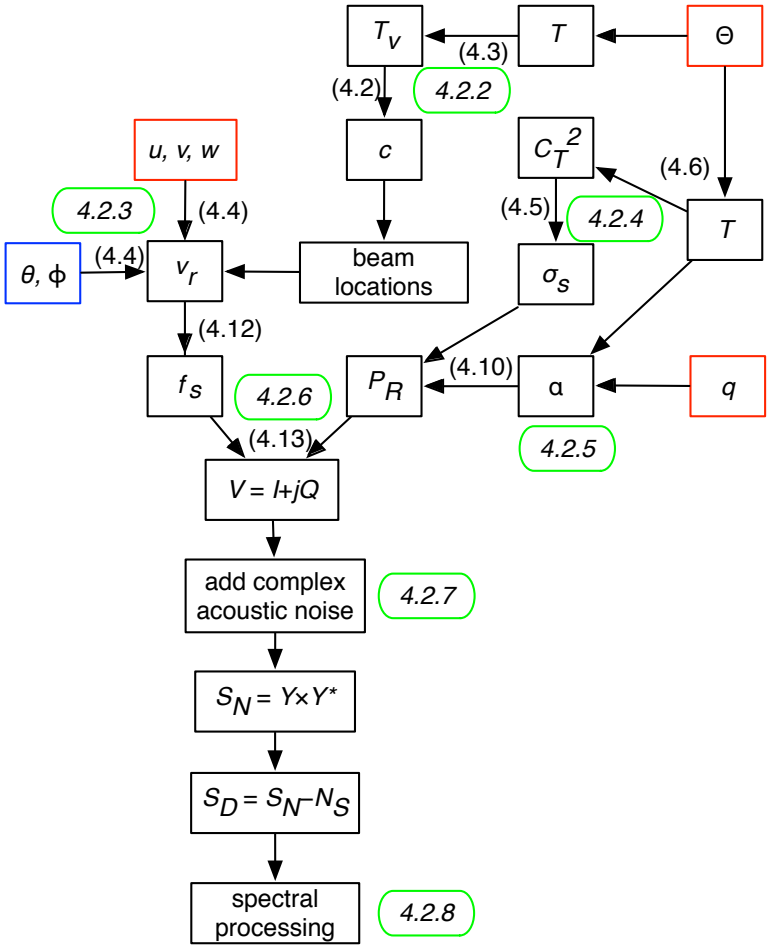

Figure 4.1: An overview of the workflow of the time series sodar simulator. The equations corresponding to the processes referred to are in parentheses. The section numbers covering each process are outlined in green. The red boxes indicate variables provided by the LES output, and the blue box indicates parameters defined by the user.

the spectra to gain time-averaged wind field estimates is applied as post-processing (Section 4.2.9).

An overview of the workflow of the time series sodar simulator is presented in Fig. 4.1. Equations corresponding to each process are denoted in parentheses, and more detail on each of these processes is presented within the follow subsections.

4.2.2 Beam locations

One of the most useful points of the sodar simulator is that many of the characteristics of the simulated sodar can be defined by the user. This allows for the emulation

Parameter name	Parameter symbol	Measurement unit	Example value
Sodar location	$[x, y, z]$	m	[200, 200, 0]
Transmit frequency	f_T	Hz	2000
Number of beams	n_b	#	5
Azimuth angle	ϕ_i i=1,2,...,n	°	[0, 90, 0, 180, 270]
Zenith angle	θ_i i=1,2,...,n	°	[15, 15, 0, 15, 15]
Range resolution	Δr	m	10
Number of range gates	M	#	40

Table 4.1: Simulated sodar operating parameters which are defined by the user, and typical values for each parameter.

of different commercially available sodar systems. The sodar operating parameters which are defined by the user are also shown in Table 4.1.

User-defined characteristics include the placement of the sodar within the LES domain, and the transmitted frequency. The number of independent beam locations, and the beamwidth, azimuth, and zenith angle of each beam are also chosen by the user. The desired approximate range resolution of the simulated sodar is also adjustable, although it should be chosen such that is is larger than the vertical grid spacing of the LES whose data the simulator is ingesting.

Once a desired range resolution is provided, the pulse duration τ is then calculated as

$$\tau = \frac{2\Delta r}{c_{\text{ref}}}, \tag{4.1}$$

where Δr is the input range resolution, and c_{ref} is a reference speed of sound of 340 m s^{-1} which corresponds to a temperature of 288 K for dry air.

Acoustic waves do not travel at a fixed speed, rather their speed depends upon the temperature of the medium through which they are traveling. The speed of sound is far more sensitive to temperature variations than electromagnetic waves are; for the

same temperature variations, the changes in refractive index are 1000 times greater for acoustic waves than for electromagnetic waves (Kallistratova 1997). The speed of sound c varies with air temperature as

$$c = 20.05\sqrt{T_v}, \tag{4.2}$$

where T_v is the virtual temperature of the air in K, and c is the sound speed measured in m s^{-1} (May et al. 1990). The virtual temperature can be calculated from the LES output field of θ, ingested by the simulator, as

$$T_v = \frac{T}{1 - \frac{e}{p}(1 - 0.622)} \tag{4.3}$$

where p is the atmospheric pressure, and e is the vapor pressure, both measured in Pa, and T is the air temperature in K.

The temperature dependence of c causes the speed of sound to change with height under typical conditions. For a constant pulse length, this causes a corresponding slight variation in the sodar range resolution with height. This is important to note, because a sodar operating without radio acoustic sounding system (RASS) capability typically measures the temperature at the instrument height only.

In order to estimate the three-component wind vectors from the simulated sodar, at least three independent and non-coplanar beams must be employed. This is typically achieved using a vertically oriented beam and two off-vertical beams which are horizontally mutually perpendicular. The DBS technique requires the assumption of horizontal homogeneity in the wind fields across the area encompassed by the beams. This assumption is typically invalid, especially during convective conditions, and as a result, the wind components derived from the sodar do not represent the winds directly above the sodar, but include contributions from the sampling domain bounded by the off-vertical beams.

The number of calculation points contained within each beam is defined as $b \times b \times n \times M$, where b is the number of calculation points in each horizontal direction,

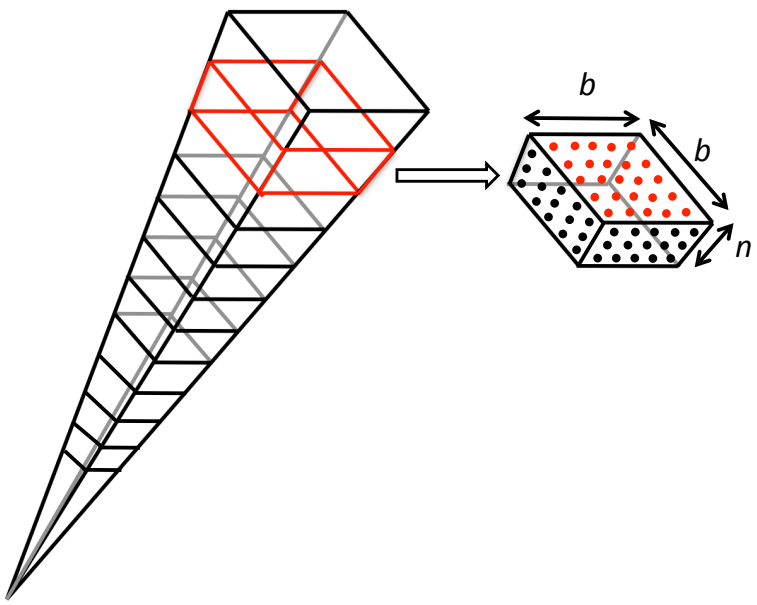

Figure 4.2: The calculation points contained within each beam. The single range gate outlined in red on the left panel is shown in detail on the right. Each range gate contains $b \times b \times n$ points at which calculations are performed. Each of n layers of beam points contains points that are equal in range from the sodar location. Thus each layer is convex in shape, while depicted as planar here for simplicity.

n is the number of layers of calculation points contained within one range gate, and M is the number of range gates within each beam. This feature is illustrated for one beam in Fig. 4.2. While each horizontal layer of beam points n is depicted as planar in Fig. 4.2 for simplicity, each point in the layer is equidistant from the sodar location, hence the layer is of convex shape. Appropriate values for b and n should be chosen with the LES resolution in mind, as this determines the maximum number of points available for interpolation.

For each of the $b \times b \times n \times M$ calculation points, the location of the calculation point relative to the LES domain must be calculated. This is performed using the input range resolution as an initial guide. The LES output temperature field is translated into a three-dimensional speed of sound field using equations (4.2) and (4.3). Using the value for τ calculated from equation (4.1), an appropriate transmit-to-receive delay, and the three-dimensional speed of sound field, the spatial distance from the simulated sodar can be calculated for the acoustic pulse. It should be noted

that the minimum range (i.e., the lowest retrieval height) of the sodar is calculated as a part of this process. The transmit-to-receive delay is the primary controller of the sodar minimum range. Currently, this value is set in the simulator to a realistic value of 0.02 s, although this can be altered by the user in order to emulate specific sodar models.

The calculation point range information is used in two ways. Firstly, the range of each of the n layers of calculation points is used, along with the beamwidth, azimuth and zenith angle of beam, and the number of points in the horizontal direction b, to calculate the distance of each and every calculation point from the simulated sodar location in spherical coordinates. The spherical coordinates are then translated into the cartesian coordinate system, and into x, y, and z values relative to the sodar location within the LES. The second way in which the range information is used is that the speed of sound field is interpolated to the location of each calculation point. The time taken for the pulse to travel between the n levels of calculation points can then be calculated. These times are subsequently summed to provide a cumulative time taken for the pulse to reach each level of calculation points. The cumulative travel times are utilized to provide a basis for the temporal interpolation between consecutive LES snapshots.

Finally, the calculation point locations in x, y, and z relative to the simulated location of the sodar are converted into indexes within the LES domain. This simplifies the step of trilinear spatial interpolation which must be performed to assign wind and temperature values to each calculation point.

4.2.3 Radial velocity

Once the spatial location of each beam point both relative to the sodar and within the LES domain has been determined, the radial velocities at the calculation points are then determined. The velocity components are first calculated at the beam points

by using trilinear spatial interpolation on the original fields of u, v, and w from the LES resolved on a $X \times Y \times Z$ numerical grid. The velocity components are also interpolated in time between concurrent LES snapshots, using the cumulative travel times calculated for each beam point as described in the preceding section, to allow for time-accurate spatial interpolation between range gates.

The radial velocity is determined from the spatially interpolated velocity components as

$$v_r(\theta, \phi) = u \sin\theta \sin\phi + v \sin\theta \cos\phi + w \cos\theta, \tag{4.4}$$

where ϕ and θ are the azimuth and zenith angles of each beam point, respectively.

4.2.4 Weighting functions

Two types of weighting functions are used in the simulator. The first type of function weights the radial velocity values based upon their location within the beam and within the range gate. The beam-weighting function employed is a multivariate normal distribution with a mean of the center of the beam in the x and y directions, and a standard deviation of the beam width. Points near boresight have the maximum value of 0 dB, with the value decreasing to -3 dB at specified beamwidth distance. The range-weighting function is based on an extended cosine pulse shape and also uses a maximum value of 0 dB at the center range of each beam and -3 dB at the extremities of each range gate. The beam- and range-weighting functions are combined to yield one single volumetric weighting function with values ranging from 0 dB to -6 dB. An example of the beam point weighting based upon beamwidth, range, and the combination of the two is shown in Fig. 4.3.

The second type of weighting function is based upon the backscattering intensity at each beam point. The use of this second weighting function is to ensure that the most acoustically reflective calculation points are more heavily weighted than the

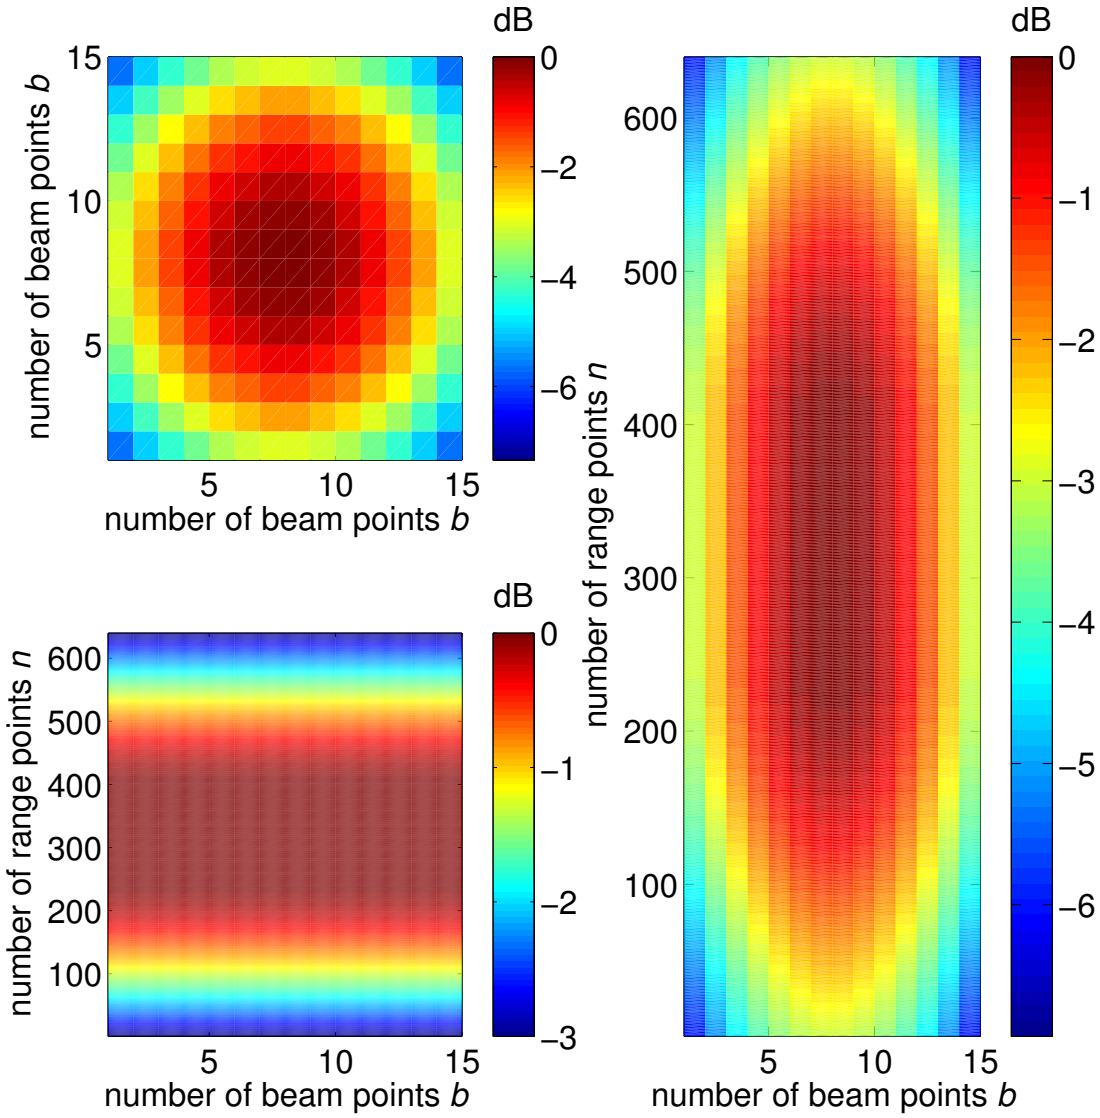

Figure 4.3: Upper left panel: an example of weighting due to beamwidth for b=15 horizontal calculation points. Lower left panel: an example of weighting due to range for b=15 and n=640. The panel on the right shows a slice of the full weighting matrix for one range gate, containing both beam- and range-weighting.

least acoustically reflective calculation points. Acoustic backscattering cross-section is calculated as

$$\sigma_s = 4 \times 10^{-3} \lambda^{-1/3} \frac{C_T^2}{T^2}, \tag{4.5}$$

(following Tatarskii 1971), where λ represents the transmit wavelength of the sodar. The structure function parameter for temperature is calculated using

$$C_T^2 = \frac{\langle [T(\mathbf{r} + \delta) - \mathrm{T}(\mathbf{r})]^2 \rangle}{|\delta|^{2/3}}, \tag{4.6}$$

where $\mathbf{r}$ is a position vector and δ is the separation vector, and $\langle \cdot \rangle$ indicates averaging across all the difference pairs. The use of (4.5) is justified at scattering lengths within the inertial subrange, where the turbulence is assumed to be isotropic and homogeneous (Tatarskii 1971).

The difference term is evaluated for all possible pairs of T values within a cube of equal size to the LES grid spacing, and then normalized by the distance δ, which here is equal to the linear size of the LES grid cell. The calculation of C_T^2 across all possible pairs within an LES grid cell is illustrated in Fig. 4.4. These pairs are then averaged to give one C_T^2 value for each grid box, which is interpolated to the beam point locations in both space and time. The second weighting function uses $\frac{C_T^2}{T^2}$ as a proxy for power. Weighting the radial velocity values with this power proxy ensures that the average radial velocity value for the range gate is weighted more heavily toward the most acoustically reflective points.

4.2.5 Atmospheric absorption

The intensity of the acoustic signal decreases as it travels through the atmosphere due to a number of factors including losses caused by spherical spreading, scattering of the acoustic signal outside of the beam volume, and atmospheric absorption. The power loss due to spherical spreading is proportional to $\frac{1}{r^2}$, and will be included in the calculation of the complex acoustic signal directly, as shown below in (4.10) and (4.13). Losses due to scattering out of the beam are neglected as they were shown

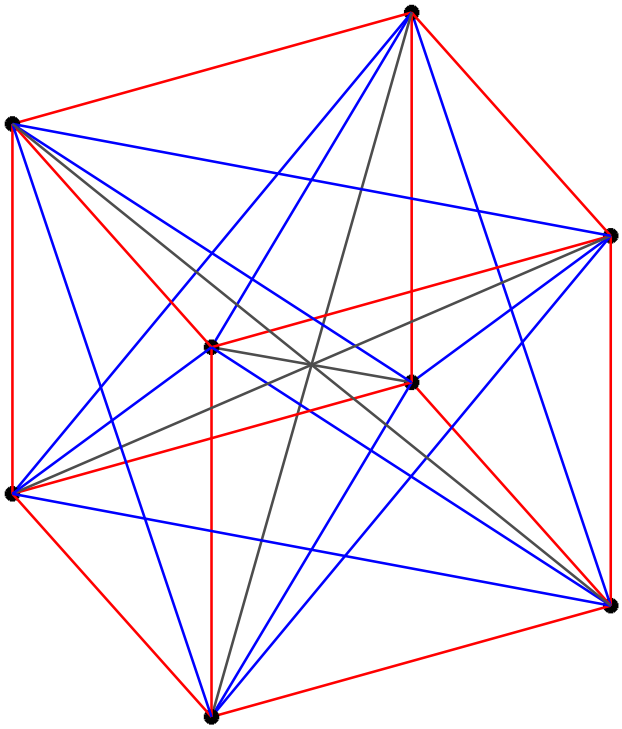

Figure 4.4: The calculation of all the possible temperature pairs across a single LES grid cell. The black dots represent the temperature measurements that bound the grid cell, and each line represent a specific pair of these measurements. The red lines show the edges of the cube, the blue lines show the outer diagonals, and the grey lines illustrate the inner diagonals. There are 28 lines per cube in all.

by Pan (2003) to be small in comparison to losses due to atmospheric attenuation for typical sodar operating frequencies. The loss due to atmospheric absorption, α, depends on the temperature, pressure, humidity, and transmitted frequency. We calculate α according to ISO9613-1 (1993).

The formula for α is as follows:

$$\alpha = 8.686 f^2 \left(\left[1.84 \times 10^{-11} \left(\frac{p_a}{p_r} \right)^{-1} \left(\frac{T}{T_0} \right)^{1/2} \right] + \left(\frac{T}{T_0} \right)^{-5/2} \times \left\{ 0.01275 \left[\exp\left(\frac{-2239.1}{T} \right) \right] \left[f_{rO} + \frac{f^2}{f_{rO}} \right]^{-1} + 0.1068 \left[\exp\left(\frac{-3352.0}{T} \right) \right] \left[f_{rN} + \frac{f^2}{f_{rN}} \right]^{-1} \right\} \right) \quad (4.7)$$

where α is measured in dB m^{-1}, and f_{r0}, the relaxation frequency of oxygen, and f_{rN}, the relaxation frequency of nitrogen, are calculated as

$$f_{rO} = \frac{p_a}{p_r} \left(24 + 4.04 \times 10^4 h \frac{0.02 + h}{0.391 + h} \right) \quad (4.8)$$

and

$$f_{rN} = \frac{p_a}{p_r} \left(\frac{T}{T_0} \right)^{-1/2} \times \left(9 + 280 h \exp\left\{ -4170 \left[\left(\frac{T}{T_0} \right)^{-1/3} - 1 \right] \right\} \right) \quad (4.9)$$

respectively. In equations (4.7)-(4.9), $p_r = 101.325$ kPa and $T_0 = 293.15$ K, and f, f_{rO}, and f_{rN} are measured in Hz. The variable h represents the molar concentration of water vapor. The units of atmospheric pressure and temperature p_a and T are kPa and K, respectively.

In essence, attenuation due to atmospheric absorption is a function of the temperature, humidity, and pressure of the air, and the frequency of the sound. The formulation presented in equations (4.7)-(4.9) produces attenuation that is accurate to within $\pm$ 10% for sounds with a frequency to pressure ratio of 4×10^4 Hz Pa^{-1} to 10 Hz Pa^{-1}, traveling through air within the temperature range -20°C to 50°C (i.e., 253.15 to 323.15 K), for atmospheric pressure below 200 kPa, and for molar concentrations of water vapor in air of between 0.05 - 5% (ISO9613-1 1993). The variation

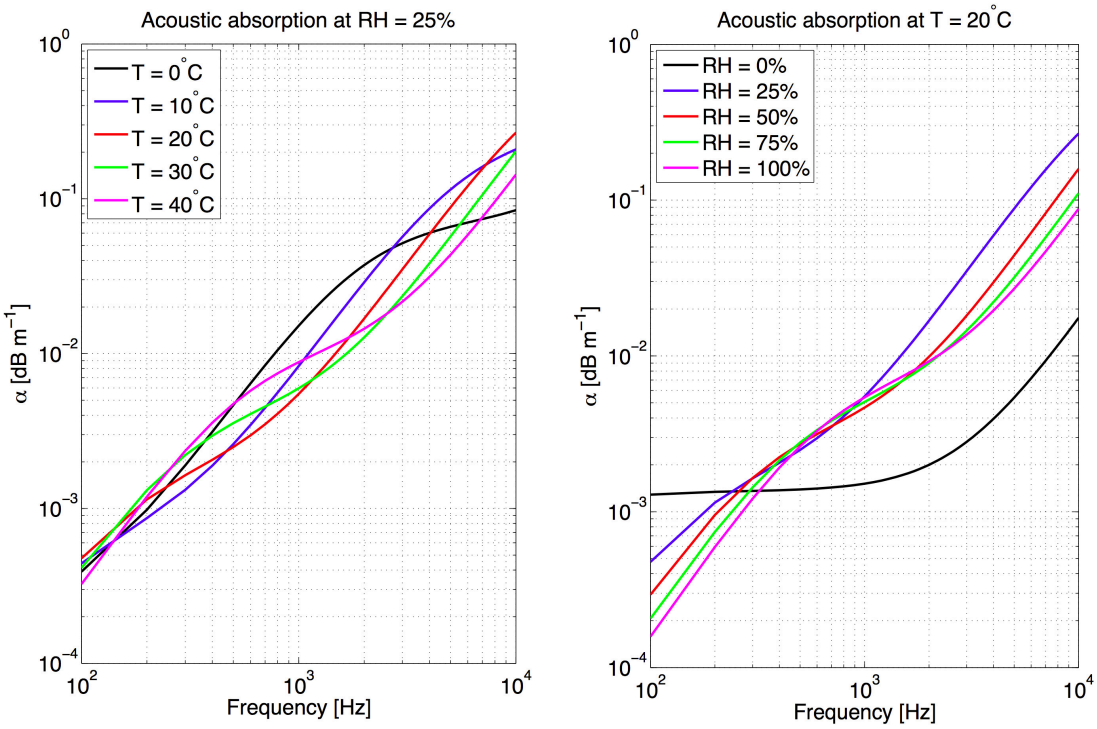

Figure 4.5: Left panel: the variation of atmospheric absorption with temperature for a constant relative humidity of 25%. Right panel: the variation of absorption with humidity for a constant temperature of 20°.

of atmospheric absorption with temperature and relative humidity is illustrated in Fig. 4.5 for a range of typical atmospheric temperature and relative humidity values.

In a similar manner to the calculation of radial velocity in the simulator, the atmospheric attenuation in the sodar simulator is calculated by spatially and temporally interpolating the LES fields of temperature and humidity to the beam point locations, before evaluating the attenuation at each point using equations (4.7), (4.8) and (4.9).

4.2.6 Complex acoustic signal

The power of the acoustic signal after time τ at each beam point is modeled following equation (3.10) (repeated here for clarity)

$$P_R = P_t G A_e \sigma_s \frac{c\tau}{2} \frac{e^{-2\alpha r}}{r^2}, \tag{4.10}$$

where P_T is the transmitted power, G is the antenna transmitting efficiency, A_e is the effective area of the antenna, σ_s is the backscattering cross-section (as calculated using equation (4.5)), α is the atmospheric absorption, and r is the range (Little 1969). In the sodar simulator, P_T, G and A_e are all set to physically realistic values, although each variable can be altered by the user to more accurately represent a particular make and model of sodar.

The phase of the signal, Φ, at each beam point is calculated as

$$\Phi = \exp(-j2\pi(f_T + f_s)t), \tag{4.11}$$

where f_T is the transmitted frequency and f_s is the frequency shift due to non-zero radial velocity, calculated as

$$f_s = \frac{2f_T v_r}{c}. \tag{4.12}$$

Within the sodar simulator, both the frequency shift f_s and the speed of sound c are calculated at each beam point, using the radial velocities that were calculated using equation (4.4) and the speed of sound field that was calculated using equation (4.2), respectively.

The phase and amplitude of the signal are next combined at each beam point to give a set of complex voltage points

$$V = I + jQ = \sqrt{P_R}\exp(-j2\pi(f_T + f_s)t). \tag{4.13}$$

The complex voltages are coherently summed across the $b \times b$ beam points at each individual range, n. This provides a complex time series of length n covering the entire range of each beam. The time series is divided into height segments for further signal processing in order to minimize computational expense.

4.2.7 Acoustic background noise

Acoustic background noise measured by sodars is not easily characterized and contains components due to localized effects, such as traffic, and echoes from nearby

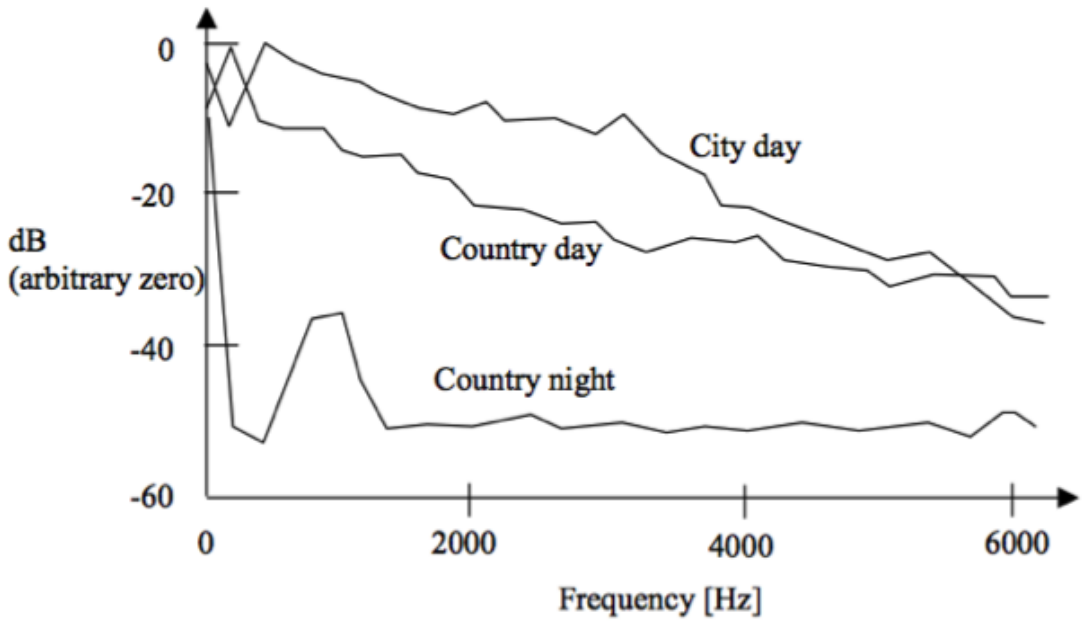

Figure 4.6: The variation of acoustic background noise between urban and rural areas, from Bradley (2008a).

stationary objects, as demonstrated in Fig. 4.6. Bradley (2012a) examined measured noise signals, separating sodar self-noise from external noise, and found that spurious peaks in the measured noise spectrum were of unknown origin, while the external noise between these peaks followed a $1/f$ frequency dependence. In this study the background noise is modeled as pink, such that it follows the $1/f$ frequency dependence seen by Bradley (2012a).

To model the pink noise, Gaussian white noise is first generated and a filter is then applied to create the $1/f$ roll-off that characterizes pink noise. Taking the FFT of the pink time-series noise provides the frequency content of the noise spectrum. An example noise spectrum is shown in red in the left panel of Fig. 4.7. To emulate a physical sodar, a noise spectrum (denoted N_S) is generated prior to each new set of measurements, where a set is defined as one completion of the pulse transmission and receipt in each of the independent beam directions. The same filter is used to generate pink time-series noise signals which are added to the complex time-series signal within each range gate. Taking the FFT of the time-series signal with added

noise, then multiplying the FFT by its conjugate and squaring the result, produces the noisy Doppler spectrum (shown in black in Fig. 4.7). This completes the simulation of the signal.

4.2.8 Spectral processing

One of the main benefits of the sodar simulator is that it allows the testing of signal processing techniques. Of particular importance is the ability to reprocess the same time-series data multiple times using different signal processing algorithms, allowing for a fair comparison of the algorithms.

To process the synthesized signals, at each range gate and for each beam, an FFT is taken of the complex time series signal within that range gate. The complex time series contains contributions from both the backscattered acoustic signal and the modeled background noise, as described in the previous section. Multiplying the result of the FFT (Y) by its conjugate (Y^*), and squaring the result, gives the Doppler power spectrum, $S = Y \times Y^*$, for each range gate. The signal processing is performed on a gate-by-gate basis, unlike a physical sodar which typically operates on a full vertical time-series profile at once. The single gate processing was chosen purely to minimize computational expense in the model. Once a full scan containing all the range gates and all beams is completed, a single separate noise spectrum, N_S, is generated. This single noise spectrum emulates the noise spectra that a sodar measures from a defined noise height, well above the range of the retrievable signal, prior to each scan.

The single generated noise spectrum is subtracted from the noisy Doppler spectrum, S, at each range gate to give an effective de-noised Doppler spectrum, S_D (illustrated by the black line in Fig. 4.7a). The same single noise spectrum, N_S, is used for each range gate and beam within one scan. This allows us to best emulate a physical sodar, which uses the measured noise spectra obtained before each scan to

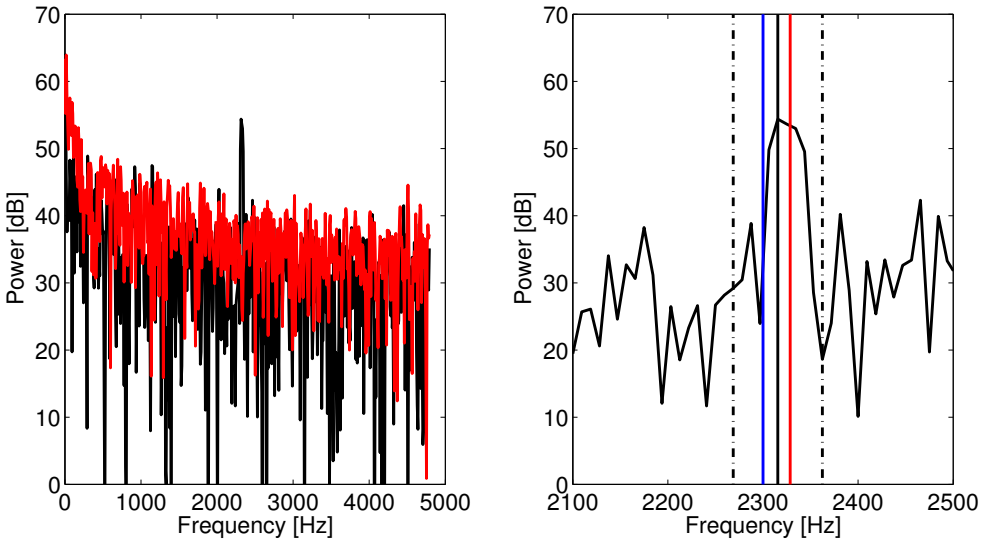

Figure 4.7: Left panel: example complete frequency spectrum with included noise output by the simulator for an off-vertical beam at a range of 450 m (black line). The generated noise spectrum is shown in red. Right panel: illustration of the spectral processing. Points within a range of $2m+1$ frequency bins are used to calculate the Doppler frequency. Here m=5 and points in-between the black dashed lines are included. The Doppler frequency used to calculate the radial velocity is marked in red. The transmitted frequency is marked by the blue line.

remove noise from all the received Doppler spectra within the scan. Further signal processing is performed on the de-noised Doppler spectrum, S_D, which mimics the frequency spectra output by a conventional sodar.

Several signal processing techniques have been tested. The one used in this study is outlined in Bradley (2008a) and is illustrated in Fig. 4.7b. First the peak of the de-noised Doppler spectrum is identified (vertical black line of Fig. 4.7b). The frequency range to be used for determination of the frequency shift is defined by a number $2m+1$ frequency bins, centered upon the identified peak, where m is a positive integer. In this study we used m=5. The bounds of this frequency range are marked in Fig. 4.7b by dashed black lines. The frequency shift within these bounds is determined by calculating the first moment of the spectrum via the use of a maximum likelihood estimator which models the frequency spectrum as normally distributed. Only spectra within the range bounded by frequencies $2m+1$ are used in the maximum likelihood

estimator. This is done to reduce the influence of background noise on determination of the frequency shift. The frequency identified as the first moment is shown on Fig. 4.7b in red. The frequency shift for that range gate is simply the difference between the first moment and the transmitted frequency.

The radial velocity for each range gate is retrieved by rearranging equation (4.12) as

$$v_r = \frac{f_s c}{2 f_T}, \tag{4.14}$$

where c is calculated from (4.2) using the mean value of temperature for beam points located within the considered range gate.

4.2.9 Post-processing

If three or more independent non-coplanar beam directions are used, then the wind components u, v and w can be retrieved from the radial velocity values as detailed in Section 3.5.4. Typically either three or five independent beams are used.

One of the advantages of the simulator is that it allows for a variety of averaging periods to be used on the same data, thus enabling an examination of the effect of the averaging period on the derived wind fields. The time averaging is performed by summing the individual Doppler spectrum contained within the desired averaging period, and recalculating the Doppler velocity based on the technique depicted in Fig. 4.7. Statistical properties of the retrieved wind profiles can also be determined for finite averaging periods, such as the vertical velocity variance, σ_w.

Aside from the retrieved wind profiles, the power of the returned signal is another important product yielded by the simulator. This may be in several forms: a simple maximum returned power (measured in dB) of each individual spectrum, if it is raw instantaneous data that are being considered; the maximum power of the summed spectra for a given averaging time; or a signal-to-noise ratio, which can be calculated from the synthesized noise spectra for both instantaneous or time-averaged signal

spectra. The calculation of SNR is especially useful, as it allows for evaluation of further statistical properties. One such property is profiles of data availability percentage using different averaging times, which are evaluated using a chosen value of SNR as a threshold.

The post-processing, in which the wind components are retrieved from the radial velocities, and the power and SNR at different averaging periods are calculated, is the final step within the sodar simulator. Once the wind profiles and power spectra have been derived, the performance of the sodar simulator can be evaluated. There are several methods used to evaluate how well the sodar simulator has performed in terms of the retrieved output parameters. The different validation techniques used on the simulator output are the subject of Chapter 5, which follows.

Chapter 5

Validation techniques

To evaluate the performance of the sodar simulator, it is necessary to compare the derived three-component wind field back to the original wind field from the LES output. We chose to compare the results in two separate ways. The first way is to compare the estimated wind field from the simulator with wind values taken directly from the LES output. This is equivalent to comparing winds from a sodar in a field campaign with winds measured on an instrumented tower (with the neglect of instrument error). One issue that can arise with this validation method is that it is difficult to separate the differences due to the spatial continuity assumption inherent in the DBS technique from the differences due to the spectral processing. To help separate these differences, we also estimate the wind fields using a so-called moment simulator, described in the following section.

5.1 Sodar moment simulator

The moment simulator uses the same beam locations and weighting as the full time-series sodar simulator, but it does not perform the spectral processing; rather, the radial velocities contained within one range gate are individual power weighted using C_T^2/T^2 and then simply averaged to gain a representative value for that range gate. Thus, a comparison of the wind fields estimated by the moment simulator and the virtual instrumented tower exposes the differences that result solely from the spatial averaging implied by the DBS technique. This is a useful technique because, as stated by Antoniou et al. (2003), there is no feasible way to make point and volume measurements directly comparable without distributing many point measurement devices throughout the volume encompassed by the sodar. While the moment simulator does

now allow for a direct comparison of results between the point and volume measurements, it does allow for an investigation of the biases between these measurement approaches, which is useful in itself. The challenges of comparing sodar- and instrumented mast-derived winds are investigated thoroughly by Bradley (2013), who states that the inherent difference between vector and scalar averaging causes an average root-mean-square (RMS) error in the measured winds of around 2%. Use of the DBS technique is based on the assumption that each of the wind components u, v, and w are constant across the region encompassed by the beam locations. Invalidity of this assumption caused by nonuniform vertical velocities or horizontal shear across the observed region can result in a bias in the wind estimates (shown for wind profiling radars by Koscielny et al. 1984).

Many field experiments have been performed to verify wind field measurements by sodars. Crescenti (1996) summarized many of the earlier studies that compared sodar-estimated winds with those from cup/vane anemometers, sonic anemometers, and lidars. It was concluded that Doppler sodars can reproduce reliable profiles of the mean wind speed and direction, with homogeneity in the wind fields leading to better comparisons between instruments. Inhomogeneity in the boundary layer adds complication and hampers interpretation of the resulting wind fields (Crescenti 1996). These effects can be examined by intercomparison of the winds retrieved from the virtual instrumented tower (LES) and the moment simulator.

5.2 UAS simulator

5.2.1 Unmanned aerial system

The University of Oklahoma operates an unmanned aerial system (UAS) known as the Small Multifunction Autonomous Research and Teaching sonde (SMARTsonde, Bonin et al. 2012; Bonin et al. 2013). The SMARTsonde is primarily used for research into the structure of the atmospheric boundary layer. The particular airframe used

in this study was a Funjet Ultra, which has a wingspan of 0.8 m and a weight of approximately 1 kg, has an endurance of 30 minutes and a cruising speed of 15 m s^{-1}. The SMARTsonde is equipped with sensors to measure the atmospheric temperature, pressure, and relative humidity. The response time for the pressure sensor is 0.5 s, and it is accurate to ± 1.5 Pa. The response time for the relative humidity and temperature sensors is 8 s and the accuracies are ±1.8 % and ±0.5 K respectively.

The SMARTsonde employs the open source Paparazzi autopilot system for autonomous flight, although a pilot with a radio controller manually operates the aircraft during takeoff and landing. A more complete description of the UAS package can be found in Bonin et al. (2013). As an update, an inertial measurement unit has recently been integrated into the autopilot for determination of the aircraft altitude.

5.2.2 Simulating the UAS

A simulator was developed in order to provide a secondary method of validation for the sodar simulator. The UAS sampling the atmosphere can retrieve both vertical profiles of the horizontal wind components and vertical profiles of the structure function parameter for temperature, C_T^2, which as shown in equation 4.10, is proportional to the sodar returned power. This means that the UAS can be used to compare against the sodar-derived wind profile and returned power.

In order to be able to use the UAS as a suitable comparison tool for the retrieved wind and thermodynamic profiles from the sodar simulator, a UAS simulator was developed. The UAS simulator ingests the same three-dimensional turbulent flow fields output by the UAS as used to populate the sodar simulator.

5.2.3 Simulated flight plans

As in the sodar simulator, many aspects of the UAS simulator are defined by the user, such that the simulator can be used to emulate desired flight plans. The OU

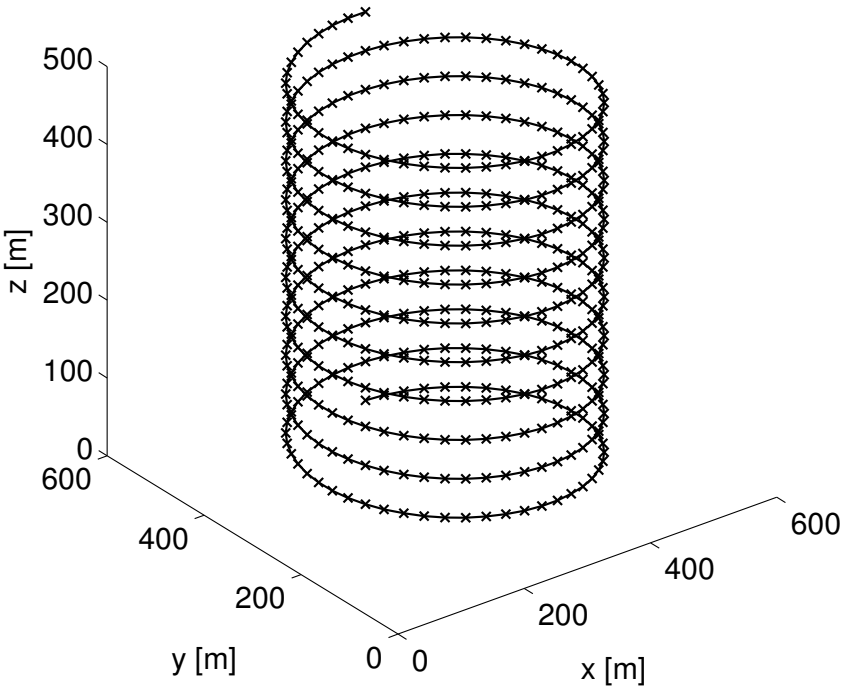

Figure 5.1: The flight path used by the UAS for measuring temperature, which can be used to retrieve a vertical profile of the horizontal wind components u and v.

SMARTSonde utilizes two main flight patterns for boundary-layer research. The first is a steady helical ascent, which is used alongside specially designed retrieval algorithms to derive the wind speed and direction Bonin et al. (2013). This flight path is illustrated in Fig. 5.1. The second flight pattern is applied to examine the spatial structure of thermodynamic quantities such as temperature and humidity. This flight plan entails the SMARTSonde flying in a number of consecutive circles at each height of interest, before ascending to the next height (illustrated in Fig. 5.2). The UAS simulator can mimic both of these flight plans.

5.2.4 Structure of the UAS simulator

User-defined flight plan characteristics in the UAS simulator include the circle radius, the desired airspeed, the height interval at which to sample, and the amount of

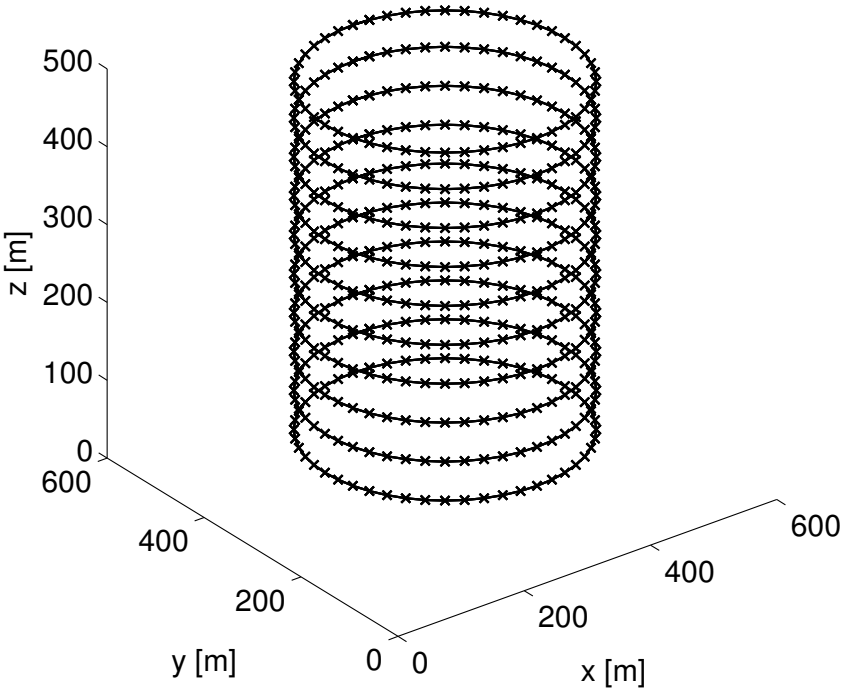

Figure 5.2: The flight path used by the UAS for measuring temperature, which can be used to calculate C_T^2.

time for the simulated UAS to spend at each height before ascending. These flight characteristics are then translated into cartesian co-ordinates within the LES grid, at times corresponding to each measurement. The LES temperature and wind data are spatially and temporally interpolated to provide temperature and horizontal wind component (u in x-direction and v in y-direction) measurements at each sampling point, just as in the sodar simulator. The ability to alter the position of the simulated UAS flight path within the LES domain allows the simulated UAS flight path to be centered over the simulated sodar location. This ensures that the sampling volumes of the two instruments are as close as possible to one another.

The method used to calculate C_T^2 from the temperature time-series data calculated in the UAS simulator is described in detail in section 5.3.2.

5.3 Methods of Evaluating C_T^2

Five methods of evaluating C_T^2 are utilized in this study: direct evaluation from LES output, evaluation using a simulated UAS flight within the LES domain, derivation from UAS data recorded during a field experiment, derivation from simulated sodar acoustic return power, and derivation from sodar data recorded during a field experiment.. A thorough description of each of these methods is provided in the following subsections.

5.3.1 Direct evaluation of C_T^2 from LES output

In evaluating C_T^2 from LES data, only the four-dimensional temperature field is required, along with knowledge of the grid spacing and time resolution of the field. The method used calculates a running sum of $(\delta T)^2$ values for a given separation distance, in a specific direction, and within a horizontal cross-section of the domain located at a particular height. Once squared differences from the entire horizontal plane have been summed, an average is calculated. This procedure is repeated for each height

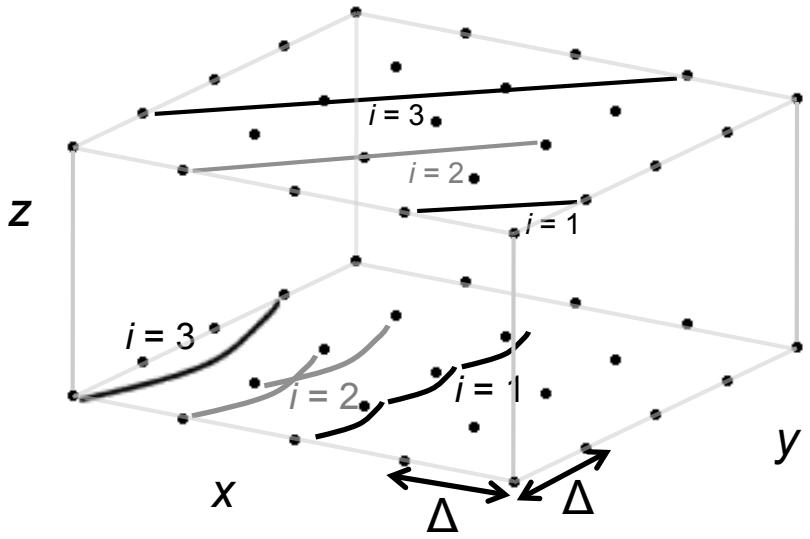

Figure 5.3: Illustration of the methods of evaluating C_T^2 directly from the LES. The lower z level illustrates the $(\delta T)^2|_y$ method and the upper z level shows the $(\delta T)^2|_{xy}$ method.

and time, before moving onto the next separation distance. This method is explained in detail in Wilson and Fedorovich (2012) for the refractive index n.

The squared temperature differences, $(\delta T)^2$, are calculated in four separate directions: x (zonal differences), y (meridional differences), the positive diagonal (henceforth referred to as xy) and the negative diagonal (henceforth yx). The equations for these calculations are

$$(\delta T)^2|_x(x, x_i, y, z, t) = [T(x, y, z, t) - T(x - x_i, y, z, t)]^2, \tag{5.1}$$

$$(\delta T)^2|_y(x, y, y_i, z, t) = [T(x, y, z, t) - T(x, y - y_i, z, t)]^2, \tag{5.2}$$

$$(\delta T)^2|_{xy}(x, x_i, y, y_i, z, t) = [T(x, y, z, t) - T(x - x_i, y - y_i, z, t)]^2, \tag{5.3}$$

and

$$(\delta T)^2|_{yx}(x, x_i, y, y_i, z, t) = [T(x, y, z, t) - T(x - x_i, y + y_i, z, t)]^2. \tag{5.4}$$

In equations (5.1)-(5.4), x_i is the separation distance in the x-direction and y_i is the separation distance in the y-direction. Note that in equations 5.3 and 5.4, $x_i = y_i$, such that the squared temperature differences are calculated on strict diagonals only. All the separation distances used to calculate $(\delta T)^2|_x$ and $(\delta T)^2|_y$ are multiples

of the LES horizontal grid spacing, $\Delta = \Delta x = \Delta y$, so no interpolation between numerical grid points in the horizontal planes is needed. The separation distances used to calculate $(\delta T)^2|_{xy}$ and $(\delta T)^2|_{yx}$ are multiples of the minimal diagonal distance ($\Delta xy = \sqrt{\Delta x^2 + \Delta y^2}$).

The method outlined above is illustrated in Fig. 5.3. The upper plane illustrates the calculation of $(\delta T)^2|_{xy}$ (equation 5.3) for separation distances of $x_i = y_i = \Delta$, 2Δ, and 3Δ, and the lower plane shows the calculation of $(\delta T)^2|_y$ (equation 5.2) for separation distances of $y_i = \Delta$, 2Δ, and 3Δ.

The $(\delta T)^2|_x$ values are calculated for each separation distance $x_i = i\Delta$ where $i = 1, 2, ..., i_{\max}$, and for each height and time of interest. These values are then spatially averaged for each separation distance $i\Delta$. The range of values for i should be selected such that the maximum range encompasses a spatial range which is comparable to the spatial volume bounded by the sodar beams at the highest height range of interest. As an example, if the sodar simulator is operated with a zenith angle of 15°, and a maximum range of 500 m, then the length and width of the spatial area bounded by the beams at 500 m height is $500 \times \sin 15° = 129.4$ m. For this example scenario, the value of $i_{\max}$ should be chosen such that $i_{\max} \times \Delta \simeq 129.4$ m.

The method outlined in the previous paragraph is repeated for equation (5.2) using separation distances $y_i = i\Delta$ where $i = 1, 2, ..., i_{\max}$ as before, and for equations (5.3) and (5.4) using $x_i = y_i = i\Delta$ where $i = 1, 2, ..., i_{\max}$ (equating to straight-line separation distances of $i\Delta xy = \sqrt{2 \times 1^2}$, $\sqrt{2 \times 2^2}$, ..., $\sqrt{2 \times i_{\max}^2}$, where Δxy is the straight-line distance between the two temperature values). The corresponding structure function counterparts to equations (5.1)-(5.4) are calculated as

$$\overline{(\delta T)^2}|_x(z, i\Delta) = \langle (\delta T)^2|_x(x, x_i, y, z, t) \rangle_{x,y,t}, \tag{5.5}$$

$$\overline{(\delta T)^2}|_y(z, i\Delta) = \langle (\delta T)^2|_y(x, y, y_i, z, t) \rangle_{x,y,t}, \tag{5.6}$$

$$\overline{(\delta T)^2}|_{xy}(z, i\Delta xy) = \langle (\delta T)^2|_{xy}(x, x_i, y, y_i, z, t) \rangle_{x,y,t}, \tag{5.7}$$

and

$$\overline{(\delta T)^2}|_{yx}(z, i\Delta xy) = \langle (\delta T)^2|_{yx}(x, x_i, y, y_i, z, t)\rangle_{x,y,t}. \tag{5.8}$$

Conceptually, C_T^2 makes sense only when the turbulence is assumed to be isotropic and homogeneous within the inertial subrange of spatial scales. Use of the LES allows for examination of the validity of the isotropy assumption by comparing the $\overline{(\delta T)^2}|$ values for different directions. If the turbulence is close to the isotropic state, there should be only minor differences between the values of $\overline{(\delta T)^2}|$ values calculated using equations (5.5)-(5.8).

A benefit of using LES output to examine C_T^2 is that it is possible to examine whether the simulated turbulence is within the inertial subrange of scale by examining the separation distance dependence. The C_T^2 data can be split into bins depending upon the separation distance, and C_T^2 can be plotted against separation distance to examine the relationship. For turbulence to be within the inertial subrange, it should follow the 2/3 law predicted by theory. Examination of the separation distance dependence allows the determination of the range over which the turbulence is considered to be within the inertial subrange. This feature is not used extensively within the current study, as separation distance dependence of C_T^2 cannot be determined from real or simulated sodar output data. This method is used in the current study primarily to identify the range of scales over which the turbulence falls within the inertial subrange.

5.3.2 Evaluation of C_T^2 from the UAS simulator

The UAS simulator provides T, u, and v values corresponding to the simulated location of the UAS within the LES domain for the duration of the simulated flight. For each individual measurement height, the temperature data are combined to create all possible data pairs (providing $n_t \times \frac{n_t-1}{2}$ unique pairs of temperature data, where n_t is

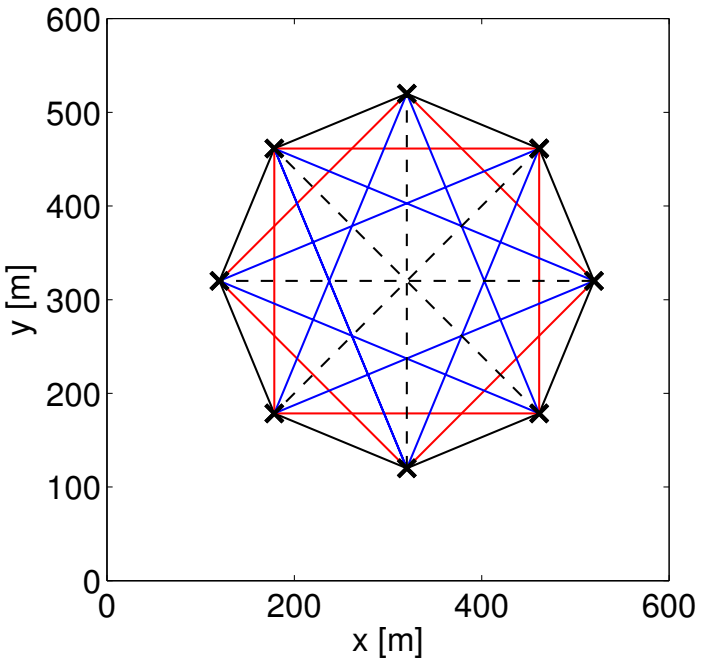

Figure 5.4: Illustration of the methods of evaluating C_T^2 from the UAS flight and simulated flight. Only eight data points around the circle are used for clarity. In this example, C_T^2 is calculated at four separation distances, 153 m (solid black line), 283 m (red line), 370 m (blue line), and 400 m (dashed black line). The radius of the circular flight path is 200 m.

the total number of temperature measurements recorded at that height). A squared temperature difference, $(\delta T)^2$, can then be calculated from each temperature pair.

As the simulated UAS co-ordinates are known at every time, the separation distance between the two data points, r, is then calculated for each temperature pair. The calculation of the temperature difference pairs and separation distance is illustrated in Fig. 5.4. The separation distances are modified to account for the advection of turbulence past the UAS by the wind, as detailed in Bonin et al. (2014). With advection correction, the separation distance is calculated as

$$r = \sqrt{(x_i - u\Delta t)^2 + (y_i - v\Delta t)^2} \tag{5.9}$$

where x_i and y_i are the separation distances in the zonal and meridional directions, u and v are the zonal and meridional winds speeds in ms^{-1} and Δt is the time between consecutive measurements.

Advection correction is not performed when C_T^2 is evaluated directly from the LES, as described in section 5.3.1, since in this case all the temperature values at a given height correspond to the same moment of time. However, the real and simulated UAS both experience advection of turbulence past the UAS by the wind. To account for the advection of turbulence on timescales equal to the time between measurements, Taylor's hypothesis of frozen turbulence is utilized (Taylor 1938).

Once the separation distance is calculated using equation 5.9, with advection correction accounted for, each squared temperature difference $(\delta T)^2$ value is divided by $r^{2/3}$ as in equation 2.6. The structure function parameter for temperature, C_T^2, is then evaluated for each height of interest by averaging all the individual $(\delta T)^2/r^{2/3}$ values recorded at that height.

5.3.3 Evaluation of C_T^2 from UAS data

The method used for evaluating (δT^2), and subsequently C_T^2, from data recorded by the UAS is very similar to that described in section 5.3.2 for the UAS simulator. The spatial co-ordinates of the UAS during the flight are recorded by a GPS sensor, at the same temporal rate that the temperature is recorded. The wind components u and v can be estimated using the heading direction, throttle, and instantaneous airspeed data recorded by the UAS. The method used to estimate C_T^2 from the UAS data is described in detail in Bonin et al. (2013).

The recorded temperature data is divided into height batches based upon the height recorded by the GPS. The separation of the UAS measurements by height is performed manually to ensure that time periods during take-off, final descent, landing, and during which the plane is ascending from one height to the next are excluded. The recorded locations, temperatures, and derived wind components are then used to calculate (δT^2) and C_T^2 as outlined in section 5.3.2.

5.3.4 Evaluation of C_T^2 from the sodar simulator

The method of retrieving C_T^2 from simulated sodar data makes use of the returned sodar power (e.g., Weill et al. 1980), P_R. The equation for returned power is provided as equation 4.10, and each of the parameters is defined in the following section. Equation 4.10 is repeated below for clarity.

$$P_R = P_t G A_e \sigma_s \frac{c\tau}{2} \frac{e^{-2\alpha r}}{r^2},$$

As a reminder, the backscattering cross-section is related to C_T^2 through $\sigma_s = 4 \times 10^{-3} \lambda^{-1/3} \frac{C_T^2}{T^2}$, (equation 4.5), where λ is the sodar wavelength in m, and T is measured in K (following Little 1969). All of the components of equation 4.10 are readily available in the sodar simulator, since each component must be calculated in order for the simulator to evaluate P_R.

C_T^2 can be evaluated using the sodar simulator either during the process of calculating P_R, or as a part of post-processing. As described in section 4.2.4, C_T^2 is interpolated on to the beam point locations within the simulator. At this stage, the averaging technique used in the moment simulator (section 5.1) can be invoked, and all the C_T^2 values contained within one range are averaged to give a mean representative value for that range gate. In order to calculate C_T^2 as a part of signal post processing, each of the parameters in equation 4.10 must be either averaged within each range gate using the technique from the moment simulator and saved, or the parameters must be estimated or measured using a different instrument. The method using different instrumentation and/or estimation is used on real sodar output data, and is discussed in the following section.

5.3.5 Evaluation of C_T^2 from sodar data

A sodar provides, as output data, values of P_R and r, whilst P_t, G, A_e, and τ are functions of the sodar system used. The value of λ is indirectly set by the sodar user, and is a function of the acoustic transmit frequency, f, and the speed of sound. In

order to estimate C_T^2 profiles from measured sodar data, first σ_s must be derived from the P_R profile. This requires estimation or knowledge of the values of G, A_e, P_t, τ, c and α.

The acoustic attenuation α can be calculated according to ISO9613-1 (1993). The full calculation of α requires input of vertical profiles of temperature (T) and specific humidity (q) encompassing the range over which α is to be calculated, and the acoustic transmit frequency (f). Typically, a sodar operating without radio acoustic sounding system (RASS) will rely on additional instrumentation to provide these measurements. The T and q profiles could be provided by a UAS, as in Wainwright et al. (2014a). In the case that vertical profiles of T and q are unavailable, measurements taken at ground level can be used, and a constant profile assumed for α.

A temperature profile is also needed in order to calculate the speed of sound (through equation 4.2), another component in equation 4.10. Again, in the case that a temperature profile is unavailable, the temperature measured by the sodar at the instrument height can be substituted, and an isothermal profile assumed.

Once all of the components of equation 4.10 have been estimated or calculated, the acoustic backscattering coefficient σ_s can be derived. Once a σ_s profile is derived, the profile of C_T^2 can be derived using equation 4.5. This derivation also requires a value for T at each range gate, which can be provided by outside instrumentation, or an isothermal profile may again be assumed.

Using constant profiles of α and c introduces uncertainty into the C_T^2 values derived from sodar data. For a sodar operating using RASS technology, vertical profiles of T and q are provided as output variables, and this source of uncertainty is mitigated.

In this chapter several validation techniques which can be used to evaluate the performance of the sodar simulator have been described. Chapter 6, which follows,

now tests the sodar simulator on three separate case studies. The validation techniques described in this chapter are used in the following chapter to evaluate the sodar simulator performance during the three case studies.

Chapter 6

Case studies

In order to test the performance of the sodar simulator, three case studies were completed. For each of the case studies, the simulator is populated using data from the University of Oklahoma LES (OU-LES, Fedorovich et al. 2004b). Details of each particular simulation are given in the corresponding sections. The first two case studies represent the convective boundary layer (CBL), while the final case study tests the sodar simulator on an LES of the nocturnal stable boundary layer (SBL).

6.1 Convective boundary layer case studies

6.1.1 31 May 2009

6.1.1.1 Experiment description

The first case study used LES output from 14-15UTC on 31 May 2009, in a domain centered upon the Atmospheric Radiation Measurement Southern Great Plains (ARM SGP) site in Lamont, OK. The OU-LES was nudged incrementally at each time step (approximately 1 second) using vertical profiles of virtual potential temperature, water vapor mixing ratio, and the horizontal wind components from the Rapid Update Cycle (RUC; Benjamin et al. 1994), as described in Gibbs et al. (2011). The surface fluxes were prescribed using temporally-interpolated sensible and latent heat fluxes from the eddy correlation flux measurement system (ECOR; Cook and Pekour 2008) at the ARM SGP site. The LES output fields that are used in the simulator are the three-dimensional fields of potential temperature Θ, the wind components, u, v and w, and the water vapor mixing ratio q.

The simulation reproduced a developing convective boundary layer and was performed in a domain of size 1.28 km $\times$ 1.28 km $\times$ 1.5 km with a uniform grid spacing of $\Delta x = \Delta y = \Delta z = 5$ m. The lowest model level was located at 2.5 m above ground level. The employed sub-grid turbulence closure scheme is based upon Deardorff (1980). For testing the sodar simulator, a subset of the LES output data described in the previous two paragraphs was used. The LES subset was of size 400 m $\times$ 400 m $\times$ 600 m. One hour of data were used with the simulated fields reported every three seconds.

The transmitted sodar acoustic frequency was set to 2300 Hz. The maximum range was set to 500 m above ground level (AGL). The minimum range is not directly defined by the user but rather determined from a preset transmit-to-receive delay, which for this test was set at 0.1 s. This resulted in an effective minimum range of 26 m, which is not atypical compared to commercial sodars. The range resolution was 11.6 m (note that this is the approximate range resolution, evaluated using a sample sound speed of 340 m s^{-1}, and the actual range resolution will vary around 11.6 m due to variance in the temperature-dependent speed of sound). Forty range gates were simulated. Five beams were used, with the four off-vertical beams at a zenith angle of 17.5°, and a beamwidth of 10° for each beam. The azimuth angles of the four off-vertical beams were set to multiples of 90°, and aligned within the LES domain such that the off-vertical beams were in the direct north, east, south and west directions. The fields of u, v, w, and Θ were updated with each new beam direction at 3-s intervals, to provide a complete set of measurements every 15 s. The timing is typical for the operational sodar we are replicating in this case study, a Metek PCS.2000. The number of points that are equidistant in range within each beam ($b \times b$) is set to 5 $\times$ 5. The number of ranges n within each range gate is set to 640, which equates to a sampling frequency close to 39 kHz.

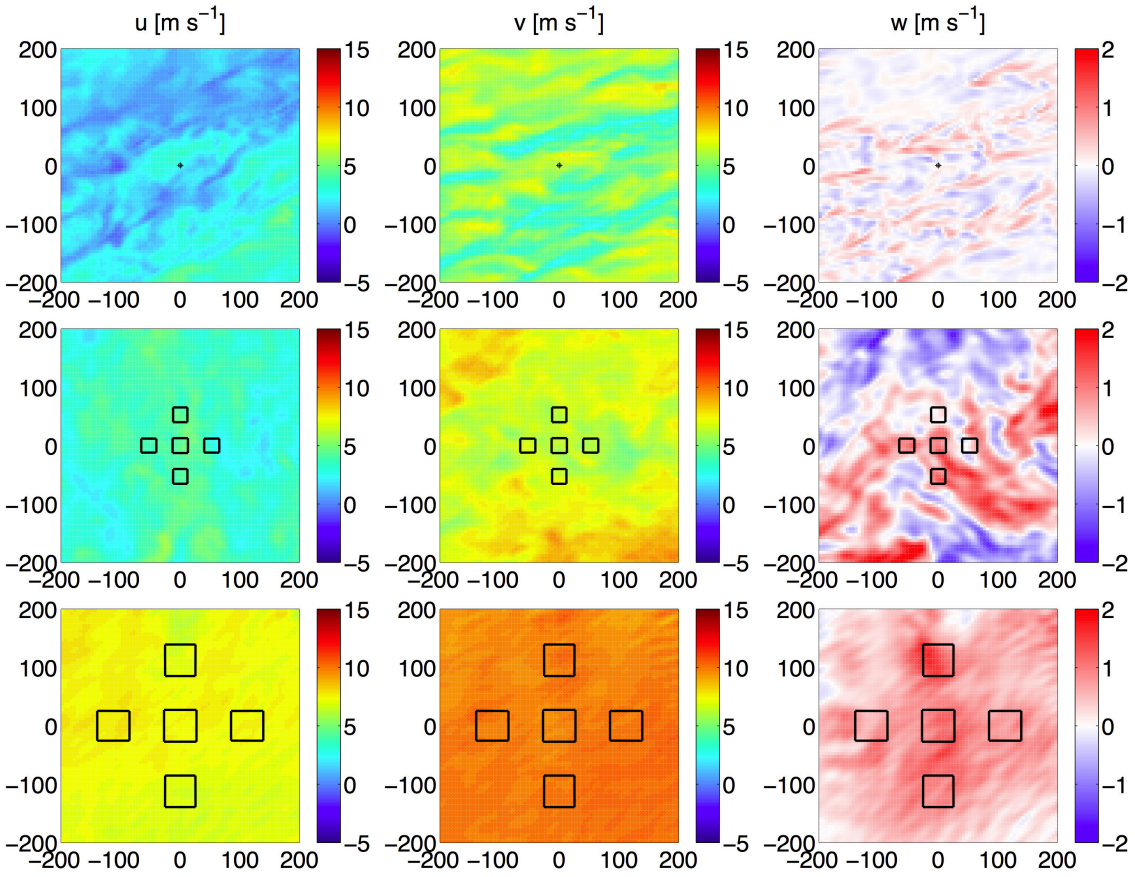

Figure 6.1: Snapshots of the LES wind component fields at 20 m (top row), 200 m (middle row) and 400 m (bottom row) above ground level for the 31 May 2009 LES output. Wind fields shown are from 30 minutes into the 60 minute simulation used to populate the sodar simulator. The areas encompassed by each sodar beam at that height are overlaid.

A sample of the numerically generated CBL flow is illustrated in Figs. 6.1 and 6.2. The boundary layer depth at the early time of the simulation is around 300 m, as indicated in Fig. 6.2. The horizontal wind components u and v are both seen to increase in magnitude above the boundary layer. Beneath 300 m, we see a pattern of updrafts (positive w) and downdrafts (negative w), with generally small magnitudes of w above the CBL top. Figure 6.1 illustrates the strong spatial variability of u, v and w within the boundary layer, and it is clear that the assumption that the wind components are homogeneous across the area encompassed by the five sodar beams is invalid.

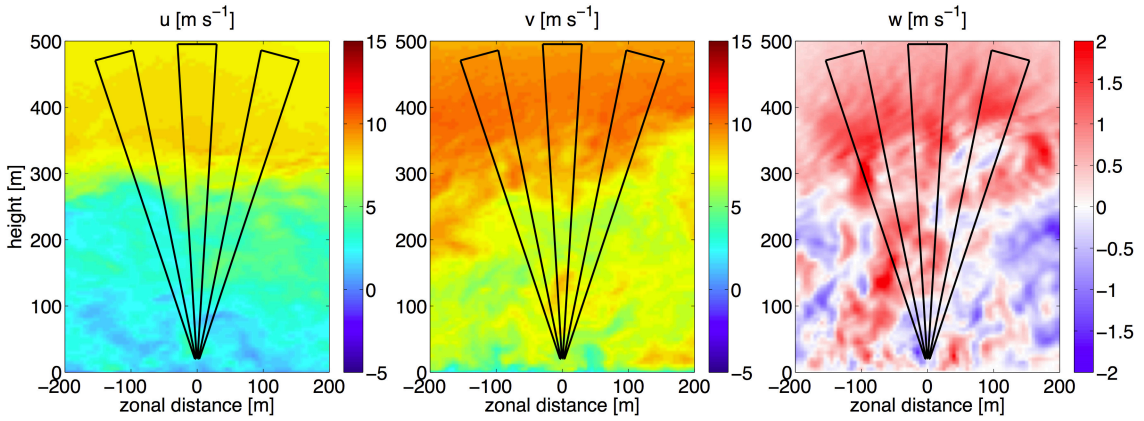

Figure 6.2: Snapshots of the LES wind component fields at $y = 0$ m for the 31 May 2009 case. The areas encompassed by three of the five beams are overlaid. The presented data illustrate the failure of the assumption of homogeneity in the wind component fields across the volume defined by the beam locations.

6.1.1.2 Results

Results based on the radial velocity calculations for the five beams of the sodar simulator are shown in Figs. 6.3 and 6.4. Figure 6.3 shows a time-averaged vertical profile of the root mean square (RMS) difference in the radial velocity calculated by spectral processing as compared to that from the moment simulator. One would expect the RMS difference to increase with height due to the effect of the assumption of spatial homogeneity inherent in the DBS technique. However, this does not seem to be the case. It can be seen that heights with increased RMS differences – namely, close to the surface and at heights close to the boundary layer top – are regions that exhibit increased spatial heterogeneity of the three-dimensional wind components.

It is clear from Fig. 6.4 that the RMS difference between u, v, and w derived from the moment and spectrally processed radial velocities does not necessarily tend to zero as the averaging period increases, as would be expected. The RMS difference in u and v, particularly for heights between 300 - 350 m, is higher for a 60 minute averaging period than for any of the shorter averaging periods, while the RMS difference decreases with a higher averaging period for time periods of up to 30 minutes. One possible cause of the larger RMS difference for heights of 300 - 350 m at the 60-minute

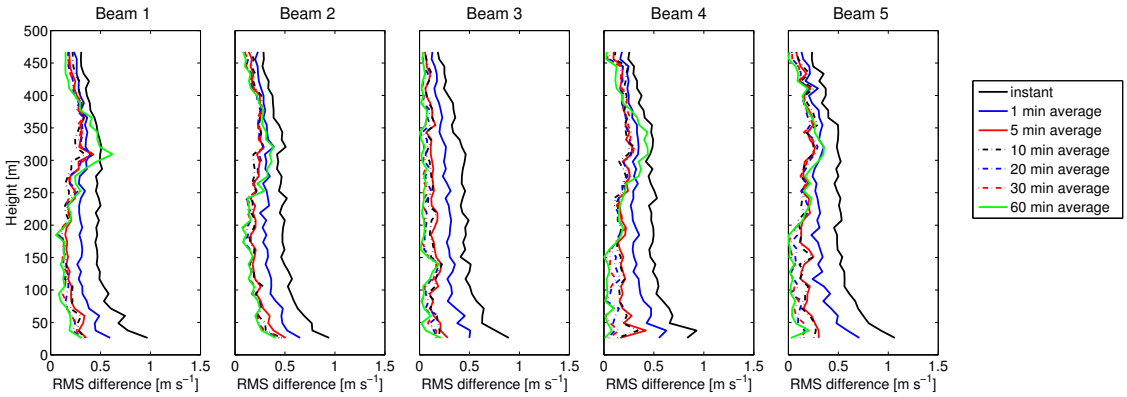

Figure 6.3: RMS difference in radial velocity between the moment and time-series sodar similar for each of the five beams for the 31 May 2009 case.

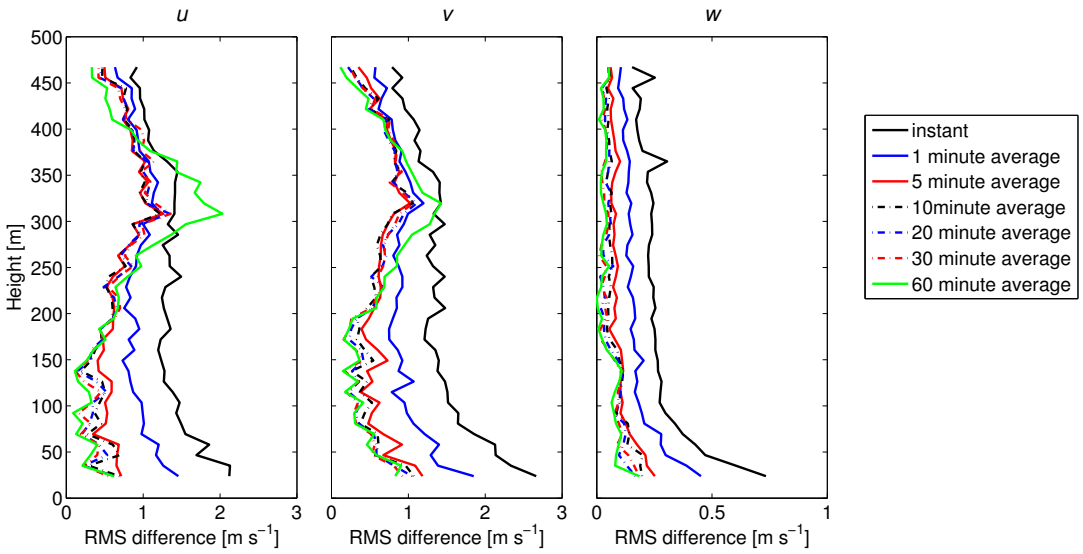

Figure 6.4: The RMS difference in the three-dimensional wind components as derived from the moment and spectrally processed radial velocities for the 31 May 2009 case.

averaging period is that this is the region in which the wind speed changes the most over the hour of simulation time. This is illustrated in the top row of Fig. 6.5, which shows how the wind components vary over the hour of simulation time. We can see in the top row of Fig. 6.5 that the greatest changes in the horizontal wind components u and v are in the height range 300 - 350 m, as the height of the boundary-layer top increases from 300 m at the beginning of the simulation to approximately 400 m by the end of the hour of simulation. The difference between u and v calculated using 30-minute and 60-minute averaging periods is illustrated in Fig. 6.6. It is clear from Fig. 6.6 that there is a considerable difference between the wind profiles in the two separate 30-minute averaging periods above 200 m, for both u and v and as calculated from the moment and spectrally-processed radial velocities. The 60-minute averaged profiles of u and v (shown by the dashed lines in Fig. 6.6) capture an average between the two 30-minute averaged wind profiles (solid line in Fig. 6.6). We note the difference between the 30-minute averaged u profiles derived from the moment and spectrally-processed radial velocities (comparing the black solid line to the closest red solid line in the left panel of Fig. 6.6) are smaller than the difference between the two sequential 30-minute average u profiles derived by either method. It can also be seen that difference between the two 30-minute average profiles of u for heights above 200 m is greater than the same difference for v, indicating that the change in u over the hour of simulation is greater than the change in v. It follows that the difference between the 60-minute average profiles of u calculated using the two methods will be greater than that for v, which can be seen both by comparing the dashed lines in Fig. 6.6 and by examining the green lines in the left and center panels of Fig. 6.4, which show that the maximum difference between the 60-minute average u profiles at 300 m is approximately 2 m s^{-1}, while that for v is lower at 1.5 m s^{-1}.

One may also expect that for all heights the RMS difference between the moment and spectrally processed radial velocities would tend to zero over a sufficiently long

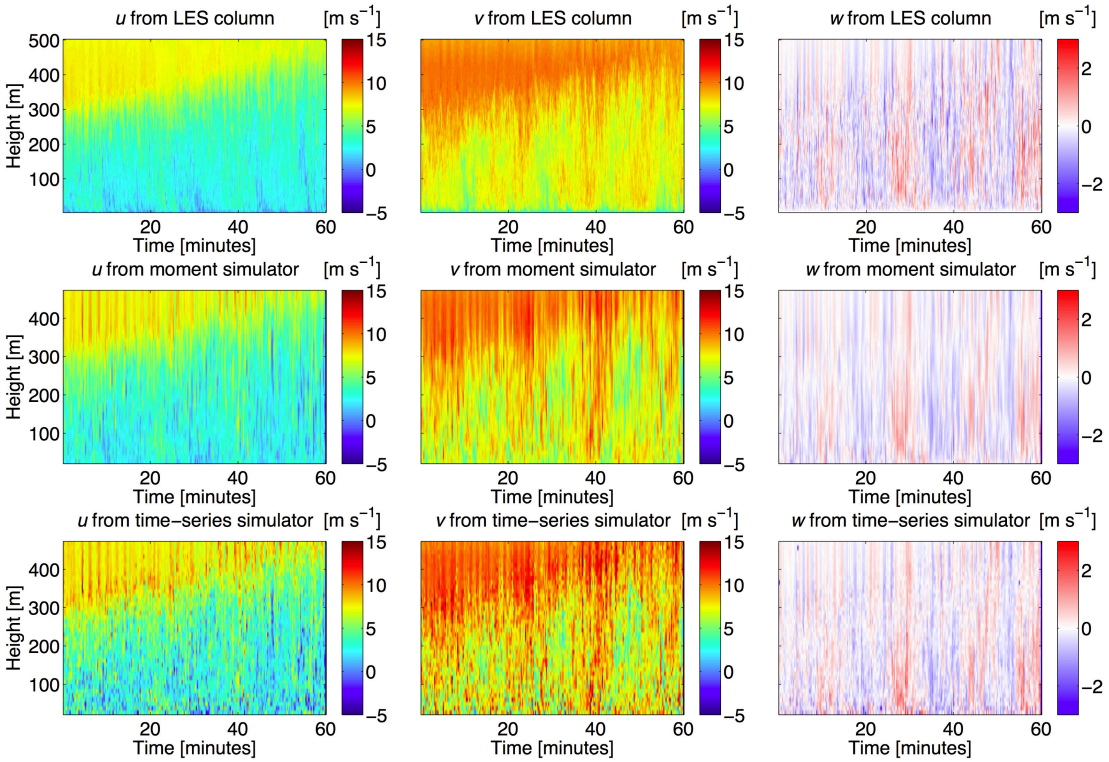

Figure 6.5: A comparison of the three-dimensional wind components u (left column), v (central column) and w (right column) as derived from the LES column directly above the simulated sodar (top row), derived from the moment simulator (middle row) and from the spectrally-processed data from the time-series simulator (bottom row) for the 31 May 2009 case.

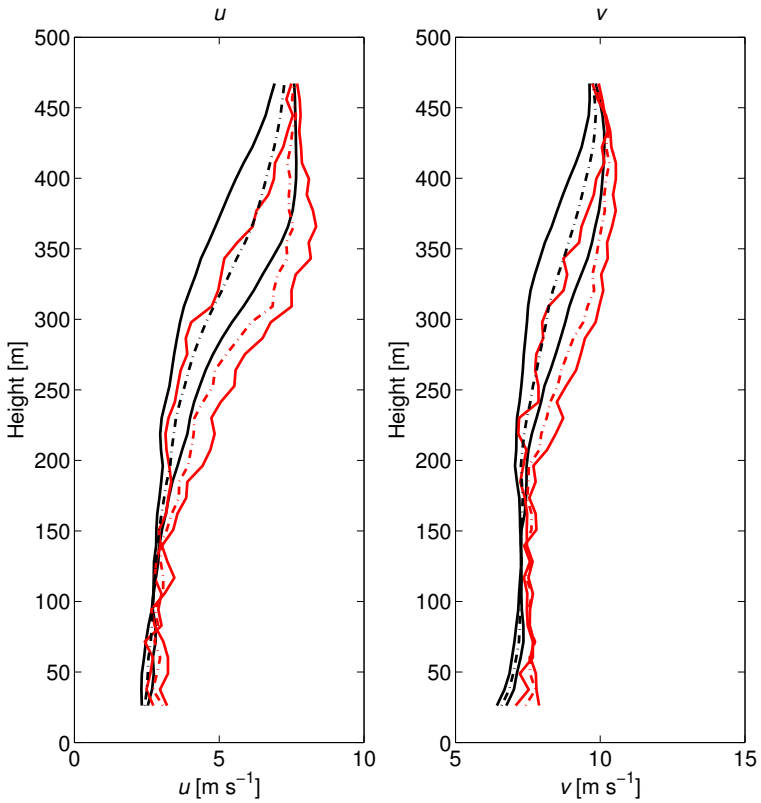

Figure 6.6: Left panel: the u profiles from the moment simulator with a 30-minute averaging period (black solid lines), and a 60-minute averaging period (black dashed line). The 30-minute averaged u profiles derived from the time-series simulator are shown by the solid red lines and the corresponding 60-minute average u profile is shown by the dashed red line. The right panel shows the corresponding profiles for v.

averaging period. We indeed see the RMS radial velocity difference decreasing as the length of the averaging period increases, but there seems to be a lower bound on the RMS difference for some beams and heights. This is most noticeable for beam 2 (the second panel in Fig. 6.3). Particularly for heights between 300 and 350 m, the RMS difference does not significantly decrease as the averaging period increases over 5 minutes, indicating a systematic bias of some kind.

The cause of this bias was thoroughly investigated. It was found that the highest RMS differences in radial velocity were collocated with range gates over which the radial velocity, and consequently frequency shift, changed the most over the height within that range gate. This is illustrated in Fig. 6.7 for a sample radial velocity from beam 3 (the vertical beam, which only contains contributions from w) for the first five minutes of simulation time. The upper panel of Fig. 6.7 shows the difference between the radial velocity from beam 3 for the moment and spectrally processed data, and the lower panel shows the corresponding maximum frequency difference (the absolute magnitude of the frequency difference is used) across the heights contained within each range gate. The correlation coefficient between the discrepancy in w and the maximum frequency shift over the range gate was calculated as 0.704. However, correlation between the discrepancy in w and the frequency shift does not imply a causation. The relation between the frequency shift across a range gate and its effect on the resulting radial velocity is explored in the following paragraphs.

If the radial velocity changes linearly (or almost linearly) with height within the range gate, then this introduces a linear variation in the frequency shift f_s in equation (4.11), resulting in the complex voltage

$$V = \sqrt{P_R} \exp\left\{-j2\pi[f_T + (f_s + At)]t\right\} \tag{6.1}$$

where A represents the frequency slope caused by the linear change in radial velocity over the depth of the range gate. Taking the FFT of the signal in equation (6.1) does not result in a spectral peak located exactly at $f_T + f_s + A/2$, which would be the

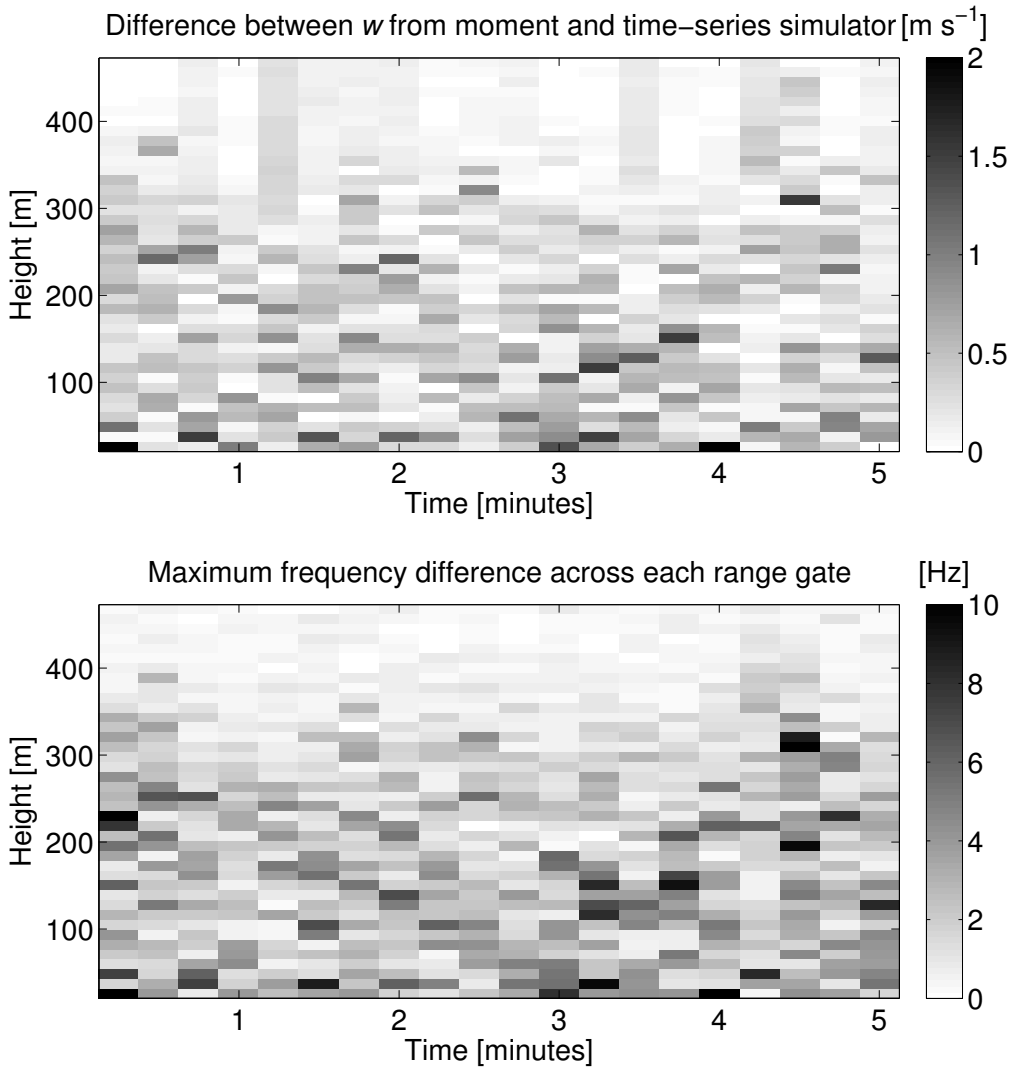

Figure 6.7: Upper panel: the absolute difference between the instantaneous radial velocity from beam 3 (vertical beam) derived from the moment and spectrally-processed data for the first five minutes of simulation time. Lower panel: the maximum frequency difference across the heights within each range gate.

instantaneous frequency at the central time moment within the FFT window, but rather results in smearing of the peak in the FFT and phase distortion (Masri and Bateman 1995). It is worth noting that the upper panel in Fig. 6.7 illustrates only the maximum frequency change.

A simple example illustrating the effect of linear frequency variation is shown in Fig. 6.8. The upper panel of Fig. 6.8 shows the variation of radial velocity with travel time for three different simulated signals. The time period shown covers the signal passing vertically through one range gate (approximately 11.6 m). Each of the three illustrated signals has a mean radial velocity of 0 m s^{-1}. The first signal (the black line) has a constant radial velocity of 0 m s^{-1} across the range gate, while the other two signals both have a 0 m s^{-1} mean radial velocity, but one signal (shown in blue) has a radial velocity which increases from -0.5 m s^{-1} to 0.5 m s^{-1} across the range gate, while the third signal (shown in red) has the same radial velocity content as the signal shown in blue, but with the radial velocity increasing rather than decreasing with height. For each of the illustrated radial velocity contents, a signal of the form $\exp\{-j2\pi[f_T + f_s]t\}$ is created. Each of the three signals with the radial velocity content illustrated in the upper panel of Fig. 6.8 is evaluated by performing a FFT on the signal, then multiplying the result of the FFT by its conjugate and squaring the result to give a power spectrum for each signal. The resulting power spectra are shown in the lower panel of Fig. 6.8. The color of the power spectra correspond to the radial velocity changes shown in the upper panel. It is clear from Fig. 6.8 that a linear radial velocity variation across a range gate can produce a distinct change in the peak frequency of the resulting spectrum. It can be seen that the resulting shift of the spectral peak corresponds to the sign of A from equation (6.1).

A fundamental assumption when interpreting the FFT results is that the signal statistics (up to second order) are wide-sense stationary, i.e., the signal statistics are stationary over the period that the FFT is taken. For sinusoidal signals of the type

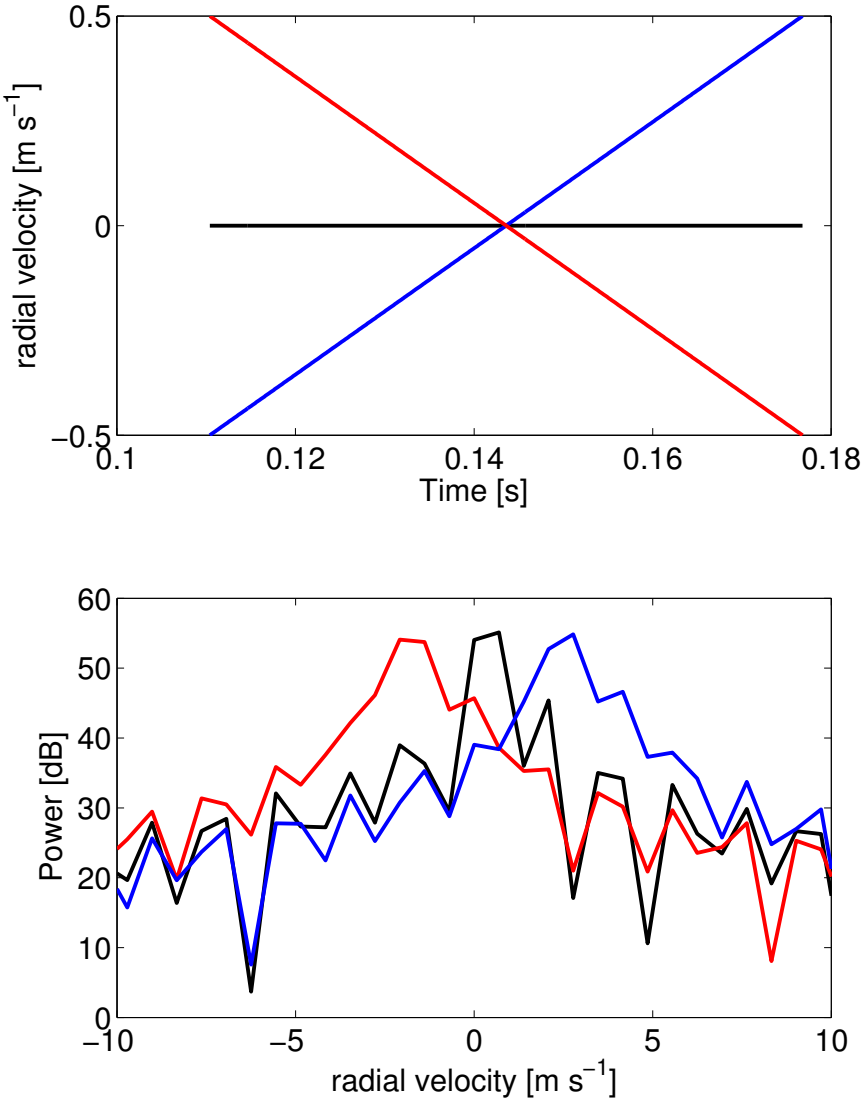

Figure 6.8: The effect of linear frequency variation within a range gate upon the derived radial velocities. Upper panel: the variation of radial velocity with signal travel time for three example signals. Lower panel: the frequency spectra resulting from processing signals with the radial velocity changes shown in the upper panel.

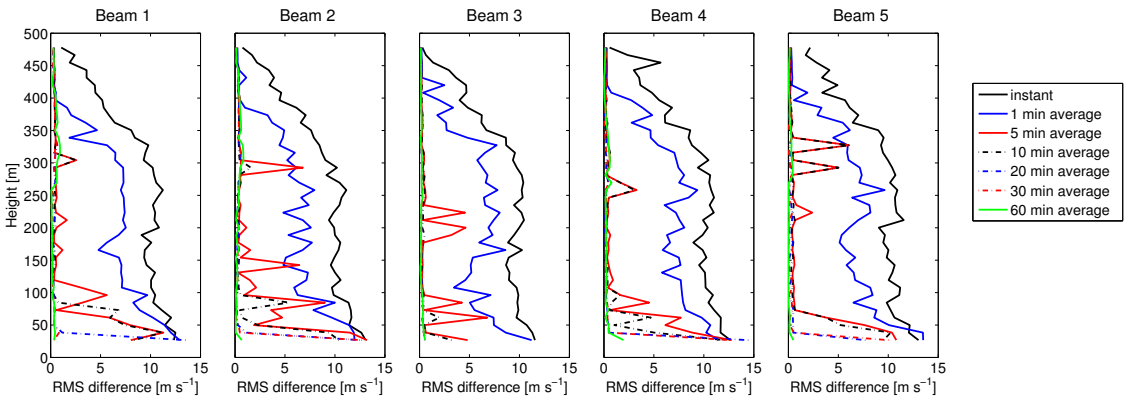

Figure 6.9: RMS difference in radial velocity between the moment and time-series sodar similar for each of the five beams for the 31 May 2009 case, when the sub-grid turbulence kinetic energy is included as in equations (6.2) - (6.4).

considered in equation (3.14), it is assumed that the amplitude and frequency remain constant within the FFT window. In our case, the FFT window corresponds to the time interval for which the signal is within one range gate. The introduction of linear (or nonlinear) frequency modulation obviously invalidates this assumption. However, the FFT method is used in most operational sodars, and all commercial sodars, as described in Section 3.5.1.

This frequency shift is not typically problematic for operational sodars, as the perfectly linear frequency modulation in the simulator is a result of the trilinear interpolation of u, v, and w to the beam point locations. For an operational sodar, the background and instrument noise would typically overshadow the linear radial velocity variation. For the sodar simulator, the effect of linear frequency modulation is more marked at lower heights, where the sodar beamwidth encompasses relatively few LES grid points. At higher levels there is rarely perfect linear frequency modulation due to radial velocity contributions from an increased number of LES grid points, which mitigates the frequency modulation effect to some degree. For this reason, it is recommended that LES output of sufficiently high spatial resolution to provide contributions from more than two grid points, even at the lowest range gate, is used.

One possible solution to overcome the linear variations in radial velocity is to include a contribution from the sub-grid turbulence kinetic energy (TKE) E in the three-dimensional wind components, in the form

$$\hat{u} = u + \epsilon\sqrt{\frac{2}{3}E}, \tag{6.2}$$

$$\hat{v} = v + \epsilon\sqrt{\frac{2}{3}E}, \tag{6.3}$$

and

$$\hat{w} = w + \epsilon\sqrt{\frac{2}{3}E}, \tag{6.4}$$

where ϵ is the output of a random number generator with a normal distribution and unity variance. The addition of TKE in this manner has previously been included in radar simulators (e.g., Cheong et al. 2008; Scipión et al. 2008). This idea was tested within the time series simulator, but it was found to introduce large errors into the spectrally processed radial velocities (Fig. 6.9), creating very large discrepancies between the moment and spectrally processed values of u, v, and w (Fig. 6.10, note the difference in the scale of the x-axis as compared to Fig. 6.4). Investigation into the cause of the discrepancies seen in Figs. 6.9 and 6.10 revealed that at heights close to the surface and the boundary layer height, values of E were as large as the resolved wind components, resulting in large errors in the spectrally derived radial velocity. A comparison of the magnitude of E against the resolved wind components is shown in Fig. 6.11.

In interpolating E down to a finer scale, the inherent scale dependence of the resolved velocity field and E are ignored. It is expected that on a spatial scale comparable to the beam point spacing used in the simulator (0.018 m in the vertical), the sub-grid contribution would be significantly lower than the resolved contribution. The issue of the scaleability of resolved velocity components and E is an important topic, and one which merits further study. However, it is outside of the scope of this dissertation. It is hoped that a thorough investigation into the scaleability of E

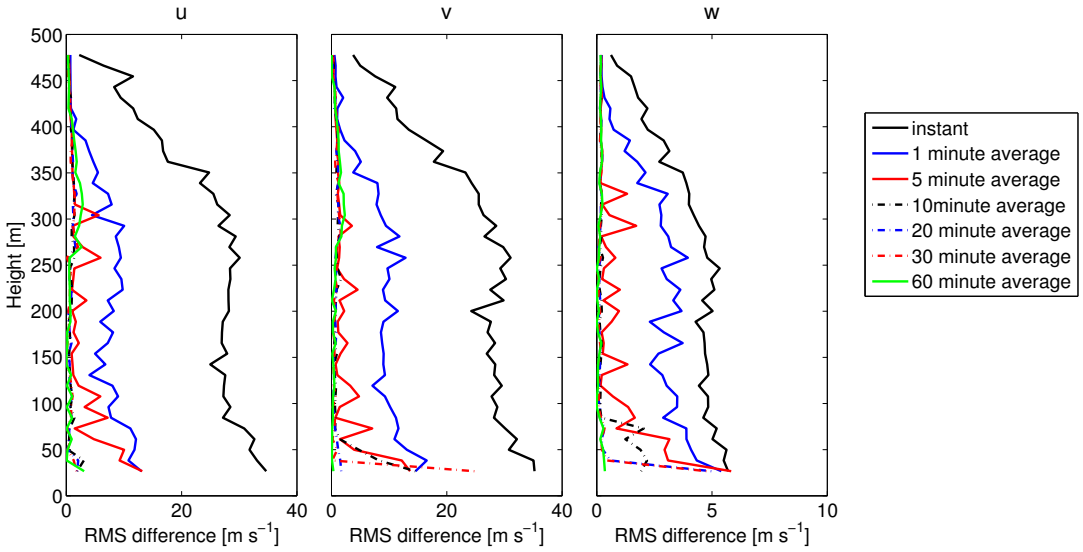

Figure 6.10: Three dimensional wind components as derived from the sodar simulator for the 31 May 2009 case, derived when the sub-grid TKE included in the wind components as in equations (6.2)-(6.4).

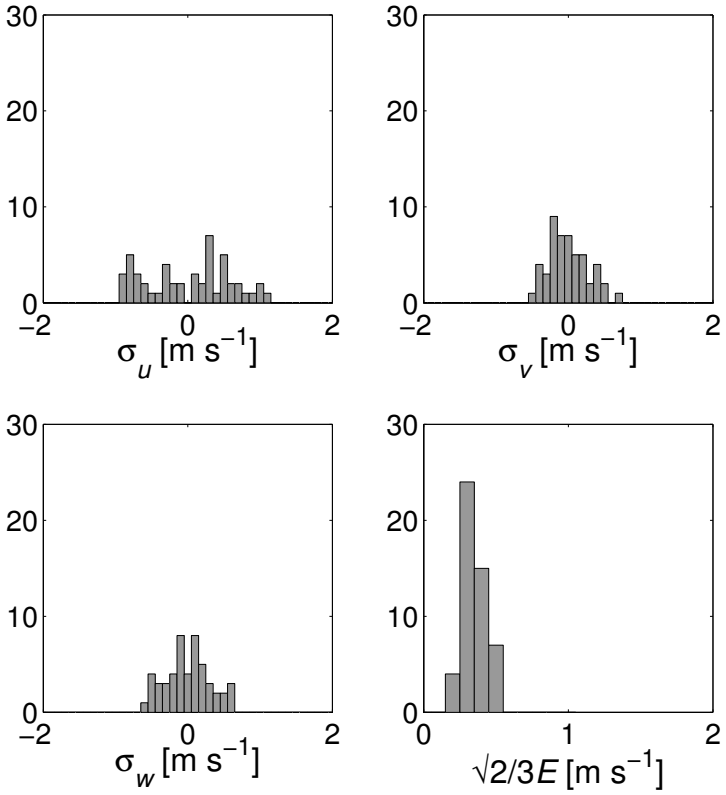

Figure 6.11: The magnitude of the variation in the resolved wind components and the contribution from sub-grid turbulence kinetic energy, at the height of the lowest range gate (∼25 m).

would allow the inclusion of contributions from E in the radial velocity, which would go some way to compensating for the linear variation in radial velocity over a range gate.

While examining the RMS difference between the moment- and spectrally-derived values of u, v, and w provides information about the error caused by the spectral processing technique itself, it does not shed any light on the issue of errors due to the spatial continuity assumption inherent in the DBS technique. To investigate this issue, the time-height plots of each of the wind components (illustrated in Fig. 6.5) can be examined. The top row in Fig. 6.5 shows the wind components taken from the LES column directly above the location of the sodar simulator. These winds correspond to what would be recorded by anemometers on an instrumented tower exactly collocated with the simulated sodar, and as such do not contain any error from the spatial continuity assumption. The simulated anemometer measurements use the LES wind components exactly, i.e., no instrument error is assumed. The middle row of Fig. 6.5 shows the corresponding wind components derived from the sodar moment simulator. These winds are derived from the radial velocities which were calculated by simply averaging the radial velocity values within each range gate. As such, differences between the top and middle rows of Fig. 6.5 can be ascribed to the fact that the DBS technique assumes spatial homogeneity of the wind components over the area encompassed by the off-vertical sodar beams (an assumption which was shown to be invalid for this case study in Fig. 6.1). The bottom row of Fig. 6.5 shows the values of u, v, and w derived from the spectrally processed radial velocities. It is clear that there is a greater difference between the spectrally-processed wind components and those from the LES directly, than between the direct LES and moment simulator-derived wind components. Comparing the top and bottom rows of Fig. 6.5 directly illustrates the full discrepancy arising from both the spectral processing technique and the invalidity of the spatial homogeneity assumption within the DBS technique.

It should be when analyzing Fig. 6.5 that the reported wind components are for instant timing (i.e., updated every 15 s, with no averaging performed), while most operational sodars use an averaging period of at least 5 minutes to reduce error caused by the high temporal variability of the acoustic background noise.

It should be noted here that even for a wind field exhibiting total spatial homogeneity in the horizontal plane, if the wind components change magnitude with height, there would still exist very minor differences between the winds taken from the LES column directly, and those derived from the radial velocity in the moment simulator. The cause of these minor differences is that when using the DBS technique, the resulting winds are ascribed to a particular height range, usually corresponding to the height of the vertical beam at that range gate. However, while the radial velocities used to calculate the wind at a particular height are all equidistant to the sodar in *range*, the radial velocities from the off-vertical beams contain contributions only from the height $H\cos\theta$, not the height H, to which the resulting wind components are ascribed. However, the differences caused by this height offset are typically minor, since the heights H and $H\cos\theta$ are often separated less than the depth contained within a single range gate. For example, for a zenith angle of 15°, at a height of 100 m the difference between H and $H\cos\theta$ will be 3.4 m, at $H = 300$ m the difference will be 10.2 m, and for $H = 500$ m, the difference is 17.0 m.

Fig. 6.12 illustrates the returned power from the vertical beam in the sodar simulator, calculated from the 10-minute averaged spectra. This power is calculated as the peak value of a sum of all the instantaneous frequency spectra contained within the averaging period. The decrease in power with range due to the $1/r^2$ term in (4.10) is clearly seen. During the first 20 minutes of simulation, power values are elevated close to the boundary-layer height due to increased spatial variation in temperature there.

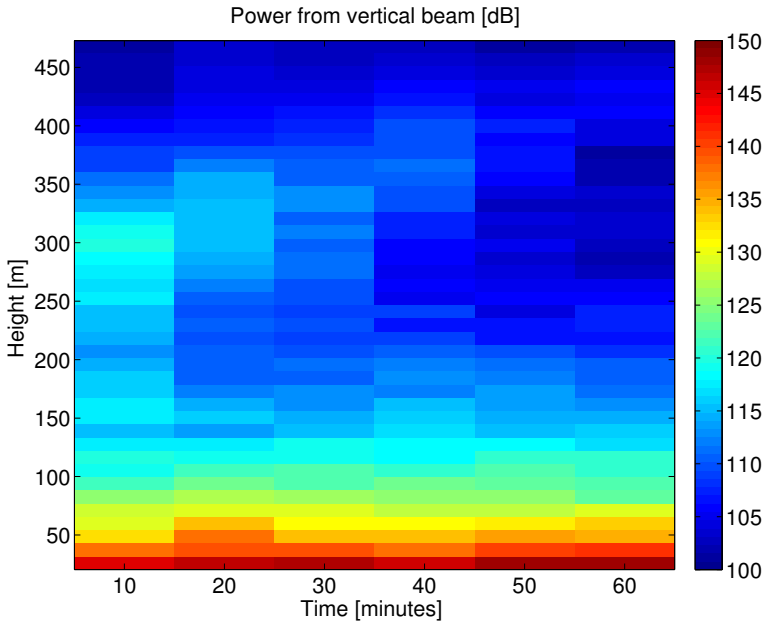

Figure 6.12: Power from the vertical beam of the sodar simulator for the 31 May 2009 case with an arbitrary set zero power value.

To further evaluate the sodar simulator, a comparison is made between the wind speed and direction derived from the time-series sodar simulator, the moment simulator, and simulated tower measurements. Figure 6.13 illustrates the wind speed and direction that would be calculated at five different heights (40, 100, 20, 300 and 400 m) corresponding to measurement levels on a virtual instrumented (i.e., fitted with ideal anemometers) tower whose base is exactly collocated with the simulated sodar (i.e., for a subset of the winds shown in the top row of Fig. 6.5). The simulated wind measurements from the anemometers are taken from the LES data in the desired location, with no additional instrument error or background noise included. The left panels of Fig. 6.13 show results from the moment simulator, and the differences between the moment simulator and virtual tower data can be interpreted as representative of the differences caused by point versus spatially averaged measurements. Clearly, the time series simulator (the right panels in Fig 6.13) produces greater outliers than the moment simulator in both wind speed and direction. This is further evidenced by the lower R^2 values in both wind speed and direction from the time series simulator

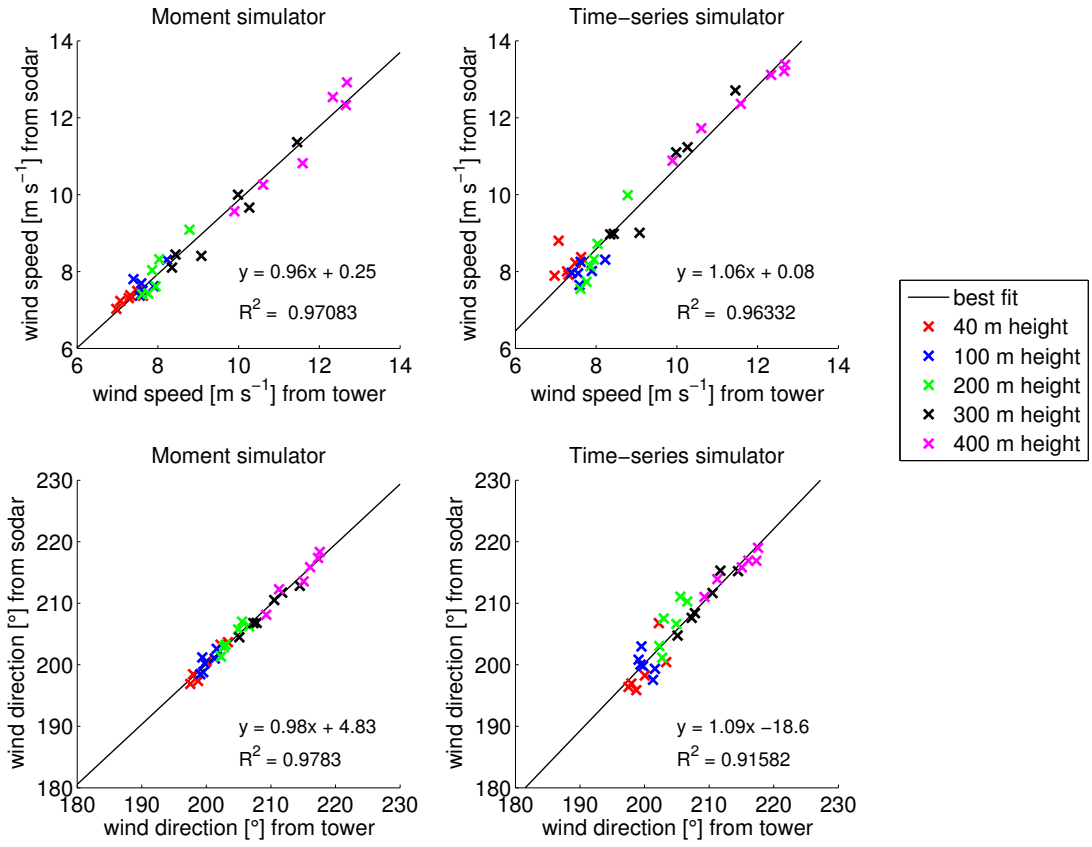

Figure 6.13: A comparison of the derived 10-minute averaged wind speeds and directions from a virtual instrumented tower, the moment simulator, and the time series simulator. The virtual tower data refer to five different heights, while the moment and time series simulator are spatially averaged.

than the moment simulator (shown on Fig. 6.13). The upper right panel of Fig. 6.13 indicates that the time series sodar simulator typically overestimates the wind speed as compared to the virtual tower, where the wind speed values are evaluated directly from the LES. Errors in the spectrally estimated radial velocity caused by the assumption of stationarity of the signal may also be contributing to this effect. The wind speed generally increases with height, and as discussed above, a linear increase in radial velocity across a range gate can result in an increase in the spectrally derived radial velocity.

6.1.2 24 April 2013

6.1.2.1 Experiment description

The second case study also uses output from the OU-LES, and represents another clear CBL. This case study used LES output from 23-24 UTC on 24 April 2013. The LES domain was centered upon the Kessler Atmospheric and Ecological Field Station (KAEFS) near Purcell, OK. The simulation was performed using a variable timestep based upon numerical stability, with an average timestep during the second hour of simulation (which is the simulation period used in this case study) of 0.4 s. The domain size was $(X \times Y \times Z) = 640$ m $\times$ 640 m $\times$ 3 km, with a uniform grid spacing of $\Delta x = \Delta y = \Delta z = 5$ m. The lowest numerical grid level was located at 2.5 m above ground. The employed sub-grid turbulence closure scheme is based upon Deardorff (1980).

The LES was nudged with output from a Weather Research and Forecasting (WRF, Skamarock et al. 2008) model run with a 4-km horizontal grid spacing, centered over Oklahoma. The WRF model was initialized using data from the North American Regional Reanalysis (NARR, Mesinger et al. 2006) output. Horizontal mean profiles of the zonal and meridional winds u and v, the specific humidity q, and the potential temperature θ were taken from the WRF output every minute. These

profiles were interpolated in time, and used to nudge the corresponding prognostic LES fields at every time step as explained in Gibbs et al. (2011).

The simulated time period of 23-24 UTC on 24 April 2013 and location of the KAEFS site were chosen as the LES domain and time frame since an intensive observation period was being carried out at the KAEFS site during this time. A Metek PCS.2000 sodar was in operation, and a UAS (see Section 5.2.1 for details) performed several flights at the site during this time period. The operational sodar is located approximately 150 m from the center of the UAS flight path. The second UAS flight occurred from 2324 to 2355 UTC (1624 to 1655 local time), and used a flight path mimicking the example in Fig. 5.2, used for examining the spatial variation of temperature. The experiment (and simulated) time window covers the onset of the evening transition regime, in which the turbulence in the boundary layer was decaying. Further details about the UAS used in this experiment, the flights performed during this experiment, and the calculation of C_T^2 from the UAS data can be found in Bonin et al. (2014) and Wainwright et al. (2014a).

The presence of the sodar and the UAS at the KAEFS site during the simulated time period allows for comparisons between the simulated and measured data. This is the first case study in which the sodar simulator has been used on LES representing a case during which an operational sodar was present, which enables a comparison between the measured and simulated sodar outputs. The presence of the UAS enables another method of calculating C_T^2, as described in Section 5.3.3. The LES corresponding to the time of the UAS flight also enables use of the UAS simulator (Section 5.3.2), such that comparisons can be made between the measured and simulated UAS output fields.

A sample of the LES output CBL flow is shown in Figs. 6.14 and 6.15. It can be seen that at the time of the simulation, the CBL was well developed and the depth of the boundary layer is at least 500 m (Fig. 6.15), so the simulated sodar will not

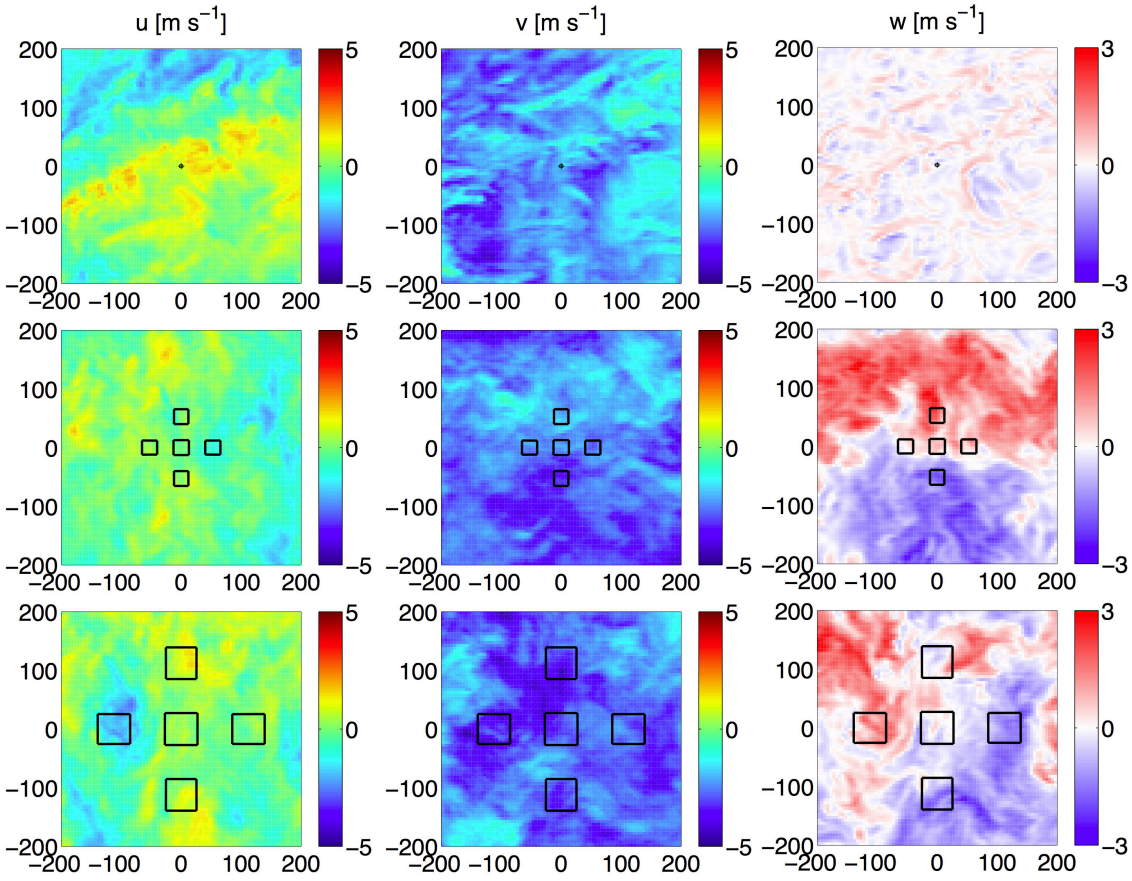

Figure 6.14: Snapshots of the LES wind component fields at 20 m (top row), 200 m (middle row) and 400 m (bottom row) above ground level for the 24 April 2013 LES output. Wind fields shown are from 30 minutes into the 60 minute simulation used to populate the sodar simulator. The areas encompassed by each sodar beam at that height are overlaid.

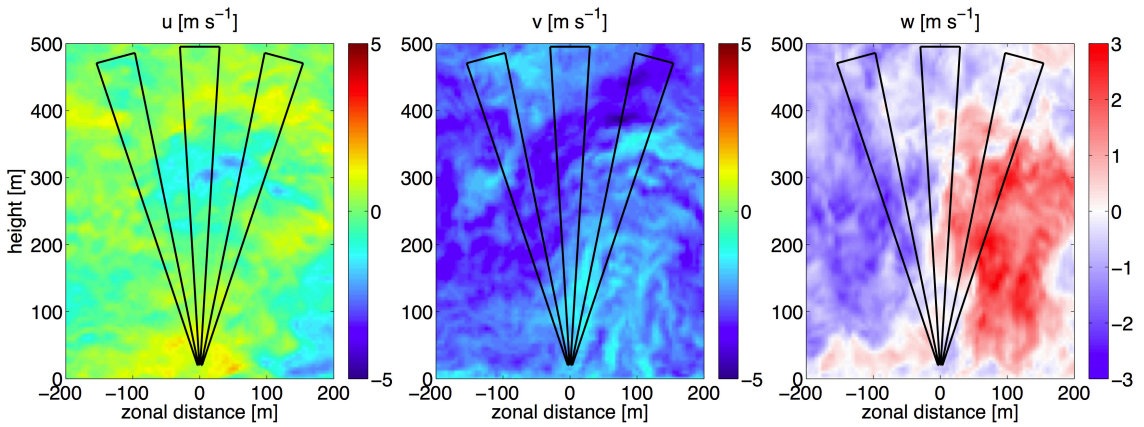

Figure 6.15: Snapshots of the LES wind component fields at $x = 0$ m for the 24 April 2013 case. The areas encompassed by three of the five beams are overlaid. As for the 31 May case, the presented data illustrate the failure of the assumption of homogeneity in the wind component fields across the volume defined by the beam locations.

encompass the full depth of the boundary layer. Throughout the boundary layer, we see a pattern of large scale updrafts and downdrafts with vales of w of up to $\pm$ 3 m s^{-1}, indicative of a fully developed CBL. It can be clearly seen from both Fig. 6.15 and the right panel on the center row of Fig. 6.14 that the assumption of homogeneity of the vertical velocity across the area encompassed by the five sodar beams is invalid.

6.1.2.2 Results

Figure 6.16 illustrates the RMS discrepancy between the radial velocities derived from the moment and spectrally processed sodar simulators. As with the 31 May 2009 case, we see that there is considerable error in the instantaneous derived radial velocities, with average differences of around 0.75 m s^{-1} for each of the beams. It is also clear from Fig. 6.16 that a longer averaging period results in a decreased RMS difference in the radial velocities between the moment and spectrally processed data, as we would expect. This is in contrast to the previous case study, in which the longest averaging period (60 minutes) did not show the smallest RMS difference in the radial velocity values. One possible reason for this is that in the previous case study, the boundary layer was growing rapidly over the course of the simulation (Fig. 6.5), while

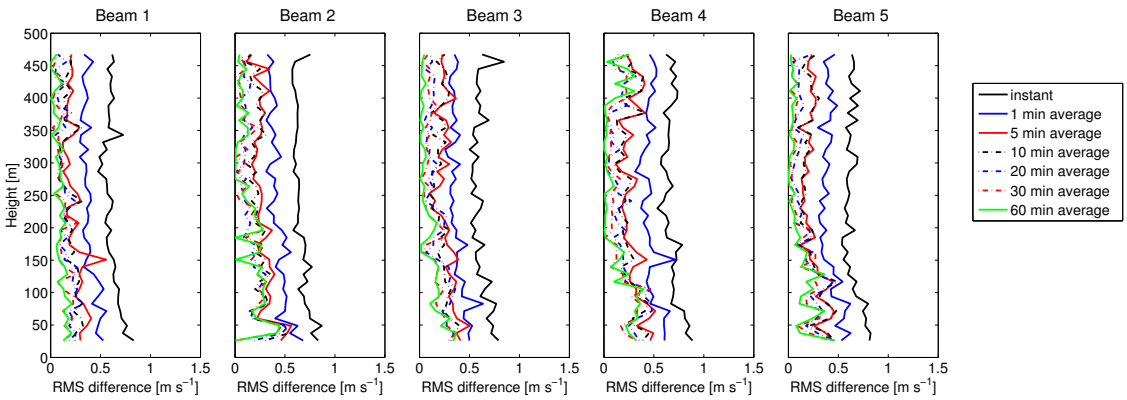

Figure 6.16: RMS difference in radial velocity between the moment and time-series sodar similar for each of the five beams for the 24 April 2013 case.

in the current case study, the boundary layer remains more static over the course of the simulated hour, at least through the 500 m depth used for the sodar simulator (illustrated in the top row of Fig. 6.18). The corresponding RMS differences between the moment and spectrally derived three-dimensional wind components is illustrated in Fig. 6.17. Comparing Fig. 6.17 to Fig. 6.4, we see that the time series sodar simulator produces smaller RMS differences in the instantaneous derived values of u, v, and w for the 24 April 2013 case study than the 31 May 2009 case study.

However, when comparing the top and middle rows of Fig. 6.18, which show the instantaneous values of u, v, and w derived from the moment simulator and those taken from a column in the LES directly above the simulated sodar location, it is clear that there are large differences in the horizontal winds when no averaging period is used. This is also illustrated in Fig. 6.19, which shows a box and whisker plot of the difference between the instantaneous wind components derived from the moment simulator as compared to those taken from an LES column above the simulated sodar location, with the LES column wind components interpolated to the height levels of the simulated sodar range gates. The blue boxes in Fig. 6.19 represent the inter-quartile range of the difference in u (left panel), v (middle panel), and w (right panel), and the red line in the center of the box represents the median error, while the red crosses represent outlying error values. Considering that the errors illustrated in

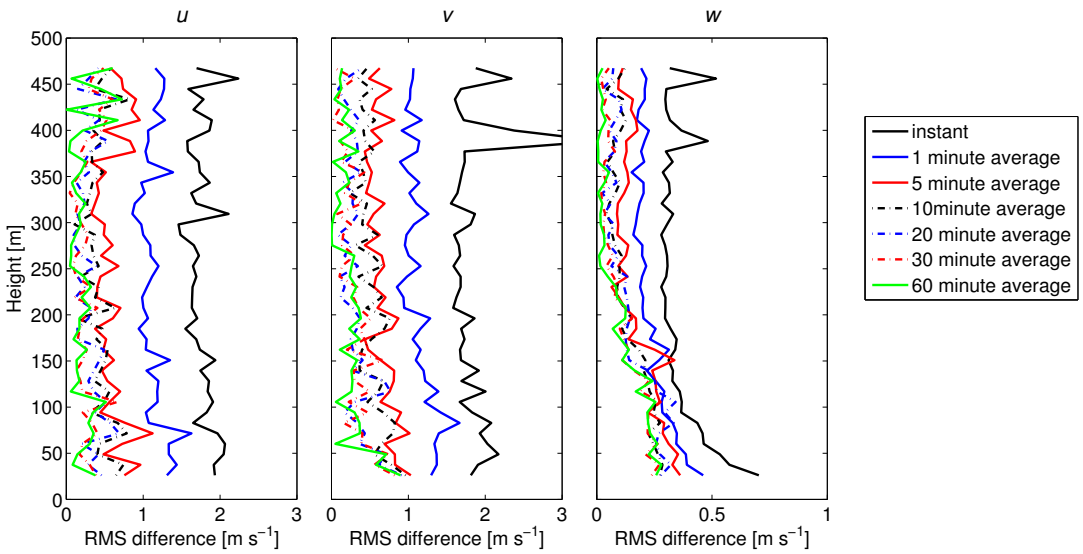

Figure 6.17: Three dimensional wind components as derived from the sodar simulator for the 24 April 2013 case.

Fig. 6.19 do not include contributions due to spectral processing method, the range of errors is high. Just by qualitatively comparing Fig. 6.18 to the corresponding figure from the previous case study (Fig. 6.5), it is apparent that the instantaneous errors differences in u, v, and w between the LES column values and those from the moment simulator are much higher for the current case study.

The wind speed differences shown in Fig. 6.19 theoretically only include errors due to the assumption of spatial homogeneity in the wind field across the area encompassed by the sodar beams. The invalidity of the assumption of spatial homogeneity across the sodar beams was illustrated in Fig. 6.14. However, the previous case study also showed similar invalidity of this assumption (see Fig. 6.1), without such large differences between the winds resulting from the moment simulator and the LES column. One possible cause of this could be the larger variation in the vertical velocity w across the area encompassed by the sodar beams (note the increased scale of the color axis for w in Fig. 6.15 as compared to in Fig. 6.2). When using the DBS technique, w is calculated from the vertical beam only, and then u and v are calculated from the off-vertical beam radial velocities, with the factor of $w \cos\theta$ removed. Thus any

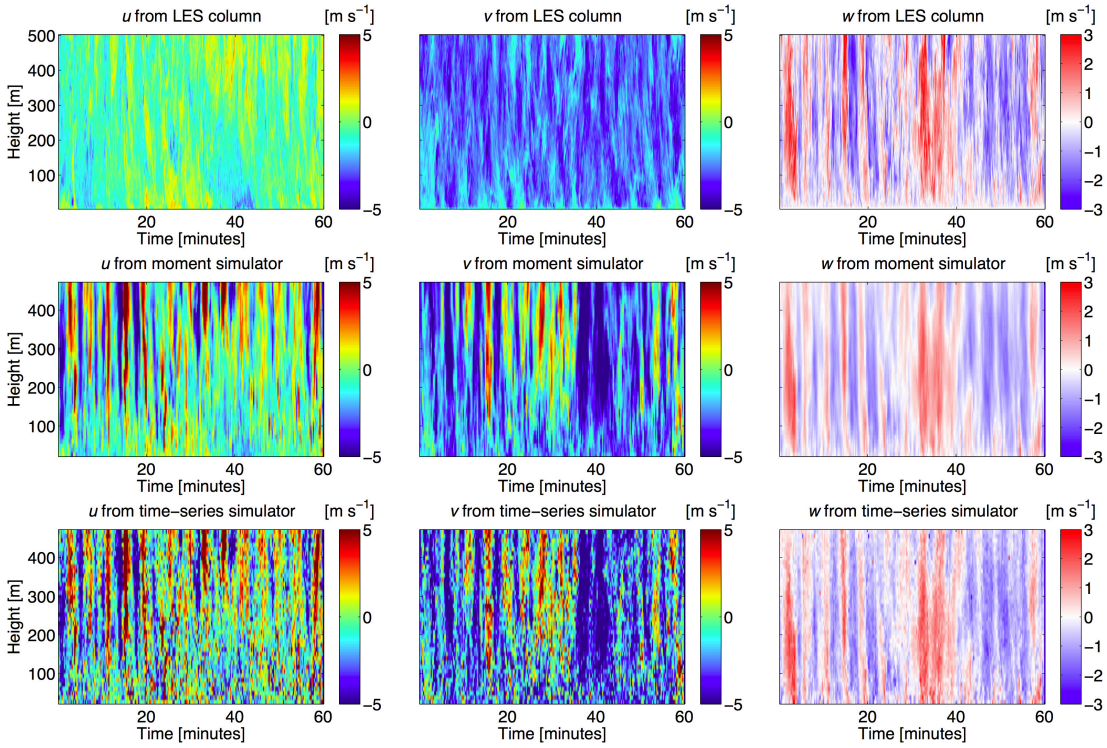

Figure 6.18: A comparison of the three-dimensional wind components u (left column), v (central column) and w (right column) as derived from the LES column directly above the simulated sodar (top row), derived from the moment simulator (middle row) and from the spectrally-processed data from the time-series simulator (bottom row) for the 24 April 2013 case.

error in the estimate of w is propagated forward into both u and v estimates. This variation of the vertical velocity on short temporal and spatial scales is one of the main reasons why sodar manufacturers tend to recommend that users operate the sodar with an averaging period of at least 10 minutes during convective conditions.

The wind components derived from an LES column, the moment simulator, and the spectrally-process radial velocity data for an averaging period of 10 minutes are shown in Fig. 6.20. Comparing Fig. 6.20 to the instantaneous derived wind components in Fig. 6.18, it is clear that for the longer 10 minute averaging period, the wind component fields from the three methods show more similarity. For a 10 minute averaging period, the differences between the moment and spectrally-derived u, v, and

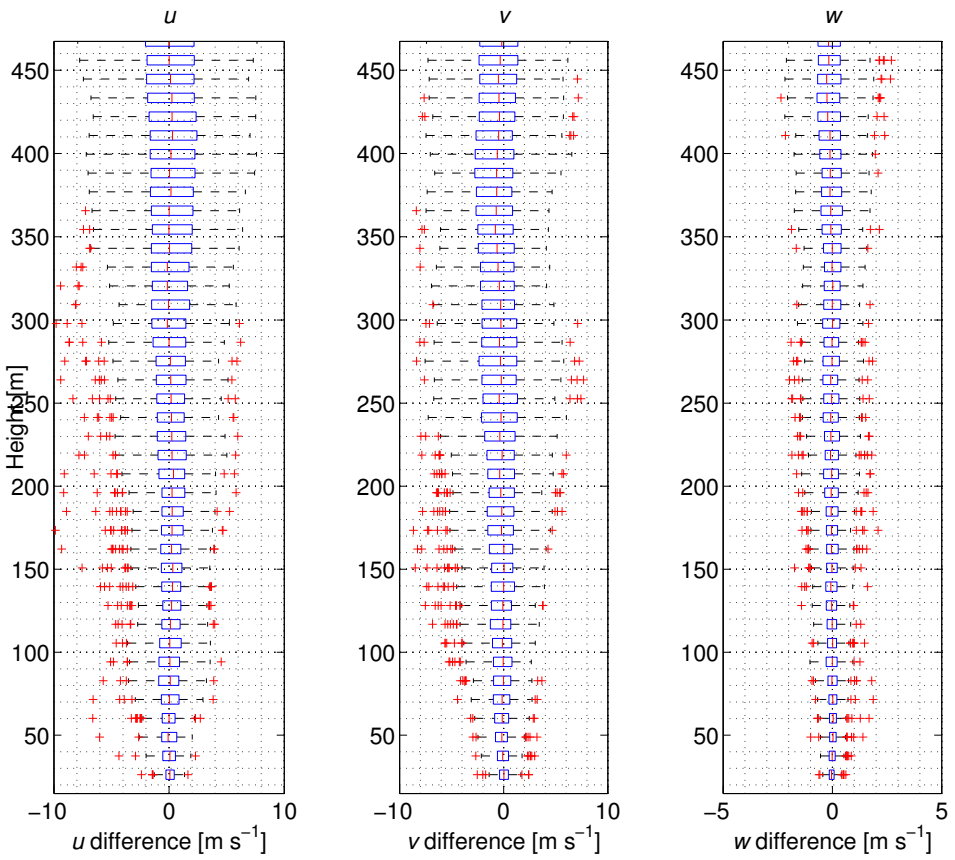

Figure 6.19: A box and whisker plot showing the error in the instantaneous wind components derived from the moment simulator.

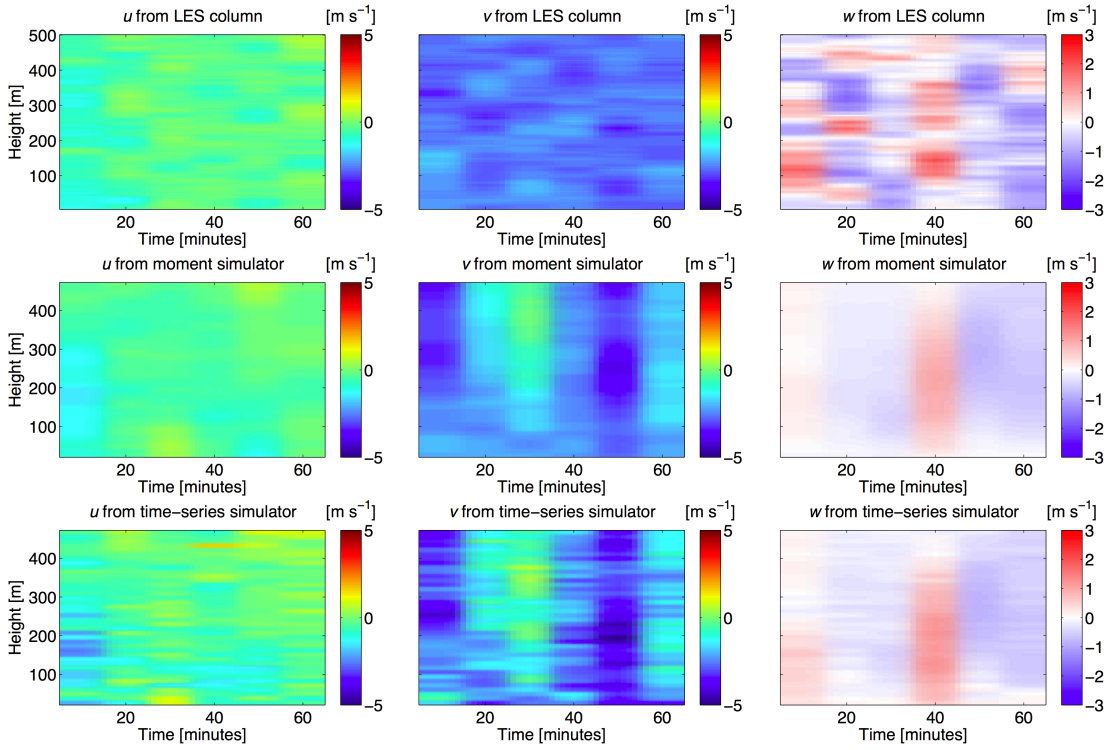

Figure 6.20: As for Figure 6.18, but for an averaging period of 10 minutes.

w (shown in the lower two rows of Fig. 6.20), which is due to the spectral processing technique used by the simulated sodar, are fairly minimal. This can be evaluated quantitatively by the green line in Fig. 6.17, which shows that the RMS difference between the moment and spectrally derived u and v values is below 0.5 m s^{-1} for all range gates in u, and for all except the lowest three range gates for v. The RMS difference in w is below 0.5 m s^{-1} for all heights, and below 0.1 m s^{-1} for range gates above the height of 200 m.

In order to investigate the discrepancies between the LES column and moment simulator radial velocities, caused by use of the DBS technique on heterogeneous wind fields, the experiment was repeated. In the second experiment, the three-dimensional vertical velocity field was set to zero everywhere. This was done to allow for an examination of the effect of the non-uniform vertical velocity on the simulator results. The resulting instantaneous wind components from the LES column, moment simulator,

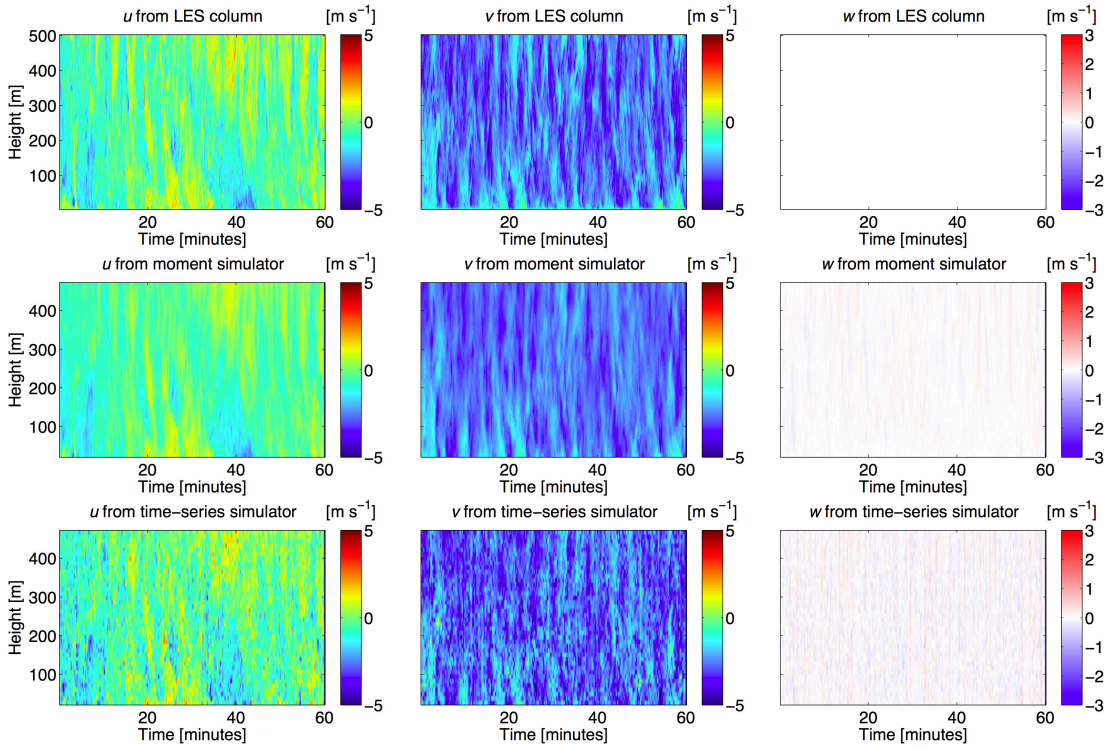

Figure 6.21: A comparison of the three-dimensional wind components u (left column), v (central column) and w (right column) as derived from the LES column directly above the simulated sodar (top row), derived from the moment simulator (middle row) and from the spectrally-processed data from the time-series simulator (bottom row) for the 24 April 2013 case with the vertical velocity set to zero.

and time-series simulator are shown in Fig. 6.21. It is clear from Fig. 6.21 that the errors in the wind components due to the use of the DBS technique are significantly decreased when the vertical velocity is set to zero (comparing Fig. 6.21 with Fig. 6.18). This suggests that the errors due to the use of the DBS technique are largely due to the non-uniform vertical velocity, and effects from nonuniform horizontal wind components are secondary.

The removal of the non-uniform vertical velocity also decreased the discrepancy between the wind components derived from the moment and time-series simulators. This is illustrated in Fig. 6.22, which shows the discrepancy between the wind components derived from the moment and time series simulators for a range of averaging

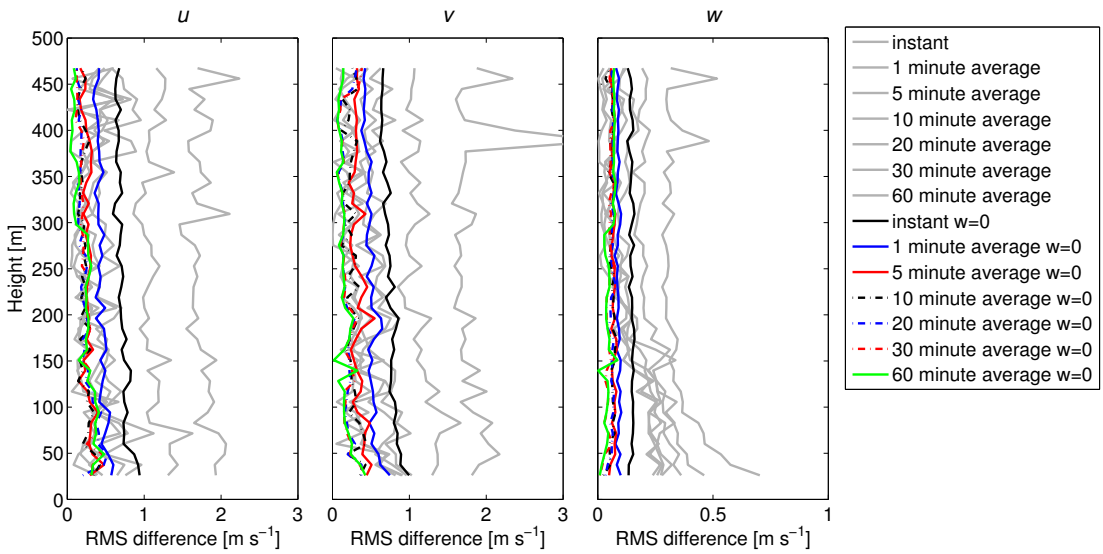

Figure 6.22: Three dimensional wind components as derived from the sodar simulator for the 24 April 2013 case when w was set to zero for a range of averaging periods. The wind component estimates from the non-zero w case are represented by the gray lines.

periods. The corresponding discrepancies from the initial experiment in which w was not set to zero are illustrated by the gray lines. It is clear that the RMS discrepancies in the wind components are significantly reduced when the nonuniform vertical velocity is removed.

As in the previous case study, the initial sodar simulator experiment was repeated with the inclusion of the sub-grid turbulence kinetic energy with the resolved three-dimensional wind components as in equations (6.2), (6.3) and (6.4). The resulting radial velocity RMS differences for each beam are shown in Fig. 6.23. By comparing Fig. 6.23 with Fig. 6.16, the difference in the radial velocity discrepancies caused solely by the inclusion of sub-grid TKE in the three-dimensional wind components can be examined. Similarly, the effect of the inclusion of sub-grid TKE on the derived wind components can be examined by comparing Figs. 6.24 and 6.17. When comparing the two sets of figures, it should be noted that the scales on the x-axis representing radial velocity (or wind speed) are quite different. In the original simulated radial velocities from each beam, shown in Fig. 6.16, the average RMS difference between the instant (i.e., 3 s) radial velocities derived from the moment and spectrally processed data

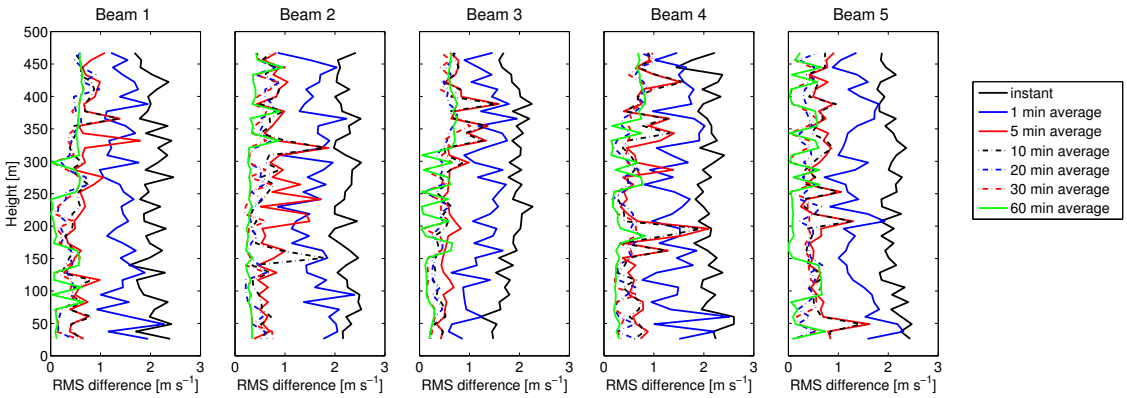

Figure 6.23: RMS difference in radial velocity between the moment and time-series sodar similar for each of the five beams for the 24 April 2013 case when TKE is included as in equations (6.2)-(6.4).

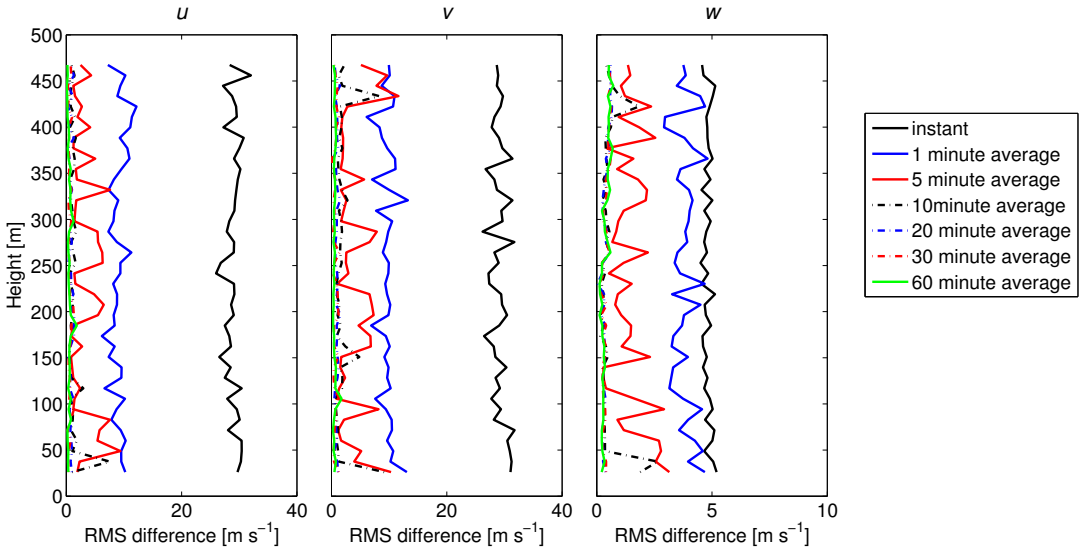

Figure 6.24: Three dimensional wind components as derived from the sodar simulator for the 24 April 2013 case when TKE is included as in equations 6.2-6.4.

was approximately 0.75 m s^{-1}, while the corresponding RMS difference when sub-grid TKE was included was between 2 - 2.5 m s^{-1} for each beam.

The higher RMS differences in the radial velocity from each beam when sub-grid TKE is included in the three-dimensional wind components propagates forward into the wind components derived from these radial velocities. For example, before including sub-grid TKE into $\hat{u}$, $\hat{v}$, and $\hat{w}$, the average error in the 3 s derived values of u across all heights and times was approximately 2 m s^{-1}, and that for w was approximately 0.5 m s^{-1}. However, when the sub-grid TKE was included, i.e., when $\hat{u}$ replaced u in the derivation of the radial velocities in the simulator, the average RMS difference in the instantaneous value of u increased from 2 m s^{-1} to over 30 m s^{-1}, and the corresponding average RMS difference in w increased from 0.5 m s^{-1} to 5.5 m s^{-1}. Not until the averaging period reached 20 minutes did the mean RMS difference in the derived u and v profiles drop below 1 m s^{-1}. For each the three wind components, with an averaging period of 60 minutes, the average RMS difference due to the spectral processing technique was approximately 0.5 m s^{-1}. The relative magnitude of the variance of the horizontal wind components, σ_u, σ_v, σ_w, and the contribution from the sub-grid turbulence kinetic energy, $\sqrt{2/3E}$, are shown in Fig. 6.25 for reference. Figure 6.25 shows the variance in u, v, and w, and the contribution from the sub-grid turbulence kinetic energy calculated from a subset of LES data of size 5 $\times$ 5 $\times$ 2 grid points which surrounds the lowest simulated range gate. It is clear from Fig. 6.25 that the contribution from TKE is approximately equal to σ_u, the variance in u, and σ_v, the variance in v.

Further investigation into this issue evaluated the relative energy contributions from both the resolved wind components and sub-gird TKE. It was found that for this case study, at the height range shown of the lowest range gate, the resolved energy for the LES subset illustrated in Fig. 6.25, calculated as $\frac{1}{2}(\sigma_u^2 + \sigma_v^2 + \sigma_w^2)$ is 0.28 m^2 s^{-2}. The corresponding sub-grid TKE, calculated as $\frac{2}{3}E$ is 0.09 m^2 s^{-2}. This

indicates that the resolved TKE is approximately 3 times higher than the sub-grid TKE. The addition of the sub-grid TKE component to the resolved wind components increases the variance in $\hat{u}$, $\hat{v}$, and $\hat{w}$, but not enough to create the 30 m s^{-1} errors in the instantaneous wind components that were seen in Fig. 6.24. The addition of sub-grid TKE to the resolved velocity components was not conducted at the LES grid spacing, rather E was interpolated to the beam point locations, and then the $\epsilon\sqrt{2/3E}$ contribution was added to u, v, and w. It is this manner of interpolation which caused the large errors in the derived wind components as seen in Fig. 6.24. Since the sub-grid contribution to the wind components was evaluated separately for each beam point (i.e., a new value of ϵ was used for each beam point), this introduced variation into the wind components at a much higher resolution than the values of E were calculated for (the beam point grid spacing is approximately 0.018 m, while the LES grid spacing at which E is calculated is 5 m). However, if the $\epsilon\sqrt{2/3E}$ contributions were calculated at the LES grid spacing and then interpolated and added to the resolved wind components at the beam point spacing, then this would act to introduce yet more linear variations into the radial velocity, due to tri-linear interpolation.

Aside from considering the winds derived from the sodar simulator, the simulator performance can also be evaluated in terms of the returned power of the sodar and associated variables. As described in Section 5.3.4, the simulated sodar power return can be used to estimate C_T^2. As described perviously, a Metek PCS.2000 sodar was also in operation at the KAEFS facility which the LES output of this case represents. The field experiment at the KAEFS facility also included flights by a UAS. One of the UAS flights fell within the time period of the simulation. The flight path used for these measurememnts, designed to mimic that in Fig. 5.2, is shown in Fig. 6.26. Since the UAS can also be used to derive C_T^2 as described in Section 5.3.3, this provides another C_T^2 dataset against which C_T^2 from the sodar simulator can be compared. The

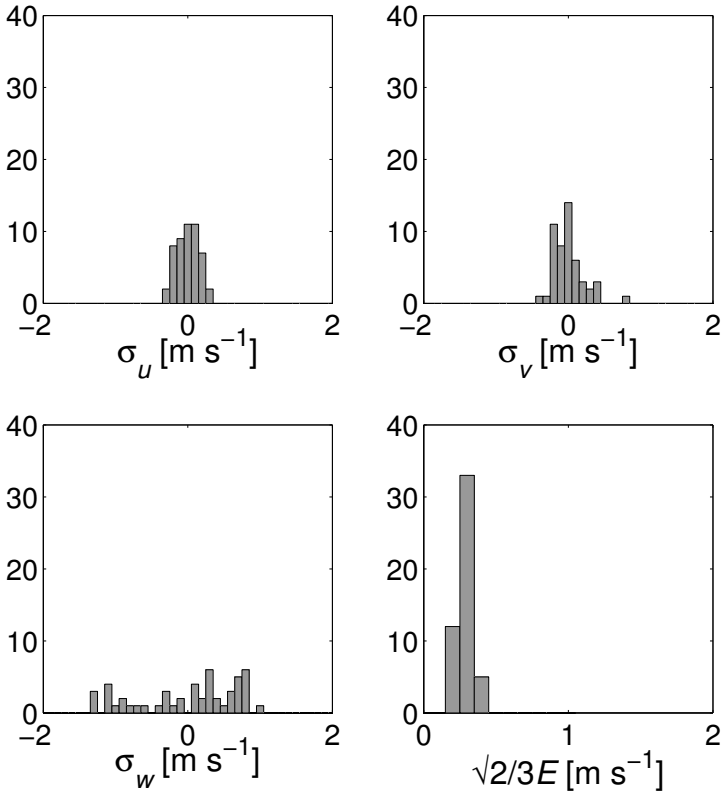

Figure 6.25: The magnitude of the variation in the resolved wind components and the contribution from sub-grid turbulence kinetic energy, at the height of the lowest range gate (∼25 m) for the 24 April 2013 case.

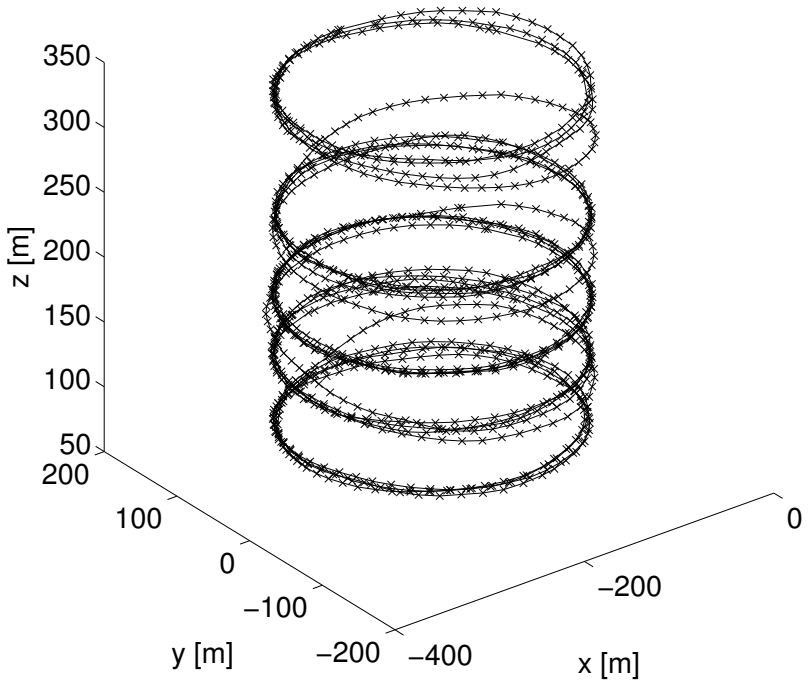

Figure 6.26: The flight path of the UAS for its second flight on 24 April 2013.

UAS simulator was populated with data from the LES output corresponding exactly to the time of the UAS flight, and the resulting simulated UAS data can be used to evaluate C_T^2 as in Section 5.3.2.

Before comparing C_T^2 values derived from each of the mentioned methods, the LES output which is used to populate the sodar and UAS simulators must first be compared with the UAS data to see how well the LES is capturing the atmospheric state observed during the field experiment. Figure 6.27 shows how the temperature recorded by the UAS during the field experiment compares with that recorded by the simulated UAS, which was populated with LES output which simulates the conditions at the KAEFS facility during the flight period. It is apparent that the LES produces a temperature field which is approximately 3°C warmer at all heights than the temperature profile recorded by the UAS. Investigations into the cause of the

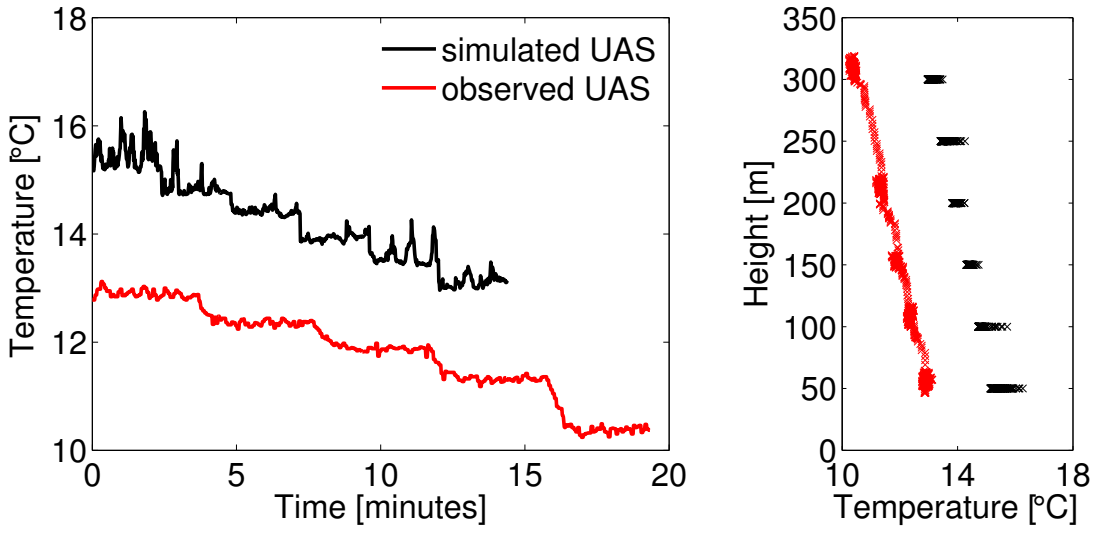

Figure 6.27: The variation of temperature recorded by the simulated UAS (black) and operational UAS (red). The left panel shows how the temperature varied in time and the right panel shows variation with height.

3°C temperature discrepancy found that the WRF simulation was 3°C warmer than the observed temperature during the field experiment, as evaluated from both temperature recordings by the UAS and measurements from the Washington Mesonet station (Brock et al. 1995), which is also located at the KAEFS facility. Despite the warmer simulated temperature, the change of temperature with height is seen to be roughly equal between the simulated and measured temperatures. It should be noted that the temperature field produced by the simulated UAS in the right panel of Fig. 6.27 does not show any data at heights between the 50 m intervals at which the UAS circled, since the simulated UAS can move from each measurement height to the next instantaneously, while the operational UAS contains data recorded while the UAS platform ascends from each measurement height to the next. The relative variation in temperature with height between the simulated and observed data can be examined qualitatively by studying Fig. 6.27. The right panel of the figure illustrates the temperature spread at each height, which appears to be greater in the simulated temperature measurements. This is confirmed by the left panel of the figure, in which

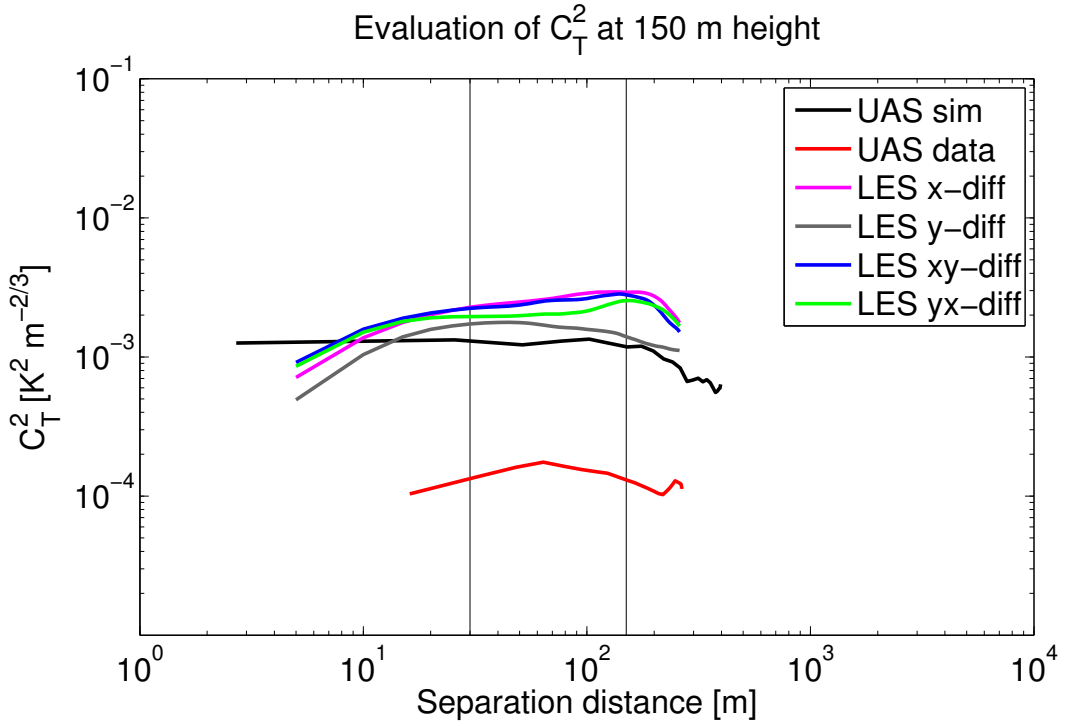

Figure 6.28: The variation of C_T^2 with separation distance from the simulated and measured UAS data, and evaluated from the LES directly in four directions. The vertical black lines indicate the bounds of the inertial subrange for turbulence.

it is also clear that the variation of temperature at each height is greater in the simulated data (the heights can be visually delineated in the left panel of Fig. 6.27 by the drop in average temperature). The left panel of Fig. 6.27 also shows that the UAS flight spent a longer portion of time at each height than the simulated UAS, close to 3 minutes (with additional time required to ascend between circling heights), while the simulated UAS remained at each height for 2.3 minutes.

The variation of temperature between the LES and the observed temperature is next evaluated quantitatively through the calculation of C_T^2. When C_T^2 is calculated from the UAS data, the simulated UAS data, and directly from the LES, the variation of C_T^2 with separation distance can be evaluated. This is illustrated in Fig. 6.28, which shows the variation of C_T^2 with separation distance for a height of 150 m when calculated from measured and simulated UAS data, and from the LES data directly. The four lines corresponding to C_T^2 calculated from the LES directly (magenta, gray,

blue, and green lines in Fig 6.28) represent C_T^2 evaluated in four separate directions: x, y, and in the positive and negative diagonal directions, as in equations (5.5) - (5.8). Evaluating C_T^2 in four separate directions allows for an examination of whether the turbulence is isotopic, in which case the C_T^2 values should be similar for each evaluated direction. It can also be examined whether or not the turbulence is within the inertial subrange, as for turbulence within this range, there should be no variation of C_T^2 with separation distance. Examining Fig. 6.28, it can be seen that there is a range within which the variation of C_T^2 with separation distance is minimal (marked by the two vertical black lines), which suggests that the inertial subrange for turbulence within this simulation extends from 30 m to 150 m, which equates to 6Δ - 30Δ, where Δ is the grid spacing used in the LES. The variation of C_T^2 evaluated in the four separate distances within the inertial subrange is seen to be relatively small. The C_T^2 evaluated in the y-direction are slightly smaller than C_T^2 evaluated in the other three directions at the upper end of the inertial subrange, but the variation is within 1.5×10^{-3} K^2 m$^{-2/3}$, so it should not have a significant impact on further results.

The C_T^2 values evaluated from the UAS simulator are close to those evaluated from the LES directly (Fig. 6.28), which is as expected since the UAS simulator is operating on the same LES data, just parsed via a different method. It should be noted that each of the C_T^2 values evaluated from the LES directly were calculated using the full hour of simulation time, while the simulated UAS remained at the height of 150 m for only 2.3 minutes.

Figure 6.28 also confirms what was noted qualitatively above, that the variation of temperature recorded by the UAS in the atmosphere is considerably lower than the simulated temperature field. The value of C_T^2 within the internal subrange calculated from the UAS (using the method outlined in Section 5.3.3) is 1×10^{-4} K^2 m$^{-2/3}$, approximately an order of magnitude smaller than that from any of the methods performed on the LES output. Another possible cause of the lower values of C_T^2

measured by the UAS could be the 8 s time constant of the temperature sensor on the UAS. The temperature sensor in the UAS records the temperature of the air as it passes the sensor, although the 8 s time constant and the inertia of the plane would allow for some mixing of the air. This effectively mixes out some of the temperature maxima and minima, which results in less variation in the recorded temperature measurements, and acts to lower the value of C_T^2. Previous experiments conducted by van den Kroonenberg et al. (2012) which measured spatially averaged C_T^2 from a UAV over a heterogeneous land surface on a midsummer morning found C_T^2 values varying from 1×10^{-2} - 8×10^{-2}, which is higher than the UAS-estimated C_T^2 values for the present case study. However, the UAS measurements presented in the van den Kroonenberg et al. (2012) study were taken over a much larger spatial area than in the current study, and the recorded temperature were on average 20°C warmer than in the current study.

Vertical profiles of C_T^2 calculated from the LES directly, from the simulated UAS measurements, and from the data recorded by the UAS and the sodar during the field experiment at the KAEFS facility shown in Fig. 6.29. The C_T^2 values from the real and simulated UAS measurements and the direct LES evaluation methods are all presented for a separation distance of 50 m, which is within the inertial subrange for turbulence, as seen in Fig. 6.29. When C_T^2 profiles are derived from sodar data, the separation distance dependence cannot be discerned, so the C_T^2 profile is assumed to be valid for separation distances which are within the inertial subrange. It is seen that the four profiles evaluated from the LES directly, in the x, y, positive and negative diagonal directions, all produce very similar vertical profiles of C_T^2. There is slight variation between the profiles evaluated from the LES directly (the magenta, green, blue, and gray lines in Fig. 6.29) and the UAS simulator (the black line), which is due to the different time periods over which C_T^2 is averaged. The C_T^2 estimates evaluated from the LES data directly are calculated using an averaging period of one hour of

simulation time, while the UAS simulator spends only 2.3 minutes at each height of interest.

It is also clear from Fig. 6.29 that the C_T^2 profile derived from the measured sodar data matches well with the C_T^2 profile derived from the measured UAS data. The C_T^2 profile derived from the sodar data is calculated using a temporal average of the sodar power return (which is recorded with a 5 minute averaging period) for the twenty minutes of the flight time, while the UAS only spends approximately three minutes at each height of interest. When calculating C_T^2 from the UAS data, the data recorded while the UAS platform ascended from one height to the next is excluded, so that only data recorded at the desired measurement heights is presented. As described in Section 5.3.5, in order to calculate C_T^2 from sodar data, a temperature profile is required. For the sodar-derived C_T^2 data shown in Fig. 6.29, the temperature profile was provided by UAS measurements. For other cases, in which the sodar is running but the UAS is not in operation, the temperature profile can be substituted with the temperature measured at the base of the sodar (used by the sodar to calculate the speed of sound, and thus assign the returned signal to range gates) can be used, and the assumption made that the temperature profile is isothermal. This would introduce a source of error in the C_T^2 estimates, with the size of the error depending upon how close to isothermal the temperature profile is.

However, there is a considerable difference between the C_T^2 profiles derived from the real data (both UAS and sodar) and those evaluated from the LES-based techniques. The real data C_T^2 and the LES-derived profiles do not match closely in terms of the magnitude of the C_T^2 values. This corresponds with Fig. 6.28, which showed much lower C_T^2 values derived from the UAS data. However, a large part of the discrepancy between C_T^2 derived from the simulated and observed data is likely due to the performance of the LES. One potential reason for the disparity in the C_T^2 values could be that local effects are not adequately captured by the relatively coarse

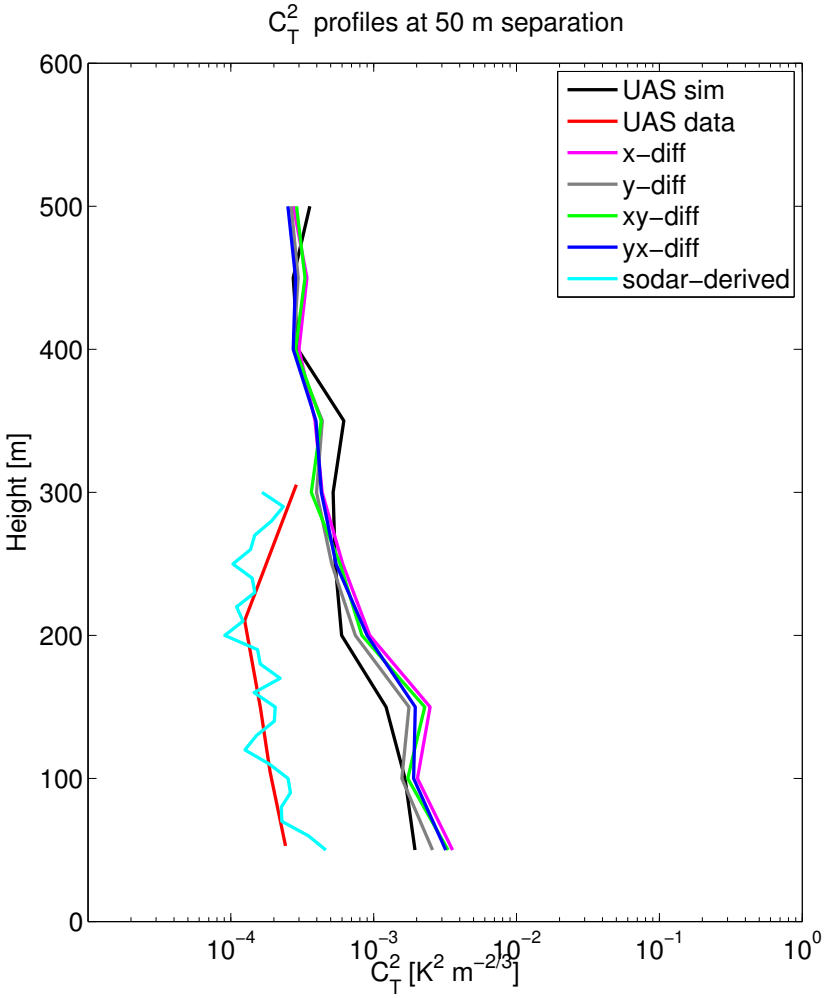

Figure 6.29: Vertical profiles of C_T^2 from the simulated and measured UAS data, evaluated from the LES directly in four directions, and calculated from sodar observations.

nudging data provided by the WRF model. In order to facilitate a comparison of the relative atmospheric stability during the field experiment and in the LES output, the bulk Richardson number, R_{iB} was calculated from both datasets as

$$R_{iB} = \frac{g\Delta\overline{\theta_v}\Delta z}{\theta_v(\Delta\overline{U}^2)} \tag{6.5}$$

where g is gravitational acceleration, $\Delta\overline{\theta_v}$ is the change in virtual potential temperature over a height of Δz, and $\Delta\overline{U}$ is the change in horizontal wind speed over the same height range, 50 - 300 m, from both the LES and the UAS data to ensure a fair comparison. The resulting values of R_{iB} are -9.85 when derived from the LES output and -5.47 when derived from the UAS measurements. This illustrates that the LES model is producing a more unstable atmosphere than indicated by the UA measurements. The increased instability in the LES may partially explain the discrepancy between the C_T^2 values derived from the LES and UAS, as seen in Figs. 6.28 and 6.29.

The similarity between the C_T^2 profiles calculated using the full LES domain (the magenta, gray, green, and blue lines in Fig. 6.29) and using the simulated UAS measurements (black line in Fig. 6.29) indicate that the method of calculating C_T^2 that is used on the measured and simulated UAS measurements can produce realistic C_T^2 values. The difference between the measured and simulated C_T^2 profiles is due to the temperature measurements recorded by the UAS showing far less variability than the simulated UAS temperature data (shown in Fig. 6.27). As mentioned previously, this may be due in part to the time constant of the temperature sensor on the UAS, which is 8 s. The time constant allows for mixing of the air during the response time of the sensor, which acts to mix out the temperature maxima and minima, reducing the temperature variance, and thus decreasing the measured values of C_T^2. The sodar is uncalibrated, so C_T^2 profiles derived from the observed sodar data should be used to monitor the relative variation of C_T^2 with time or with height, as the magnitude of the C_T^2 values is dependent on the chosen calibration constant (the same calibration constant should be used for all cases utilizing data from the

same sodar). The temperature variation recorded by the UAS may not produce C_T^2 profiles which are as accurate as those produced using LES data, for the reason that C_T^2 values represent temperature variations on a very small scale, and the stated accuracy of the temperature probe used on the UAS flights for the 24 April 2013 field experiment is 0.3 K (Bonin et al. 2014). However, as mentioned above, the simulated UAS measurements produce a realistic C_T^2 profile, and as such, the method of estimating C_T^2 using UAS measurements is sound, but the accuracy of the UAS temperature probe should be taken into account when producing C_T^2 profiles from measured UAS temperature data.

Figure 6.30 shows the average 5-minute sodar returned power profiles for the thirty minutes encompassing the duration of the UAS flight. Note that the color scales on the three power profiles are not the same, although they encompass the same range (40 dB). It is clear from Fig. 6.30 that the range of power values with height is similar for the observed sodar data, the moment simulator, and the time series sodar simulator. The return power observations are considerably more variable in time than either of the simulated power return measurements. This is also evident in the C_T^2 profiles derived from the real sodar data, shown in the right panel of FIg. 6.30, which show more variation in the profiles derived from the observed sodar power (black lines) than from the simulated power from either the moment simulator (red lines) or time series simulator (blue lines). Aside from the increased temporal variation in the C_T^2 values derived from observed data, there is general good agreement between the C_T^2 profiles derived using the three methods. The power values output from the moment simulator do not include any spectral processing, rather the power values calculated for each beam point location using equation (4.10) are simply averaged across the range gate. The power values from the time series sodar simulator are instead calculated using the area under the power spectrum, bounded by the frequencies at which the power decreases from the spectral peak by 6 dB. While some of the temporal variation in

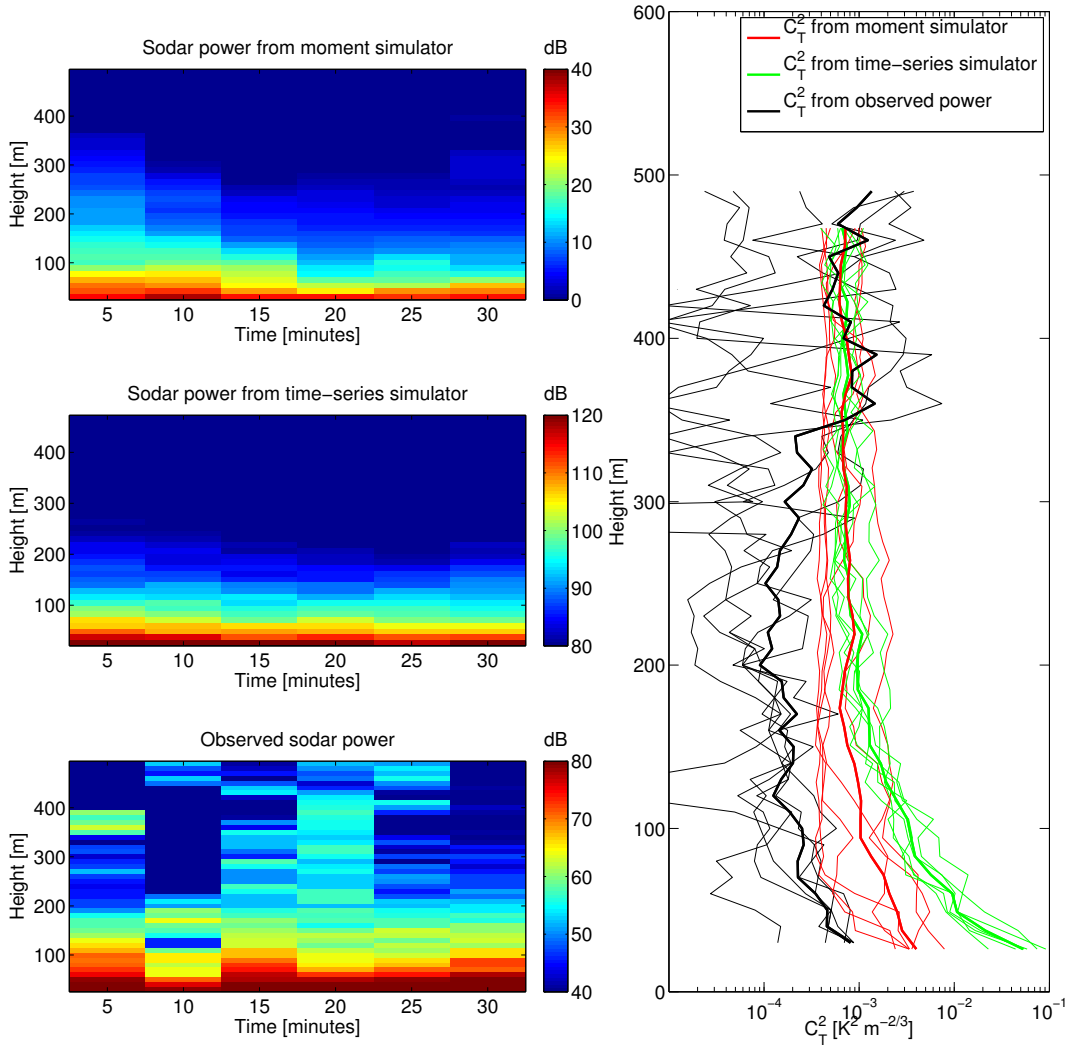

Figure 6.30: Simulated sodar return power (upper left) and observed sodar power (lower left) during the 30 minutes encompassing the time of the UAS flight. The right panel shows C_T^2 values derived from the moment simulator (red), time series simulator (blue) and observed sodar power (black).

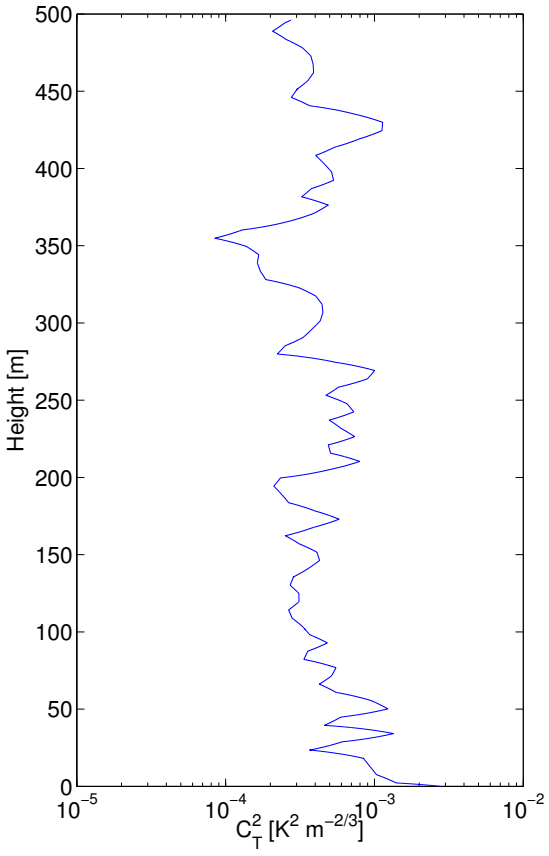

Figure 6.31: The C_T^2 profile calculated as in Fig. 4.4 for a single timestep in the sodar simulator.

the observed power profiles may be due to variation in the signal, it may also be that some of this variation is due to variance in the acoustic background noise. The lack of variation in the simulated background noise is due to the way that noise is included in the simulated sodar data (described in Section 4.2.7). The algorithm used to simulate the background noise requires a scaling factor for the noise, to bring the simulated pink noise to the correct power. The same scaling factor is used for each timestep, which effectively limits the temporal variation in the simulated background noise.

The C_T^2 values are also calculated within the sodar simulator as part of the calculation of the returned power, through equation (4.10). The C_T^2 profile calculated as part of the sodar simulator processing is shown for a single example time in Fig.

6.31. The C_T^2 profile appears to oscillate in height, but this is an artifact of the way that C_T^2 is evaluated in the sodar simulator: one value is calculated for each grid box of the LES, and these values are then tri-linearly interpolated in space and linearly interpolated in time between LES snapshots to match the beam location and timing. Due to the tri-linear spatial interpolation, the C_T^2 profile appears to oscillate at a frequency of twice the LES vertical grid spacing, 10 m for the current case study. The variation of C_T^2 with height shown in Fig. 6.31 compares reasonably well to the C_T^2 profiles calculated from both the power spectrum from the time-series sodar simulator (shown by the blue lines in Fig. 6.30), and the power profile resulting from the moment simulator (red lines in Fig. 6.30).

6.2 Stable boundary layer case study

6.2.1 2 July 2006

6.2.1.1 Experiment description

In order to test how the sodar simulator performs for a stable boundary layer (SBL) case, LES output of the SBL was required. Compared to unstable and neutral boundary layers, large eddy simulations of the stable boundary layer are still a relatively under examined topic. Fewer experiments comparing LES output of the SBL with experimental data have been performed, and it is generally accepted that there are many challenges relating to LES of the SBL which remain unsolved. The first study of a LES of the SBL was reported by Mason and Derbyshire (1990), who found difficulty in sustaining turbulence in the model when plausible initial conditions were used. More recent LES SBL studies identified further problems including unphysical runaway cooling at the surface, excessive sensitivity to the model resolution, and unexpected crashing of the model (Basu et al. 2011). In order to examine the challenges and issues faced in modeling the SBL in the LES format, a large collaborative modeling experiment, GABLS1, was performed in 2003 (Beare et al. 2006; Holtslag 2006).

The GABLS1 experiment was successful as an intercomparison study, but whether or not some of the features seen in the LES results were physically realistic remained unanswered, as the LES output could not be compared directly with observational data.

To further examine the performance of the LES SBL output, a follow-up experiment, GABLS3, was devised. The GABLS3 experiment involved 10 LES modeling groups simulating the same case study. The chosen case study was selected from a multi-year observational dataset (Baas et al. 2008). The chosen case for GABLS3 was a strongly stably stratified nocturnal boundary layer including a low level jet event, observed at the Cabauw tower, in the Netherlands on 1-2 July 2006 (Basu et al. 2011).

The setup of the GABLS3 LES test case was highly specialized, in order to ensure the fairest comparison between all of the participating modeling groups. The required simulation time was nine hours, encompassing 0000 - 0900 UTC on 2 July 2006. The meteorological conditions that were observed during this period include the development of a stably stratified boundary layer, and the transition into the daytime CBL. Initial conditions were provided for the modeling experiment, which included merged data from the instrumented 200 m tower at Cabauw, wind profiling radar data and high-resolution sounding data. Also prescribed were time-height-dependent geostrophic wind forcing data, derived from a surface pressure station network combined with analyzed forecast model data. Time-height-dependent advection of momentum, heat, and moisture derived from mesoscale model forecasts were also prescribed. The lower boundary conditions were prescribed as a provided dataset of potential temperature and specific humidity, extrapolated from near-surface observations taken during the case study period.

Whilst the OU-LES code was not one of those tested during the initial GABLS3 numerical modeling experiment, the LES output used for the current study is based on

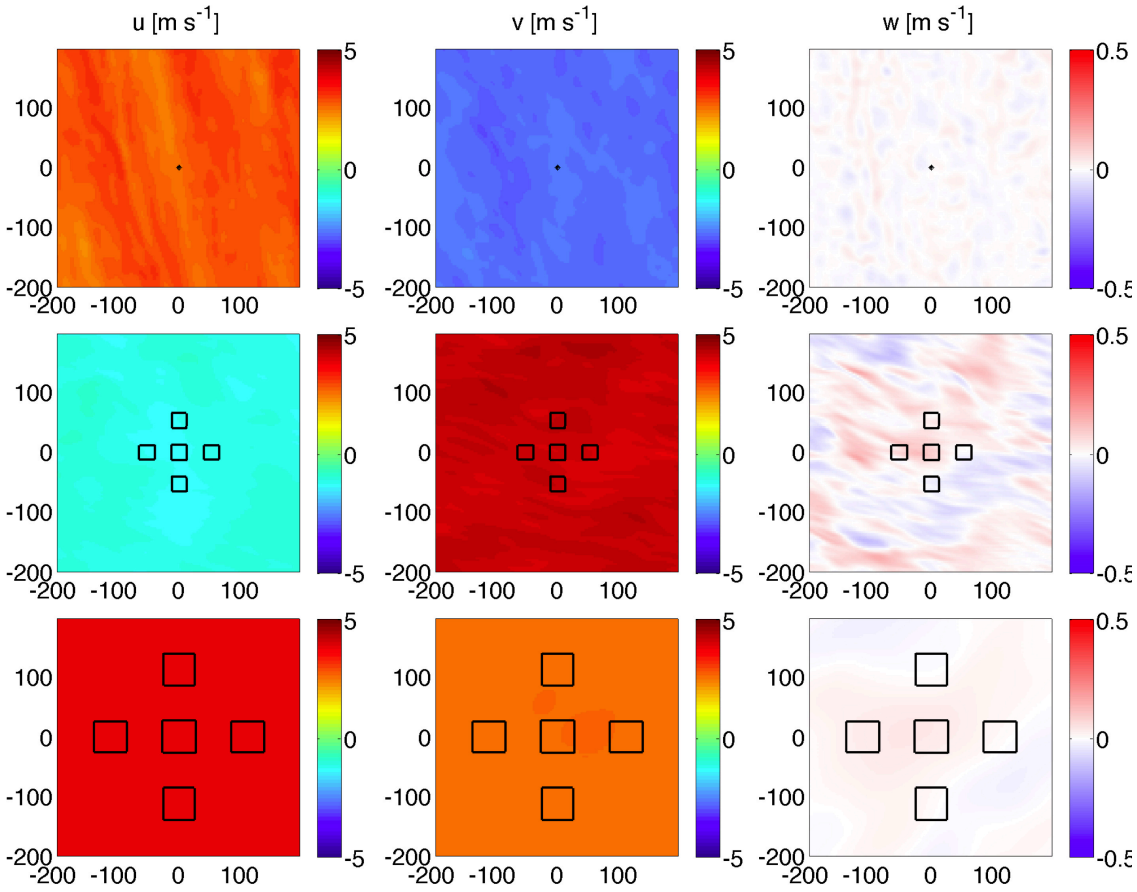

Figure 6.32: Snapshots of the LES wind component fields at 20 m (top row), 200 m (middle row) and 400 m (bottom row) above ground level for the July 2nd 2006 LES output. Wind fields shown are from 30 minutes into the 60 minute simulation used to populate the sodar simulator. The areas encompassed by each sodar beam at that height are overlaid.

the GABLS3 experiment. The OU-LES utilized all the prescribed initial and boundary conditions from the GABLS3 study, with a horizontal and vertical grid spacing of 3.125 m (the required grid spacing for the GABLS3 intercomparison experiment was 6.25 m in both the horizontal and vertical).

Snapshots of the output of the numerical simulation are shown in Figs. 6.32 and 6.33, with the location of the simulated sodar beams overlaid. The strong stability is clear from Fig. 6.32, where the horizontal variation in temperature is seen to be much smaller than for either of the previous cases - compare to Fig. 6.14, which uses the same color scale as Fig. 6.32 for u and v, with a reduced color scale for w.

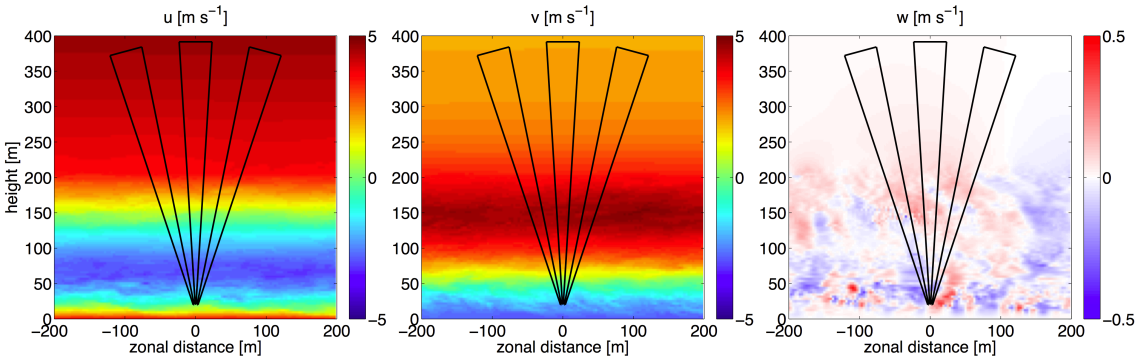

Figure 6.33: Snapshots of the LES wind component fields at $y = 0$ m for the 2 July 2006 case. The areas encompassed by three of the five beams are overlaid.

Examining Fig. 6.33, the boundary layer depth is seen to be approximately 200 m at the onset of the simulation time. Fig. 6.33 also illustrates that while the wind components are almost entirely horizontally homogeneous across the LES domain, the vertical profile of u shows the zonal wind decrease in magnitude from 4 m s^{-1} at the surface to -3 m s^{-1} at a height of 75 m, then increase to return to 4 m s^{-1} by 200 m in height, above which the change in u with height is small. The vertical profile of v shows an approximately linear increase from -3 m s^{-1} at the surface to 5 m s^{-1} at a height of 150 m, above which v steadily decreases with height to reach 2.5 m s^{-1} at 400 m height. The right panel of Fig. 6.33 shows that there are small scale variation in w within the SBL, but on a scale of less than ± 0.25 m s^{-1}, and above the boundary layer, w is close to zero everywhere.

6.2.1.2 Results

The results of the sodar simulator are illustrated in Figs. 6.34 and 6.35. Fig. 6.34 shows the RMS difference between the radial velocity profiles from the moment simulator and the full spectral processing. Across all averaging periods, the RMS difference in the radial velocities between the two methods is considerably smaller than it was for either of the two CBL case studies. The average RMS radial velocity difference

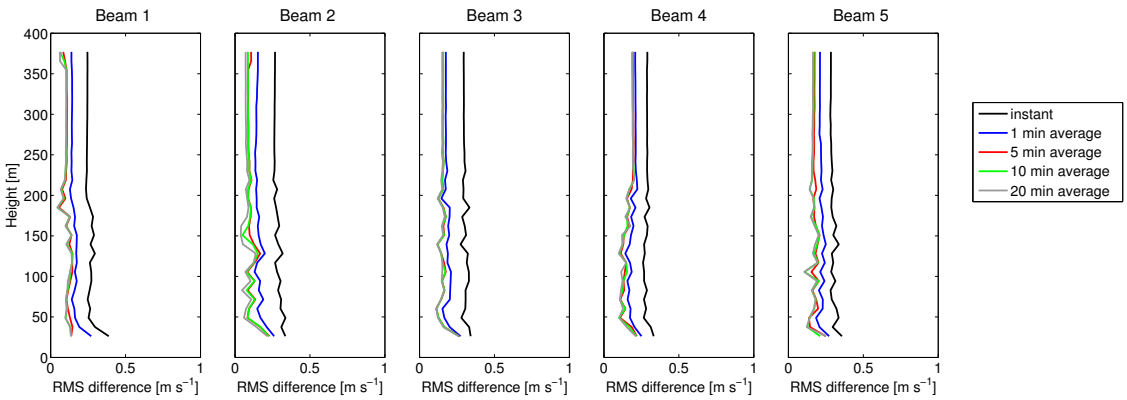

Figure 6.34: Radial velocity for each of the five sodar simulator beams for the 2 July 2006 case.

across all heights and beams for the current stable boundary layer case study is approximately 0.2 m s^{-1}, while it was 0.6 m s^{-1} and 0.8 m s^{-1} for the 31 May 2009 and 24 April 2013 cases, respectively. The smaller RMS radial velocity difference propagates through to the derived three-dimensional wind components, which shows an average RMS difference for the instantaneous wind component estimates between the two methods of 0.6 m s^{-1} for u, 0.7 m s^{-1} for v, and 0.25 m s^{-1} for w. This is in contrast to the higher RMS differences seen for the two convective cases, as the 31 May 2009 case had RMS differences of 1.5 - 2 m s^{-1} for the horizontal wind components and 0.4 m s^{-1} for w, and the 24 April 2013 case showed RMS differences of approximately 2 m s^{-1} for u and v and 0.5 m s^{-1} for w.

The instantaneous wind components derived from an LES column above the simulated sodar, and from the moment and spectrally-processed radial velocities are shown in Fig. 6.36. It should be noted that the temporal scale of Fig. 6.36 is different from the two previous CBL cases, as the SBL case was performed using LES output of a shorter time duration (19 minutes as compared to one hour for the CBL cases). As can be seen in Fig. 6.36, the LES output (top row) shows that the structure of the SBL remained temporally constant during the simulation period. The moment simulator is seen to capture the wind components well (middle row). The higher values of u and higher negative values of v at the surface do not appear in the moment or

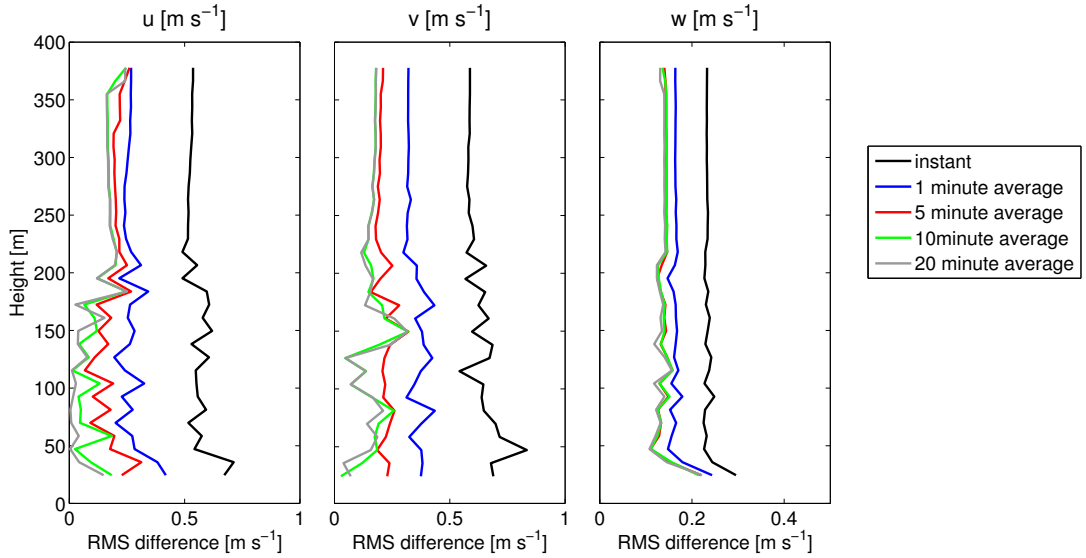

Figure 6.35: Three dimensional wind components as derived from the sodar simulator for the 2 July 2006 case.

time-series simulator results because the minimum range of the sodar simulator is 26 m, while the LES output shown in the top row of Fig. 6.36 is from heights of 3.125 m upwards.

The wind components derived from the spectrally processed radial velocities also capture the profiles of u and v reasonably well, although the bottom right panel of Fig. 6.36 shows that there is considerable temporal variation in the w spectrally-derived w profile, especially at heights above the boundary layer, which was not present in the LES column (top right panel) or the moment-derived w values (center right panel). The spectra from beam 3, the vertical beam, are shown for five consecutive times near the end of the simulation in the upper panel of Fig. 6.37, and it can be seen that there is considerable variation in the frequency of the spectral peak, despite w being close to zero for the LES column during this time period (upper right panel of Fig. 6.36). The resolution of the spectra shown in Fig. 6.37 is 9.3 Hz. The temporal variation seen in the instantaneous w measurements from the spectrally-processed data are relatively constant with height, which is surprising, given the almost constant w field at heights above boundary layer. The spectra generated by range gates between

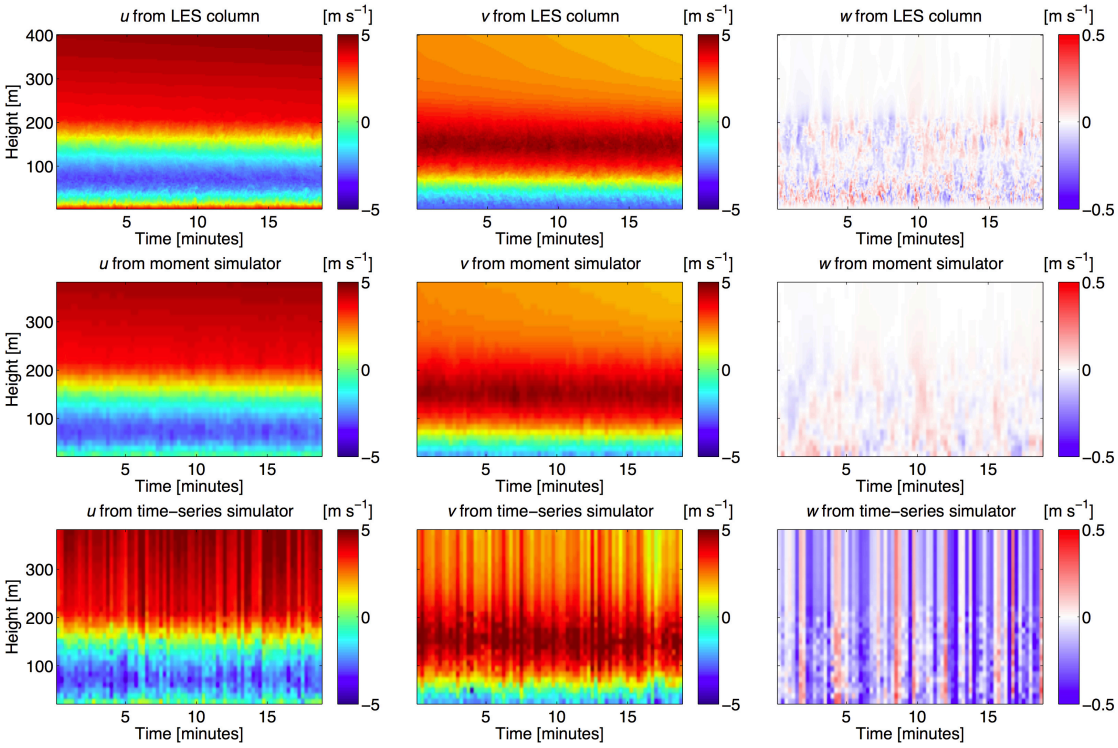

Figure 6.36: A comparison of the three-dimensional wind components u (left column), v (central column) and w (right column) as derived from the LES column directly above the simulated sodar (top row), derived from the moment simulator (middle row) and from the spectrally-processed data from the time-series simulator (bottom row) for the 2 July 2006 case.

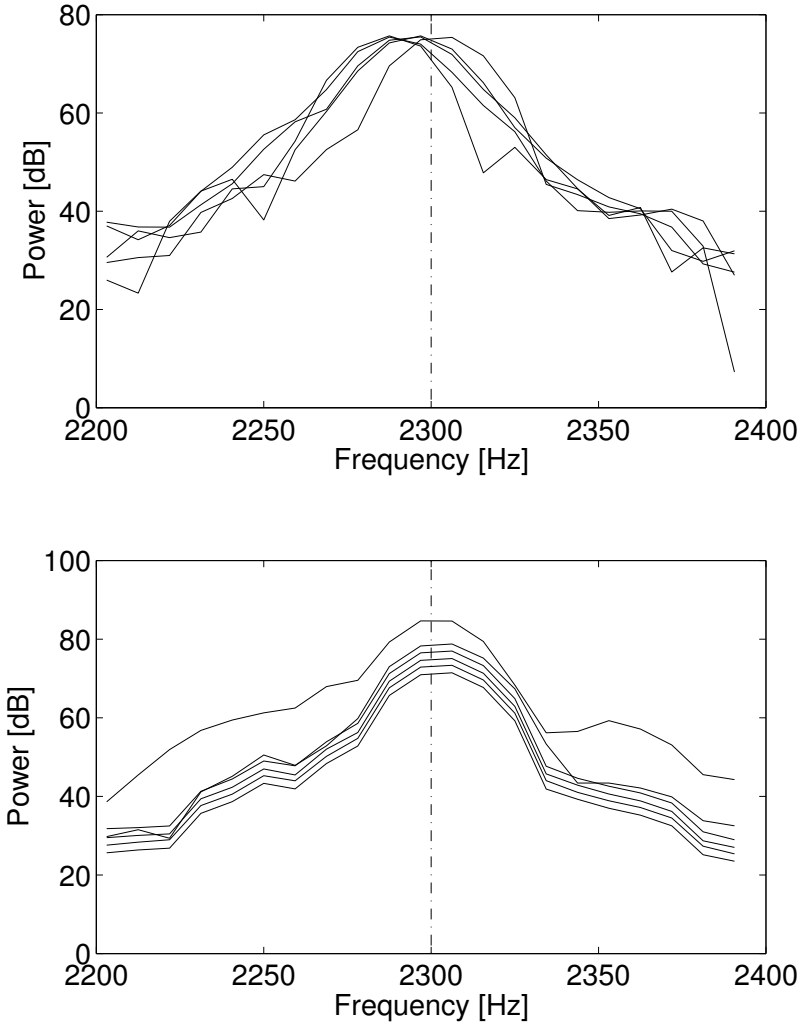

Figure 6.37: Frequency spectra from the vertical beam for the 2 July 2006 case. The upper panel shows five consecutive spectra from a height of 300 m from within the final two minutes of simulation time, and the lower panel shows spectra from several heights between 200 m and 375 m for one of the times included in the the upper panel.

200 - 375 m for a single time, which was one of the times included in the upper panel of Fig. 6.37, is shown in the lower panel. The spectra from the various heights all show a consistent spectral peak, so the relative constancy of w with height for each individual estimate time seen in Fig. 6.36 is consistent with the frequency spectra. Examining the spectrally-derived profiles of u and v in Fig. 6.36, it is clear that the error in the instantaneous w profiles carries through to u and v, which both show wind speeds above the boundary layer which are relatively constant with height, but vary with time more than shown by the moment simulator. This variation in the wind components during the simulation time is illustrated for a single range gate at 300-m height in Fig. 6.38. One possible partial cause of the constancy of w with height in the lower right panel of Fig. 6.36 could be the overlap of height between the range gates. The radial velocity contributions to each range gate are smeared in height due to oversampling in range by the sodar. However, while this may be responsible for some of the constancy with height of the w estimates produced by the time-series simulator, the strength of this effect was not shown in either of the previous two case studies. Although, it is worth noting that both of those case studies had w fields which showed greater magnitudes of variation with time than the stable case presented in this section. This is illustrated by comparing the color scales with which w is shown in Fig. 6.36, where the total magnitude of w varies by ± 0.5 m s^{-1}, as compared to Fig. 6.5 for the 31 May 2009, which showed w values varying between approximately ± 2 m s^{-1}, and Fig. 6.18, in which the w values varied by ± 3 m s^{-1}. As such, it may be that the effect of oversampling in range discussed above could also be present in the results from these case studies, but it is overshadowed by the larger magnitude of the vertical velocity.

Figure 6.38 compares the difference between the variance of u, v, and w with time when evaluated from the moment and time-series simulators. The solid lines in Fig. 6.38 represent the wind components derived from the moment simulator (with

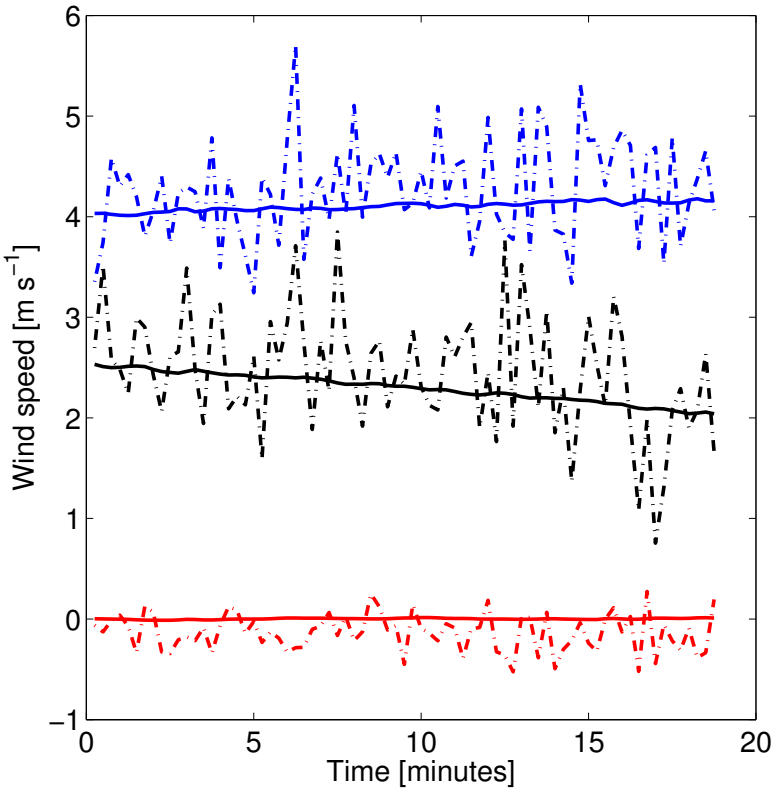

Figure 6.38: The variation with time of instantaneous wind components u (blue), v (black) and w (red) when evaluated from the moment simulator (solid lines) and the spectrally-processed radial velocities (dashed lines) for the range gate at a height of 300 m.

u shown in blue, v shown in black and w shown in red) for the range gate at 300 m height, and for the duration of the simulation. It is seen that there is little variation in the moment-derived wind components with time, with w close to zero at all times, and a slight increase in u and decrease in v over the simulation - with the total variation of either horizontal wind component no higher than 0.25 m s^{-1}. The corresponding wind components derived from the spectrally-processed radial velocities for the same height are shown by the dashed lines in Fig. 6.38. It is clear that the spectrally derived wind components exhibit far greater temporal variation than those from the moment simulator, although the temporal variations from the time-series wind components seem to vary around a central value that is close to the DBS 'truth' captured by the moment simulator.

Further investigation into the cause of the relative constancy of the errors with height, and the high temporal variation of the errors revealed that the primary cause of the errors in the time series simulator were variable speed of sound. As discussed in Section 4.2.2, the spatial location of the beam must be determined prior to interpolation of the LES fields to the calculation point locations. As such, a constant value for the speed of sound, c, is used to determine the spatial location of each calculation point. Once the calculation point locations are determined, the LES fields are interpolated to these calculation points. In the calculation of the received complex acoustic voltage, the t used in $\exp(-j2\pi(f_t + f_s)t)$ was taken as the time since the release of the acoustic pulse. This was calculated as $t = 2r/c$, with r the range of the signal at that time and c the full temperature-dependent speed of sound. The use of the temperature-dependent speed of sound means the value of t will change slightly with each timestep, since r remains constant (since the range of the beam points are calculated using a constant speed of sound). The change of the variable t with time has the same effect on the complex signal as slight temporal variations in the frequency, since the two are multiplied in the signal. In order to test this

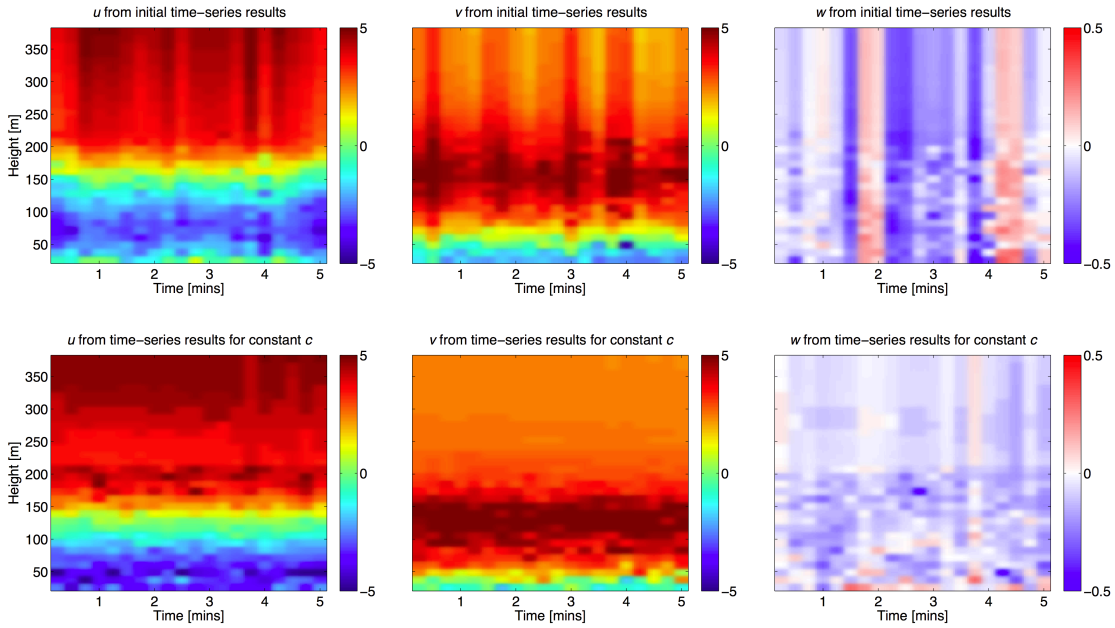

Figure 6.39: A comparison of u, v, and w from the initial time-series simulation (top row) and when forced to use the speed of sound and time values from the first timestep.

suggestion, the first five minutes of the simulation were repeated, with the speed of sound profile used to calculate t held constant following the first simulation timestep. The results of this five minute simulation are illustrated in Fig. 6.39, where the top row represents the wind components from the original simulation, and the bottom row represents the simulation with the same values of t used for every timestep. It is clear from FIg. 6.39 that by forcing t to remain identical at each timstep, the unphysical constancy of w with height but strong variation in time, previously witnessed at heights above 200 m, is reduced. This carries forth into the estimates of u and v, which in the previous simulation contained unphysical temporal variations above the boundary layer. This temporal variation is seen to be mitigated by using the same values of t for each timestep (bottom row of Fig. 6.39).

Chapter 7

Discussion and conclusions

In this study, a sodar simulator was designed which is capable of ingesting numerically simulated turbulent flow fields and reproducing the output parameters that a sodar operating under the atmospheric conditions represented by the LES would produce. The simulator has been developed in a way that many of its operating parameters are defined by the user, in order to facilitate comparison with existing operational systems. To evaluate the results of the time series sodar simulator, a moment simulator has been developed in parallel. The moment simulator uses the same beam locations as the time series sodar simulator, but it does not include the spectral processing. In the moment simulator, the radial velocities contained within a range gate are simply averaged to provide a representative velocity value. The use of the moment simulator as a comparison tool ensures that any differences between winds derived from the moment and time series simulators are due to the sodar measurement and signal processing techniques and not caused by the invalidity of the assumption of horizontal homogeneity of the three-component wind fields across the area encompassed by the sodar beams.

The parallel development of the moment simulator also allows for an examination of the differences in the wind components derived from an LES column collocated with the simulated sodar, and those derived from the moment simulator. The difference between the winds resulting from the moment simulator and those from the LES column act to quantify the discrepancy in the measured wind field caused by the use of the DBS technique despite the invalidity of the assumption of a horizontally homogeneous wind field. The errors caused by the use of DBS in horizontally heterogeneous wind fields is an interesting research topic in itself, and it has implications

for the accuracy and representativeness of wind fields measured by instruments (such as sodars and lidars) which make use of the DBS technique.

The performance of the time series sodar simulator was tested on three distinct case studies using LES output fields of temperature, humidity, horizontal and vertical wind components, and sub-grid turbulence kinetic energy.

The first case study, 31 May 2009, used LES data representative of developing CBL in the morning (0900-1000 CDT). It was seen in Fig. 6.1 that the wind components were not horizontally homogeneous across the area encompassed by the sodar beams, most noticeably for the vertical velocity, w. The development of the boundary layer in the LES data from a depth of 300 m at the start of the simulation time to 400 m after one hour of simulation time can be clear seen in Fig. 6.5. Figure 6.5 also shows good agreement between the winds derived from the moment and time series sodar simulators. Calculation of the average RMS difference in u, v and w between the moment series and time series derived winds for averaging periods ranging from 3-s to 1-hour are shown in Fig. 6.4. The average RMS discrepancy between the wind components derived from the two simulators generally decreased for longer averaging periods, but increased for u and v at the longest averaging period, 1-hour. An investigation into the cause of this higher discrepancy at the 1-hour averaging period revealed that the primary cause was the significant development of the boundary layer over the hour of simulation. The greatest RMS difference between the u and v estimates from the two methods was located between 300 - 350 m, which is the height range through which the top of the boundary layer grew throughout the simulation. It was also noted from Fig. 6.4 that at some heights, the RMS differences in the u and v estimates remained the same for averaging periods greater than 5 minutes, rather than decreasing as would be expected. Investigation into the cause of the systematic bias in the u and v estimates from the time series simulator found that the range gates with the highest RMS differences between the moment and time series simulator were

those over which the radial velocity varies the most with height within the range gate. The simulated sodar, like most commercial sodars, processes the simulated returned complex acoustic signal using an FFT. A fundamental assumption when performing an FFT on a signal is that the signal statistics are wide sense stationary. Violations of this assumption result in the peak of the resulting frequency spectra not being located at the mean frequency content of the signal (as it is if the signal is stationary). This effect was illustrated for a simple example signal in Fig. 6.8. It was surmised that this effect is partially due to the effect of trilinear interpolation introducing an unphysical perfect linear radial velocity variation across the height contained within a range gate. The relatively good agreement in estimated horizontal wind speed and direction between the LES column (i.e., simulated instrumented tower data), the moment simulator, and the time series simulator is shown in Fig. 6.13. It was seen that estimated wind speed and direction from both the moment and time series simulators agree to the simulated instrumented tower measurements with $R^2 > 0.9$ (Fig. 6.13), indicating good agreement between the three wind speed and direction estimates.

One potential way to mitigate the undesirable effect of trilinear interpolation causing linear radial velocity variation is to include the addition of sub-grid turbulence kinetic energy to the resolved three-dimensional wind components. This method was tested on the 31 May 2009 case study and it was found that the addition of sub-grid TKE introduced unphysical large errors into the radial velocity (as in Fig. 6.10).

The second case study tested the sodar simulator on LES data representing 24 April 2013. The hour of LES output used for the test represents the fully developed convective boundary layer at 1800-1900 CDT (local time). As with the previous case study, it was seen that the assumption of horizontally homogeneous flow fields across the spatial area bounded by the sodar beams was invalid (Figs. 6.14 and 6.15). The LES slices shown in Fig. 6.15 also indicate that the boundary layer depth was greater than the height of the LES subset used to populate the sodar simulator.

Unlike the previous case study, the RMS difference between the wind components from the moment and time series simulator were seen to consistently decrease with longer averaging periods (Fig. 6.17). The effect of linear variations in radial velocity causing systematic biases in the radial velocities derived from the time series simulator was not as strong as for the previous case study. The reason that the effect of linear frequency modulation was not as strong in the second case study is that the boundary layer extends through the depth of the LES subset used to populate the simulator. As such, the area encompassed by the simulator contains turbulent variations around a mean wind speed, rather than the strong vertical shear in the horizontal wind fields that exists at the height of the boundary-layer top in the previous case study (compare Figs. 6.2 and 6.15).

The second case study also showed considerable differences between the instantaneous wind components observed in the LES column above the simulated sodar location and those derived from the moment simulator, as illustrated in Fig. 6.18. These differences are solely due to the use of the DBS technique on horizontally inhomogeneous wind fields. Instantaneous horizontal wind component estimates from the moment simulator showed deviations from the corresponding LES column winds of up to 8 m s^{-1} (illustrated in Fig. 6.19). One possible cause of the increased error in the wind component estimates resulting from use of the DBS technique (as compared to the previous case study) was suggested to be increased spatial variation in w as compared to the previous case study (note the variation in the color scales used for w in Figs. 6.1 and 6.14). Since w is calculated from the vertical sodar beam only, the estimated value of w is then used in the derivation of the estimates for u and v. As the mean value of w tends to 0 m s^{-1} for longer averaging periods, it follows that the errors in w caused invalid use of DBS will decrease at longer averaging periods, and the resulting errors in the u and v estimates will also thus decrease. Figure 6.20 shows that this is the case. The addition of the sub-grid turbulence kinetic energy

to the resolved wind components as in equations (6.2)-(6.4) was tested on this case study. Further investigation into the poor results from the inclusion of the sub-grid TKE found that the magnitude of the sub-grid TKE was not the cause of the high errors seen in Figs. 6.16 and 6.17. The additional variation in the wind components resulting from the inclusion of sub-grid TKE was shown to be of similar magnitude to the variance in the resolved wind components over a small area of the LES data ($5 \times 5 \times 2$ grid points). Rather, the large error in the u, v, and w estimates that results from the addition of sub-grid TKE multiplied by a normally distributed random number, applied at scales much smaller than the LES grid spacing. This introduces additional variance into the wind components $\hat{u}$, $\hat{v}$, and $\hat{w}$. While the variance resulting from the addition of TKE is of a physically reasonable magnitude, the addition of this variance at the scale of the sodar beam points is unphysical. The LES grid spacing is 5 m, and fluctuations in u, v, and w resulting from sub-grid TKE are representative of this scale. The addition of this variance at scales of a few centimeters is not physically realistic.

Several methods of calculating C_T^2 were tested on the simulated temperature field from the 24 April 2013 case study, and the associated measurements made by a UAS and sodar during a field experiment at the site represented by the LES data. Comparisons of the temperature measurements recorded by a simulated and operational UAS (shown in Fig. 6.27) revealed that the simulated boundary layer on 24 April 2013 was 3° warmer than indicated by measurements from the UAS flight and the Mesonet station located at the field site. The variation of temperature recorded by the UAS at several heights was also considerably smaller than the simulated temperature variation. This is due in part to the LES representing a more unstable boundary layer than the UAS observation indicate (when stability was evaluated using the bulk Richardson number). It may be partly due to the accuracy of the temperature probe installed on the UAS platform. The method of deriving C_T^2 from UAS measurements

was tested on simulated UAS measurements based upon LES output. It was found that C_T^2 values resulting from the simulated UAS matched well with those evaluated from the LES fields directly, which suggest that the method is sound. It is suggested that future attempts to use UAS observations to estimate C_T^2 are made using a temperature probe with higher sensitivity. The shape of the C_T^2 profile derived from sodar observations matches reasonable well with both the UAS observations and the LES-based methods, although the magnitude of C_T^2 values derived from sodar observations should be used only to examine relative spatial and temporal changes, as the sodar is uncalibrated.

The sodar simulator was also tested on a third case study, using LES data representing the nocturnal stable boundary layer on 2 July 2006 over Cabauw, in the Netherlands. The discrepancy between the LES column wind components and those derived from the sodar simulator is minimal, since the wind field is horizontally homogeneous, as can be seen in Fig. 6.32. The wind components estimated from the time series radial velocities show temporal variation not seen in those estimated by the moment simulator. Additionally, the temporal variation remains relatively constant with height. It is suggested that oversampling in range may be partially responsible for the relative constancy of the error in the retrieved w estimates with height. These errors are due to the spectral processing technique and not the use of DBS, as they are not seen in the w estimate derived from the moment simulator (Fig. 6.36). Investigation showed that these errors resulted from the temporal variation of the time variable in equation 4.11, which was caused by the variation of the speed of sound with temperature, since t in equation 4.11 was calculated through $t = 2r/c$. It is thus suggested that the same constant value for c be used in the calculation of t as is used for the initial beam location calculation.

Overall, the sodar simulator provides a useful tool for the investigation of the way the boundary layer is viewed by acoustic remote sensors. It allows for the emulation

of operational sodar systems and for the testing of difference scanning strategies and signal processing techniques on the same turbulent flow data. It also provides a way to compare wind data that would be obtained by point and spatially averaged measurements in a variety of simulated atmospheric environments.

7.1 Future work

Planned future work on the topic of sodar simulation should include an investigation into an appropriate scaling for the inclusion of spatially interpolated values of resolved velocity and sub-grid turbulence kinetic energy. This should allow for a reduction in the RMS differences in wind components between the moment and spectrally-derived values for u, v, and w. This is one of the most pressing issues that has surfaced from the work presented in this dissertation. The addition of subgrid TKE to the resolved wind components would mitigate the linear frequency modulation caused by trilinear interpolation of the LES radial velocities to the sodar beam points. However, this additional must be performed in a physically realistic manner. The question of interpolating subgrid components resulting from large eddy simulations to finer scales is complex and worthy of further study.

There are several additional features which could be added to the sodar simulator that would allow for additional analysis. The calculation of sodar-estimated vertical velocity variances would provide a means of comparison for observed vertical velocity variance profiles, which are not easily verified in the field without the use of several different remote sensing instruments. This would facilitate a comparison between velocity variances derived using the instantaneous velocity values and those derived using spectrum width. Some form of automated instantaneous rejection of spectra with high noise levels (as is implemented in operational sodars) would also add an interesting facet to the sodar simulator technique, as it would allow for an analysis

of the percentage of data availability at different heights under certain atmospheric conditions.

Future work on the UAS simulator should include the addition of a simulated time constant for the temperature measurements, which would allow for an examination of the effects of the time constant of the temperature sensor on the derived C_T^2 values. The addition of a time constant to the UAS simulator would also allow for a more fair comparison between stimulated and observed UAS C_T^2 profiles.

Overall, the technique of sodar simulation has proven to be a useful tool for the thorough investigation of sodar signal processing techniques and scanning strategies. Results from initial tests on three case studies using LES data to populate the sodar simulator have shown that the simulator can reproduce realistic wind profiles and turbulence intensity estimates. As such, the sodar simulator could be used to test experimental signal processing techniques or algorithms repeatedly on the same dataset. This provides a useful tool for sodar research and development. It is hoped that in the future the time series simulator will be extended to simulate a wider range of sodar output parameters (such as velocity variance profiles and data availability) and text novel signal processing techniques such as testing Wavelet analysis and testing new cluster analysis techniques on the wind speed data.

Bibliography

Anderson, P. S., 2003: Fine-scale structure observed in a stable atmospheric boundary layer by sodar and kite-borne tethersonde. *Boundary-Layer Meteorol.*, **107**, 323–351.

Anderson, P. S., R. S. Ladkin, and I. A. Renfrew, 2005: An autonomous doppler sodar wind profiling system. *J. Atmos. Oceanic Tech.*, **22**, 1309–1325.

Antoniou, I., H. E. Joergensen, F. Ormel, S. Bradley, S. von Hunerbein, S. Emeis, and G. Warmbier, 2003: On the theory of sodar measurement techniques. Rep. RISO-R-1410 (EN), 60 pp.

Argentini, S., G. Mastrantonio, I. Petenko, I. Pietroni, and A. Viola, 2012: Use of a high-resolution sodar to study surface-layer turbulence at night. *Boundary-Layer Meteorol.*, **143**, 177–188.

Baas, P., F. C. Bosveld, H. K. Baltink, and A. A. M. Holtslag, 2008: Towards a third intercomparison case for GABLS using Cabauw data. *Proceedings of the 18th Symposium on Boundary Layers and Turbulence*, American Meteorological Society, Stockholm, Sweden.

Balsley, B. B. and K. Gage, 1982: On the use of radars for operational profiling. *Bull. Amer. Meteorol. Soc.*, **63**, 1009–1018.

Banta, R. M., R. K. Newsom, J. K. Lundquist, Y. L. Pichugina, R. L. Coulter, and L. Mahrt, 2002: Nocturnal low-level jet characteristics over Kansas during CASES-99. *Boundary-Layer Meteorol.*, **105**, 221–252.

Barth, M. F., R. B. Chadwick, and D. W. van de Kamp, 1994: Data processing algorithms used by NOAAs wind profiler demonstration network. *Ann. Geophys.*, **12**, 518–528.

Basu, S., A. A. M. Holtslag, and F. C. Bosveld, 2011: GABLS3-LES intercomparison study. *Proceedings of the ECMWF GABLS workshop on diurnal cycles and the stable boundary layer*.

Batchelor, G. K., 1957: *Symposium on naval hydraulics*, National Academy of Sciences, chapter Wave scattering due to turbulence. 409–423.

Beare, R. J., M. K. MacVear, A. A. M. Holtslag, J. Cuxart, I. Esau, J. C. Golaz, M. A. Jiménez, M. F. Khairoutdinov, B. Kosovic, D. Lewellen, T. S. Lund, J. K. Lundquist, A. McCabe, A. F. Moene, Y. Noh, S. Raasch, and P. Sullivan, 2006: An intercomparison of large-eddy simulations of the stable boundary layer. *Boundary-Layer Meteorol.*, **118**, 247 – 272.

Behrens, P., J. O'Sullivan, R. Archer, and S. Bradley, 2012: Underestimation of mono-static sodar measurements in complex terrain. *Boundary-Layer Meteorol.*, **143**, 97–106.

Benjamin, S. G., K. J. Brundage, and L. L. Morone, 1994: Implementation of the Rapid Update Cycle. part I: Analysis/model description. Technical Report 416, NOAA/NWS Tech. Procedures Bull.

Beyrich, F., 1994: Sodar observations of the stable boundary layer height in relation to the nocturnal low-level jet. *Meteorol. Z.*, **3**, 29–34.

—, 1997: Mixing height estimation from sodar data - a critical discussion. *Atmos. Environ.*, **31**, 3941–3953.

Bingöl, F., J. Mann, and D. Foussekis, 2009: Concially scanning lidar error in complex terrain. *Meteorol. Z.*, **18**, 189–195.

Blokhintzev, D., 1946: The propagation of sound in an inhomogeneous and moving medium I, II. *J. Acoust. Soc. Amer.*, **18**, 322–334.

Bonin, T., P. Chilson, B. Zielke, and E. Fedorovich, 2012: Observations of early evening boundary layer transitions using a small unmanned aerial system. *Boundary-Layer Meteorol.*, **2**, 1–14.

Bonin, T. A., P. B. Chilson, B. S. Zielke, P. M. Klein, and J. R. Leeman, 2013: Comparison and application of wind retrieval algorithms for small unmanned aerial systems. *Geosci. Instrum. Method. Data Syst.*, **2**, 177–187.

Bonin, T. A., D. Goines, A. Scott, C. E. Wainwright, J. A. Gibbs, and P. B. Chilson, 2014: Measuring structure function parameters with a small unmanned aerial system. *Boundary-Layer Meteorol.*, **to be submitted**.

Botnick, A. M. and E. Fedorovich, 2008: Large eddy simulation of atmospheric convective boundary layer with realistic environmental forcings. *Quality and Reliability of Large-Eddy Simulations*, J. Meyers et al., ed., Springer Verlag, 193–204.

Bradley, S., 1999: Use of coded waveforms for sodar systems. *Meteorol. Atmos. Phys.*, **71**, 15–23.

—, 2008a: *Atmospheric Acoustic Remote Sensing*. CRC Press, 271 pp.

—, 2008b: Wind speed errors for lidars and sodars in complex terrain. *IOP Conf. Ser.: Earth Environ. Sci.*, **1**, 1–7.

—, 2012a: The noise part of sodar signal-to-noise. *Proceedings of the 16th Symposium for the Advancement of Boundary-Layer Remote Sensing*, 271–274.

—, 2012b: A simple model for correcting sodar and lidar errors in complex terrain. *J. Atmos. Oceanic Tech.*, **29**, 1717–1722.

—, 2013: Aspects of the correlation between sodar and mast instrument winds. *J. Atmos. Oceanic Tech.*, **30**, 2241–2247.

Bradley, S., I. Antoniou, S. von Hünerbein, D. Kindler, H. E. Joergensen, and M. D. Noord, 2005: SODAR calibration for wind energy applications. Technical report, University of Salford.

Bradley, S., S. von Hünerbein, and T. Mikkelsen, 2012: A bistatic sodar for precision wind profiling in complex terrain. *J. Atmos. Oceanic Tech.*, **29**, 1052–1061.

Brewer, G. A., 1954: Calibrating the PBY-6A airplane for meteorological purpose. *Aero. Eng. Rev.*, **13**, 24–28.

Brock, F. V., K. C. Crawford, R. L. Elliott, G. W. Cuperus, S. J. Stadler, H. L. Johnson, and M. D. Eilts, 1995: Oklahoma Mesonet: A technical overview. *J. Atmos. Ocean. Tech.*, **12**, 5–19.

Busse, J. and K. Knupp, 2012: Observed characteristocs of the afternoon-evening boundary layer transition based on sodar and surface data. *J. Appl. Meteorol. Climat.*, **51**, 571–582.

Capsoni, C., M. D'Amico, and R. Nebuloni, 2001: A multiparameter polarimetric radar simulator. *J. Atmos. Oceanic Tech.*, **18**, 1799–1809.

Chanin, M. L., A. Gariner, A. Hauchecorne, and J. Porteneuve, 1989: A Doppler lidar for measuring winds in the middle atmosphere. *Geophys. Res. Lett.*, **16**, 1273–1276.

Cheong, B. L., M. W. Hoffman, and R. D. Palmer, 2004: Efficient atmospheric simulation for high resolution radar imaging applications. *J. Atmos. Ocean. Tech.*, **21**, 374–378.

Cheong, B. L., R. D. Palmer, and M. Xue, 2008: A time series weather radar simulator based on high-resolution atmospheric models. *J. Atmos. Ocean. Tech.*, **25**, 230–243.

Clothiaux, E., T. Ackerman, and D. Babb, 1996: *Radiation and water in the climate system: remote measurements.*, NATO, chapter Ground-based remote sensing of cloud properties using millimeter-wave radar. 323–366.

Conzemius, R. J., 2004: *The effects of wind shear on convective boundary layer entrainment*. Ph.D. thesis, University of Oklahoma.

Cook, D. R. and M. S. Pekour, 2008: Eddy correlation flux measurement systems handbook. Technical report, DOE Tech. Reo. DOE/SC-ARM/TR-05.

Cooley, J. W. and J. W. Tukey, 1965: An algorithm for machine calculation of complex fourier series. *Math. Comput.*, **19**, 297–301.

Coulter, R. L., 1979: A comparison of three methods for measuring mixing-layer height. *J. Appl. Meteorol.*, **18**, 1495–1499.

Coulter, R. L. and B. L. Li, 1995: A technique using the wavelet transform to identify and isolate coherent structures in the planetary boundary layer. *Proceedings of the 11th Symposium on Boundary Layers and Turbulence*, Charlotte, NC. Preprints, 291–294.

Crescenti, G. H., 1996: A look back on two decades of Doppler sodar comparison studies. *Bull. Amer. Meteorol. Soc.*, **78**, 651–673.

—, 1998: The degradation of Doppler sodar performance due to noise: a review. *Atmos. Environ.*, **32**, 1499–1509.

Dabbert, W. F., R. Shellhorn, H. Cole, A. Paukkunen, J. Hörhammer, and V. Antikainen, 2002: Radiosondes. Technical report, NCAR Earth Observing Laboratory.

Deardorff, J. W., 1980: Stratocumulus-capped mixed layer derived from a three-dimensional model. *Boundary-Layer Meteorol.*, **18**, 495 – 527.

Detwiler, A., J. Scannell, D. Kliche, and S. Williams, 2012: Creating the long-term T-28 instrumented research aircraft data archive. *Bull. Amer. Meteor. Soc.*, **93**, 1817–1820.

Edner, H., P. Ragnarson, S. Svanberg, E. Wallinder, R. Ferrara, R. Cioni, B. Raco, and G. Taddeucci, 1994: Total fluxes of sulfur dioxide from the Italian volcanoes Etna, Stromboli, and Vulcano measured by differential absoption lidar and passive differential optical absoption spectroscopy. *J. Geophys. Res.*, **99**.

Emeis, S., 2010: *Measurement methods in atmospheric sciences: in situ and remote*. Gebrüder Borntraeger.

—, 2011: *Surface-based remote sensing of the atmospheric boundary layer*. Springer.

Emeis, S., K. Schäfer, and C. Münkel, 2008: Surface-based remote sensing of the mixing-layer height - a review. *Meteorol. Z.*, 621–630.

Emeis, S. and M. Türk, 2004: Frequency distributions of the mixing height over and urban area from SODAR data. *Meteorol. Z.*, **13**, 361–367.

Eymard, L. and A. Weill, 1982: Investigations of clear air convective structures in the PBL using a dual Doppler radar and a Doppler sodar. *J. Appl. Meteorol.*, **21**, 1891–1906.

Fedorovich, E., R. Conzemius, I. Esau, F. K. Chow, D. Lewellen, C.-H. Moeng, D. Pino, P. Sullivan, and J. V.-G. de Arellano, 2004a: Entrainment into sheared convective boundary layers as predicted by different large eddy simulation codes. *Preprints, 16th Symp. on Boundary Layers and Turbulence, Amer. Meteor. Soc., 9-13 August, Portland, Maine, USA*, CD–ROM, P4.7.

Fedorovich, E., R. Conzemius, and D. Mironov, 2004b: Convective entrainment into a shear-free, linearly stratified atmosphere: bulk models reevaluated through large eddy simulations. *J. Atmos. Sci.*, **61**, 281–295.

Fedorovich, E., F. T. M. Nieuwstadt, and R. Kaiser, 2001: Numerical and laboratory study of horizontally evolving convective boundary layer. Part I: Transition regimes and development of the mixed layer. *J. Atmos. Sci.*, **58**, 70–86.

Fedorovich, E., R. Rotunno, and B. Stevens, eds., 2004c: *Atmospheric turbulence and mesoscale meteorology*. Cambridge University Press, 280pp pp.

Gaynor, J. E., C. B. Baker, and B. Templeman, 1991: Fine time-scale comparisons between doppler sodar and sonic anemometer-derived winds. *7th Symposium on Meteorological Observations and Instrumentation*, New Orleans, LA, USA. Amer. Meteorol. Soc. Preprints, 401–404.

Gera, N., N. C. Gupta, V. Mohanan, and B. S. Gera, 2013: Sodar studies of air pollution meteorology over Delhi. *Int. J. Sci. Eng. Research*, **4**, 1805–1811.

Gibbs, J. A., E. Fedorovich, and A. M. J. van Eijk, 2011: Evaluating weather research and forecasting (WRF) model predictions of turbulent flow parameters in a dry convective boundary layer. *J. Appl. Meteorol. Climat.*, **50**, 2429–2444.

Giez, A., G. Ehret, R. Schwiesow, K. J. Davis, and D. H. Lenschow, 1999: Water vapor flux measurements from ground-based vertically pointed water vapor differential absorption lidars. *J. Atmos. Oceanic Tech.*, **16**, 237–250.

Greenhut, G. K. and G. Mastrantonio, 1989: Turbulence kinetic energy budget profiles derived from Doppler sodar measurements. *J. Appl. Meteorol.*, **28**, 99–106.

Hadelberg, C. R. and N. K. K. Gamage, 1994: Structure-preserving wavelet decomposition of intermittent turbulence. *Boundary-Layer Meteorol.*, **70**, 216–246.

Hall, F. F., Jr. and E. J. Owens, 1975: Atmospheric acoustic echo sounding investigations at the south pole. *Antarct. J. U.S.*, **10**, 191–192.

Hill, R. J., 1997: Algorithms for obtaining atmospheric surface-layer fluxes from scintillation measurements. *J. Atmos. Oceanic Tech.*, **14**, 456–467.

Holland, G. J., P. J. W. adn J. A. Curry, G. Tyrell, D. Gauntlett, G. Brett, J. Becker, R. Hoag, and W. Vaglienti, 2001: The aerosonde robotic aircraft: A new paradigm for environmental observations. *Bull. Amer. Meteorol. Soc.*, **82**, 889–901.

Holtslag, A. A. M., 2006: GEWEX atmospheric boundary-layer study (GABLS) on stable boundary layer. *Boundary-Layer Meteorol.*, **118**, 243–246.

ISO9613-1, S., Geneva, 1993: Acoustics – attenuation of sound during propagation outdoors – part 1: Calculation of the absorption of sound by the atmosphere. ISO9613-1:1993(E).

Ito, Y., Y. Kobori, M. Horiguchi, M. Takehisa, and Y. Mitsuta, 1989: Development of wind profiling sodar. *J. Atmos. Oceanic Tech.*, **6**, 779–784.

Jorgensen, D. F., 1984: Mesoscale and convective-scale characteristics of mature hurricanes. Part I: General observations by research aircraft. *J. Atmos. Sci.*, **41**, 1268–1286.

Jung, Y., G. Zhang, and M. Xue, 2008: Assimilation of simulated polarimetric radar data for a convective storm using the ensemble Kalman filter. Part I: Observation operators for reflectivity and polarimetric variables. *Mon. Wea. Rev.*, **136**, 2228–2245.

Kaimal, J. C., H. N. Baynton, and J. E. Gaynor, 1980: The boulder low-level inert-comparison experiment. Technical report, BAO Rep. 2, Wave Propulsion Laboratory, NOAA/ERL.

Kallistratova, M. A., 1997: Physical grounds for acoustic remote sensing of the boundary layer. *Acoustic remote sensing application*, S. P. Singal, ed., Springer Verlag, volume 69, 3–35.

—, 2002: Acoustic waves in the turbulent atmpshere: a review. *J. Atmos. Oceanic Tech.*, **19**, 1139–1150.

Kallistratova, M. A. and R. L. Coulter, 2004: Application of sodars in the study and monitoring of the environment. *Meteor. Atmos. Phys.*, **85**, 21–37.

Kallistratova, M. A. and R. D. Kouznetsov, 2012: Low-level jets in the Moscow region in summer and winter obeserved with a sodar network. *Boundary-Layer Meteorol.*, **143**, 159–175.

Kallistratova, M. A., R. D. Kouznetsov, V. F. Kramar, and D. D. Kuznetsov, 2013: Profiles of wind speed variances within nocturnal low-level jets observed with a sodar. *J. Atmos. Oceanic Tech.*, **30**, 1970–1977.

Kallistratova, M. A., R. D. Kouznetsov, D. D. Kuznetsov, I. N. Kuznetsova, M. Nakhaev, and G. Chirokova, 2009: Summertime low-level jet characteristics measured by sodars over rural and urban areas. *Meteorol. Z.*, **18**, 289–295.

Kalogiros, J. A. and C. G. Helmis, 1999: Automatic detection of fixed echoes in sodar facsimile records using the Wavelet transform. *J. Atmos. Oceanic Tech.*, **16**, 734–738.

Kay, S. M. and S. L. J. Marple, 1981: Spectrum analysis - a modern perspective. *Proc. IEEE*, **69**, 5–51.

Keder, J., 1999: Detection of inversions and mixing height by REMTECH PA2 sodar in comparison with collocated radiosonde measurements. *Meteorol. Atmos. Phys.*, **71**, 133–138.

Kolmogorov, A. N., 1941: Dissipation of energy in isotropic turbulence. *Dokl. Akad. Nauk SSSR*, **31**, 538–541.

Konrad, T. G., M. L. Hill, R. R. Rowland, and J. H. Meyer, 1970: A small, radio-controlled aircraft as a platform for meteorological sensors. *Appl. Phys. Lab. Tech. Digest*, **10**, 11–19.

Koscielny, A. J., R. J. Doviak, and D. S. Zrnić, 1984: An evaluation of the accuracy of some radar wind profiling techniques. *J. Atmos. Oceanic Tech.*, **1**, 309–320.

Kouznetsov, R. D., 2009: The multi-frequency sodar with high temporal reoslution. *Meteorol. Z.*, **18**, 169–173.

Kraichnan, R. H., 1953: The scattering of sound in a turbulent medium. *J. Acoust. Soc. Amer.*, **25**, 1096–1104.

Lighthill, M. J., 1953: On the energy scattered from the interaction of turbulence with sound or shock waves. *Proc. Cambridge Phil. Soc.*, **49**, 531–555.

Little, C. G., 1969: Acoustic methods for the remote probing of the lower atmosphere. *Proc. IEEE*, **57**, 571–578.

Marshall, J. M., A. M. Peterson, and A. A. Barnes, 1972: Combined radar-acoustic sounding system. *Appl. Opt.*, **11**, 108–112.

Mason, P. J. and S. H. Derbyshire, 1990: Large-eddy simulation of the stably stratified atmospheric boundary layer. *Boundary-Layer Meteorol.*, **53**, 117 – 162.

Masri, P. and A. Bateman, 1995: Identification of non-stationary audio signals using the FFT, with application to analysis-based synthesis of sound. *IEEE Colloquium on Audio Engineering*, 11.1–6.

May, P. T., R. G. Strauch, K. P. Moran, and W. L. Ecklund, 1990: Temperature sounding by RASS with wind profiler radars: A preliminary study. *IEEE Trans.*, **28**, 19–28.

May, R. M., M. I. Biggerstaff, and M. Xue, 2007: A Doppler radar emulator with an application to the detectability of tornadic signatures. *J. Atmos. Oceanic Tech.*, **24**, 1973–1996.

Mayer, S., A. Sandvik, M. O. Jonassen, and J. Reuder, 2011: Atmospheric profiling with the UAS SUMO: a new perspective for the evaluation of fine-scale atmospheric models. *Meteorol. Atmos. Phys.*, **116**, 15–26.

McPherson, R. A., C. A. Fiebrich, K. C. Crawford, J. R. Kilby, D. L. Grimsley, J. E. Martinez, J. B. Basara, B. G. Ilston, D. A. Morris, K. A. Kloesel, A. D. Melvin, H. Shrivastava, J. M. Wolfinbarger, J. P. Bostic, D. B. Demko, R. L. Elliott, S. J. Stadler, J. D. Carlson, and A. J. Sutherland, 2007: Statewide monitoring of the mesoscale environment: a technical update on the oklahoma mesonet. *J. Atmos. Oceanic Tech.*, **24**, 301–321.

Meijninger, W. M. L., O. K. Hartogenesis, W. Koshiek, J. C. B. Hoedjes, R. M. Zuurbier, and H. A. R. D. Bruin, 2002: Determination of area-averaged sensible heat fluxes with a large aperture scintillometer over a heterogeneous surface - Flevoland field experiment. *Boundary-Layer Meteorol.*, **105**, 37–62.

Mesinger, F., G. DiMego, E. Kalnay, K. Mitchell, P. C. Shafran, W. Ebisuzaki, J. Woollen, D. Jović, E. Rogers, E. H. Berbery, M. B. Ek, Y. Fan, R. Grumbine, W. Higgins, H. Li, Y. Lin, G. Manikin, D. Parrish, and W. Shi, 2006: North American Regional Reanalysis. *Bull. Amer. Met. Soc.*, **87**, 343–360.

Monin, A. S., 1962: Characteristics of the scattering of sound in a turbulent atmosphere. *Sov. Phys. Acoust.*, **7**, 370–373.

Morlet, J., 1983: Sampling theory and wave propagation. *Issues in acoustics signal/image processing and recognition*, C. H. Chen, ed., Springer, 233–261.

Moulsley, T. Y. and R. S. Cole, 1979: High frequency atmospheric acoustic sounders. *Atmos. Environ.*, **13**, 347–350.

Muschinski, A., P. P. Sullivan, D. B. Wuertz, R. J. Hill, S. A. Cohn, D. H. Lenschow, and R. J. Doviak, 1999: First synthesis of wind-profiler signals on the basis of large-eddy simulation data. *Radio Sci.*, **34**.

Neff, W. D., 1978a: Beamwidth effects on acoustic backscatter in the planetary boundary layer. *J. Appl. Meteorol.*, **17**.

—, 1978b: Boundary layer research at South Pole Station using acoustic remote sensing. *Antarct. J. U.S.*, **13**, 179–181.

—, 1988: Observations of complex terrain flows using acoustic sounders: Echo interpretation. *Boundary-Layer Meteorol.*, **42**, 207–228.

Newton, K. E., 1955: Instrumentation of a B-17 airplane for clud physics research. Technical report, Dept. meteorol., University of Chicago.

Obukhov, A. M., 1941: Scattering of sound in turbulent flow (in russian). *Dokl. Akad. Nauk SSSR*, **30**, 611–614.

Organization, W. M., 1996: *Guide to meteorological instruments and methods of observation*. WMO, Geneva., 6th edn. edition.

Owens, E. J., 1974: Development of a portable acoustic echo sounder. NOAA Tecnhical Report ERL298-WPL31, NOAA.

Palmer, R. D., M. F. Larsen, E. L. Sheppard, S. Fukao, M. Yamamoto, T. Tsuda, and S. Kato, 1993: Poststatistic steering wind estimation in the troposphere and lower stratosphere. *Radio Sci.*, **28**, 261–271.

Pan, N., 2003: Excess attenuation of an acoustic beam by turbulence. *J. Acoust. Soc. Amer.*, **114**, 3102–3111.

Parish, T. R. and D. H. Bromwich, 1989: Instrumented aircraft observations of the katabiatic wind regime near terra nova bay. *Mon. Wea. Rev.*, **117**, 1570–1585.

Pekour, M. S. and M. A. Kallistratova, 1993: Sodar study of the boundary layer over Moscow for air-pollution application. *Appl. Phys. B*, **57**, 49–55.

Petenko, I. V., 2001: Advanced combination of spectral and wavelet anlysis ('spavelet' analysis). *Boundary-Layer Meteorol.*, **100**, 287–299.

Petenko, I. V. and V. A. Bezverkhnii, 1999: Temporal scales of convective coherent structures derived from sodar data. *Meteorol. Atmos. Phys.*, **71**, 105–116.

Renfrew, I. A. and P. S. Anderson, 2006: Profiles of katabatic flow in summer and winter over Coats Land, Antarctica. *Quart. J. Roy. Meteorol. Soc.*, **132**, 779–802.

Reynolds, O., 1876: On the refraction of sound by the atmosphere. *Phil. Trans. R. Soc. Lond.*, **166**, 315–324.

Russell, P. B. and E. E. Uthe, 1978: Regional patterns of mixing depth and stability: Sodar network measurements for input to air quality models. *Bull. Amer. Meteor. Soc.*, **54**, 1275–1287.

Scipión, D., 2011: *Characterization of the convective boundary layer through a combination of large-eddy simulations and a radar simulator*. Ph.D. thesis, University of Oklahoma.

Scipión, D. E., P. B. Chilson, E. Fedorovich, and R. D. Palmer, 2008: Evaluation of an LES-based wind profiler simulator for observations of a daytime atmospheric convective boundary layer. *J. Atmos. Ocean. Tech.*, **25**, 1423–1436.

Scipión, D. E., R. D. Palmer, P. B. Chilson, E. Fedorovich, and A. M. Botnick, 2009: Retrieval of convective boundary layer wind fields statistics from radar profiler measurements in conjuction with large eddy simulations. *Meteorol. Z.*, **18**, 175–187.

Seibert, P. and M. Langer, 1996: Deriving characteristic parameters of the convective boundary layer from sodar measurements of the vertical velocity variance. *Boundary-Layer Meteorol.*, **81**, 11–22.

Shuqing, M., C. Hongbin, W. Gai, P. Yi, and L. Qiang, 2004: A miniature robotic plane meteorological sounding system. *Adv. Atmos. Sci.*, **21**, 890–896.

Singal, S. P., B. S. Gera, and D. R. Pahwa, 1994: Application of sodar to air pollution meteorology. *Int. J. Remote Sensing*, **15**, 427–441.

Skamarock, W. C., J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, and J. G. Powers, 2008: A description of the advanced research WRF Version 3. Technical report, NCAR, USA.

Soler, M. R., J. Hinojosa, M. Bravo, D. Pino, and J. V.-G. de Arellano, 2003: Analyzing the basic features of different complex terrain flows by means of a doppler sodar and a numerical model: Some implications for air pollution problems. *Meteor. Atmos. Phys.*, **85**, 141–154.

Song, J., K. Liao, R. L. Coulter, and B. M. Lesht, 2005: Climatology of the low-level jet at the Southern Great Plains atmospheric boundary layer experiments site. *J. Appl. Meteorol.*, **44**, 1593–1606.

Stull, R. B., 1988: *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers, Dordrecht, Holland.

Taconet, O. and A. Weill, 1983: Convective plumes in the atmospheric boundary layer as observed with an acoustic Doppler sodar. *Boundary-Layer Meteorol.*, **25**, 143–158.

Tatarskii, V. I., 1961: *Wave Propagation in a Turbulent Medium*. McGraw-Hill, New York.

—, 1971: *The Effects of the Turbulent Atmosphere on Wave Propagation*. Kefer Press, 472 pp.

Taylor, G. I., 1938: The spectrum of turbulence. *Proc. Roy. Soc., Ser. A.*, **164**, 476–490.

Travouillon, T., M. Schöck, S. Els, R. Riddle, and W. Skidmore, 2011: Using a sodar to measure optical turbulence and wind speed for the Thirty Meter Telescope site testing. Part I: reproducability. *Boundary-Layer Meteorol.*, **141**, 273–288.

Turner, D. D., R. A. Ferrare, L. A. H. Brasseur, W. F. Feltz, and T. P. Tooman, 2002: Automated retrievals of watre vapor and aerosol profiles from an operation Raman lidar. *J. Atmos. Oceanic Tech.*, **19**, 37–50.

van den Kroonenberg, A. C., S. Martin, F. Beyrich, and J. Bange, 2012: Spatially-averaged temperature structure parameter over a heterogeneous surface measured by an unmanned aerial vehicle. *Boundary-Layer Meteorol.*, **142**, 55–77.

Wainwright, C. E., P. B. Chilson, J. A. Gibbs, T. A. Bonin, E. Fedorovich, and R. D. Palmer, 2014a: Methods for evaluating structure function parameters using unmanned aerial systems and large eddy simulation. *Boundary-Layer Meteorol.*, **to be submitted**, .

Wainwright, C. E., P. M. Stepanian, P. B. Chilson, R. D. Palmer, E. Fedorovich, and J. A. Gibbs, 2014b: A time series sodar simulator based on large-eddy simulation. *J. Atmos. Oceanic Tech.*, **31**, 876–889.

Walczewski, J., 1997: Application of sodar in urban air-quality monitoring systems. *Acoustic remote sensing applications*, S. P. Singal, ed., Narosa Publishing House.

Weill, A., F. Baudin, J. P. G. amd P. Van Grunderbeeck, and P. Leberre, 1978: Turubulence structure in temperature inversions and in convection fields as observed by Doppler sodar. *Boundary-Layer Meteorol.*, **15**, 375–390.

Weill, A., C. Klapisz, B. Strauss, F. Baudin, C. Jaupart, P. V. Grunderbeeck, and J. P. Goutorbe, 1980: Measuring heat flux and structure functions of temperature fluctuations with an acoustic Doppler sodar. *J. Appl. Meteorol.*, **19**, 199–205.

Wilson, C. and E. Fedorovich, 2012: Direct evaluation of refractive-index structure functions from large-eddy simulation output for atmospheric convective boundary layers. *Acta Geophys.*, **60**, 1474–1492.

Wulfmeyer, V. and J. Bösenberg, 1990: Ground-based differential absorption lidar for water-vapor profiling: assessment of accuracy, resolution, and meteorological applications. *Appl. Opt.*, **37**, 3825–3844.

Yu, T.-Y. and R. D. Palmer, 2001: Atmospheric radar imaging using spatial and frequency diversity. *Radio Sci.*, **36**, 1493–1504.

Yushkov, V. P., M. A. Kallistratova, R. D. Kouznetsov, G. A. Kurbatov, and V. F. Kramar, 2007: Experience in measuring the wind-velocity profile in an urban environment with a Doppler sodar. *Izv. Acad. Sci. USSR Atmos. Oceanic Phys.*, **43**, 168–180.