Project funded by the European Commission under the 6th (EC) RTD Framework Programme (2002- 2006) within the framework of the specific research and technological development programme “Integrating and strengthening the European Research Area”

Project UpWind

Contract No.:
019945 (SES6)

“Integrated Wind Turbine Design”

Remote Sensing (UpWind WP6)
QinetiQ Lidar Measurement Report

AUTHORS:	Chris Hill
AFFILIATION:	QinetiQ
ADDRESS:	Malvern Technology Centre, St Andrews Road, Malvern, Worcs WR14 3PS U.K.
TEL.:	+44 1684 894161
EMAIL:	chill@QinetiQ.com
FURTHER AUTHORS:	Michael Harris, Natural Power
REVIEWER:	Project members

Document Information

DOCUMENT TYPE	Technical report (deliverable 6.1.1)
DOCUMENT NAME:	Lidar measurement report (QINETIQ/TS/FPPS/TR0900813)
REVISION DATE:	15 September 2010
CLASSIFICATION:	R1: Restricted to project members
STATUS:	Released

Abstract: This report provides a detailed description of the measurement process for a Doppler lidar that measures wind speed and direction. Emphasis is on continuous-wave lidars, with QinetiQ's ZephIR^{TM} lidar as the main example. Pulsed lidar operation is more briefly discussed.

Contents

STATUS, CONFIDENTIALITY AND ACCESSIBILITY									
Status				**Confidentiality**				**Accessibility**	
S0	Approved/Released	**x**		**R0**	General public			Private web site	**x**
S1	Reviewed			**R1**	Restricted to project members	**x**		Public web site	
S2	Pending for review			**R2**	Restricted to European. Commission			Paper copy	
S3	Draft for comments			**R3**	Restricted to WP members + PL				
S4	Under preparation			**R4**	Restricted to Task members +WPL+PL				

PL: *Project leader* **WPL:** *Work package leader* **TL:** *Task leader*

1. Introduction

In order to reduce costs associated with the siting of tall masts, the wind energy industry needs methods such as lidar for remotely obtaining accurate wind profiles. However, widespread acceptance by the industry requires that this technique be extensively validated. A longer-term aspiration must be a certifiable and traceable measurement standard leading to formal accreditation. The inclusion of remote sensing techniques in accredited testing is a challenge, and this report makes steps towards that aim by providing a detailed description of all stages of the measurement process.

Section 2 provides an overview of lidar techniques and technology. Different types of lidar system are surveyed, and the generic physical principles underlying their operation are reviewed. *Wind profiling by a ground-based conically-scanned continuous-wave lidar* is the specific case treated in most detail here: it is rapidly becoming a powerful tool in the wind energy industry, and is exemplified by QinetiQ's ZephIR^{TM} lidar. Other cases can be important – the sensor can be mobile or airborne or fitted on a turbine, the scan pattern can be much more complicated, and the modulation format may be any of a multitude familiar in radar – and we mention these more briefly.

Section 3 reviews some factors and assumptions behind our lidar measurements of wind speeds, before section 4 details the practical steps needed to reach a value of wind speed. It is important to understand any sources of error and uncertainty, and these are reviewed and analysed in section 5. Section 6 gives some conclusions and recommendations from this study.

The Appendix contains worked examples of beam geometry and focus range for ZephIR and WindCube wind profiling lidars.

2. Principles of lidar operation and system description

2.1 Brief survey of lidar types

There are many different types of lidar [1] and these perform diverse tasks (e.g. 3D imaging and rangefinding, gas species detection, remote measurement of vibrations). Here we specifically discuss systems for the measurement of wind speed in the atmosphere [2]. Such systems fall into two broad categories: coherent lidar and direct detection lidar. Coherent lidar measures Doppler shifts by comparing the frequency of backscattered radiation *with that of a reference beam via a light beating process*, whereas direct detection lidar [3] performs its frequency-shift measurements *by passing the light through an optical filter*, such as a Fabry-Perot etalon. Coherent wind lidar systems can be categorised according to their emission waveform (pulsed or continuous), waveband (visible, near-IR, far-IR), and their transmit/receive geometry (monostatic or bistatic). This report concentrates on unmodulated continuous-wave (CW) coherent monostatic lidar systems that operate in the telecommunications near-IR band around 1.5 μm [4]. Pulsed all-fibre lidar has also been developed [5], and several groups (including UpWind colleagues from Leosphere/ONERA and Sgurr/HALO) have recently progressed this approach with promising results.

2.2 Principles underlying anemometry by coherent laser radar (CLR)

The principle by which coherent lidar measures the velocity of a target is simple: a beam of coherent radiation illuminates the target, and a small fraction of the light is backscattered into a receiver. Motion of the target along the beam direction leads to a change $\delta\nu$ in the light's frequency via the Doppler shift. This frequency shift is accurately measured by mixing the return signal with a reference beam (usually a portion of the original beam), and sensing the resulting beats at the difference frequency on a photodetector. Figure 1 illustrates a simplified generic CLR as a bistatic system, in which the transmit and receive optics are separate and distinct. In practice a monostatic geometry, in which the transmit and receive paths share common optics, is more usual.

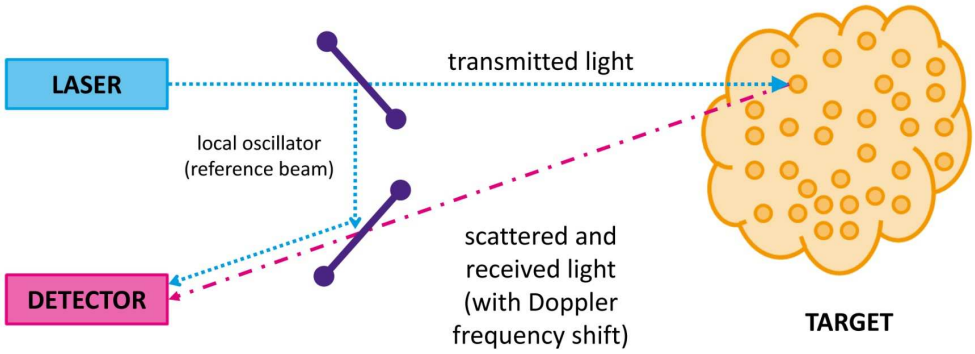

Figure 1: Generic bistatic lidar system. A small fraction of the transmitted light is tapped off by a beamsplitter to form a reference beam. This is superimposed at a second beamsplitter with the weak return scattered from moving particles. The detector's electrical output contains a resulting beat signal; this undergoes spectral analysis to determine particle velocity.

2.3 Role of local oscillator and range selection by focus

The reference beam, or local oscillator (LO), plays a crucial role in the operation of a CLR [6]:

1 It helps to define a three-dimensional region of space ("probe volume") from which light must be scattered for detection of the beat signal.
2 It provides a stable reference frequency to allow very precise velocity determination; as a consequence the Doppler shift measurement by a CLR system is inherently calibrated.
3 It amplifies the signal via the beating process to allow operation at a sensitivity that approaches the shot-noise (or quantum) limit. This very high sensitivity permits the operation of CLR systems in an unseeded atmosphere, relying only on detection of weak backscattering from natural aerosols.

The detector output faithfully reflects the modulations and fluctuations of *one* mode (in space and polarisation) of the scattered light – the mode that matches the LO. With little or no loss of information, these modulations are transferred from the optical to the electrical regime, and can now be conveniently filtered and recorded. Radiation from other sources is rejected spatially and spectrally. Sunlight, for example, comes from a source that is (usually) very far from the probe volume, and very thinly spread across the spectrum; in fact CLR systems are usually completely immune to the effect of background light.

CW systems have the advantage of reduced cost and complexity. A CW lidar achieves operation at a given range by beam focusing. Wind profiling is achieved by focusing at a number of chosen ranges in turn. Focusing of the lidar beam brings about a Lorentzian spatial weighting function (along the beam axis) with its peak located at the beam waist [4, 6]. This function has a half-width given by the Rayleigh range (the distance from the waist at which the beam area has doubled). Roughly, the *beam diameter* at the waist increases linearly with range; the *Rayleigh range* (or range resolution) increases roughly as the square of range; hence the effective probe *volume* varies as the 4^{th} power of the range. This strong dependence on focal range has some implications for the signal statistics at shorter ranges [7].

There are many modulated and non-CW options. In principle, most of the modulation formats and waveform strategies from conventional radar are now available to the lidar designer, because there have been large advances in telecommunications-type hardware (amplifiers, modulators, splitters etc.) and in processor price/performance. Hence the literature is thick with pulsed / chirped / FM / AM formats. In practice, designers for the Doppler wind lidar applications are concentrating on relatively simple repetitive-pulse formats, and seeking the right balance of cost, reliability, output power, sensitivity and range discrimination.

There is no simple overall trade-off "pulsed *versus* CW", but in some common applications the range-resolving considerations are almost decisive by themselves. As said above, the CW lidar relies on beam focusing to define its sensitive volume, and this volume has a length that scales approximately quadratically with focus range at first; for ZephIR this length is about 10 m at 100 m range. At "long" ranges, beyond several hundred metres, diffraction limits mean that the lidar cannot bring its beam to a well-defined focus, so there is little or no inherent range resolution, and some form of pulsed system (where the pulse time-of-flight can be related to the target range) must be preferred.

But at “short” ranges, below a few tens of metres, typical pulsed lidars are ineffective: the pulse length is at least as long as the range, and/or the receiver has a certain “dead time”, during and after pulse transmission, in which (if allowed to receive) it would be dominated by noise from the high-energy transmission. The net result – stated crudely – is that CW lidars are necessary at short ranges; pulsed lidars are necessary at long ranges; there is an intermediate region where they both work fairly well; and the “crossover” from one to the other may be debated but is generally agreed to fall within the range (say 80 m to 150 m) most relevant to modern wind turbines. Section 5 discusses some different meanings of “range resolution”, and the Appendix gives examples of the different focus choices and the associated spatial weightings, for ZephIR and its pulsed colleagues.

2.4 Doppler frequency analysis and signal processing

The stages of signal processing for CLR wind signals are discussed in Section 4.7. The detector output, containing the beat signal information embedded in broadband noise, is typically digitised by an analogue-to-digital converter (ADC). The recorded time series can be analysed as they stand, or transformed (e.g. by fast Fourier transform methods) into “frequency spectra”. When a number of these spectra are averaged in order to reduce fluctuations and improve estimation accuracy – and this may be a large number, hundreds or thousands – we hope a Doppler-shifted peak stands clearly above a flat shot-noise floor. A value for the line-of-sight wind speed can then be computed via a velocity estimation algorithm (finding the peak of the Doppler spectrum, or perhaps the centroid, or some other quantity). There is a large literature on frequency-domain and time-domain lidar Doppler estimators, often drawing on the decades of experience in Doppler radar measurement, and we make only brief remarks in section 4 below.

2.5 Wind profiling in conical scan mode

With a single lidar, since a single measurement provides only the component of wind speed along the beam direction, it is necessary to vary the direction of the beam in order to generate a measurement of the wind speed vector. A conical or VAD (velocity-azimuth-display) scan pattern has been widely used [8] (see Figure 2); as the beam moves, it intercepts the wind at different angles, thereby building up a series of measurements around a disk of air from which the wind speed vector can be derived. In uniform flow, a plot of the measured line-of-sight wind speed (V_{LOS}) *versus* scan azimuth angle (ϕ) takes the form of a cosine wave, with the peak Doppler shifts corresponding with measurements when the azimuth scan angle aligns with the upwind and downwind directions. Doppler shifts close to zero are obtained when the azimuth angle is perpendicular to the flow.

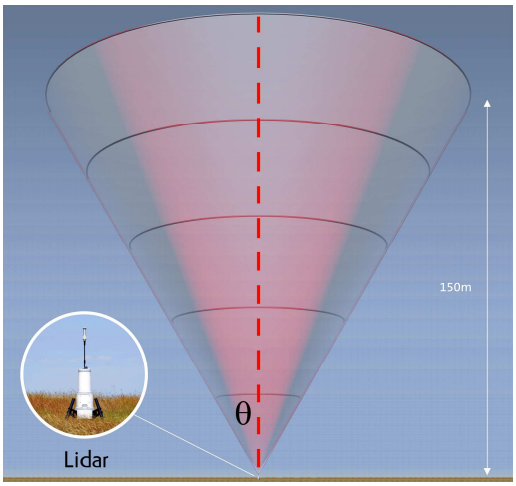

Figure 2: Conical scan pattern as used for lidar wind profiling. The cone half-angle (θ) is typically of order 30°. The lidar can operate successfully even when part of its scan is obscured, e.g. by an adjacent met mast.

For a lidar system that cannot distinguish the sign of the Doppler shift, the plot of V_{LOS} *versus* ϕ follows a rectified cosine wave. There is an ambiguity of 180° in the derived value of wind bearing, but this is easily resolved with reference to a conventional anemometer (cup, sonic etc.) at a height of only a few metres.

2.6 QinetiQ ZephIR lidar

Many different groups have built and successfully deployed wind lidars over the past 30 years. However, commercial lidar products have been available from only a few companies. In 2001 QinetiQ (formerly the government-funded establishment RSRE, later DRA then DERA) began a programme to develop a commercial fibre-based lidar, exploiting decades of research in the coherent lidar area. It launched the ZephIR product in 2003, and over 60 systems have now been deployed successfully around the world, in several demanding applications that illustrate the flexibility and robustness of this wind profiling solution. Initial deployment (March 2003) was on the nacelle of a 2.3 MW wind turbine (Figure 3a), measuring the wind speed up to 200 m in front of the blades. The lidar consisted of a rack unit containing laser source, detector and signal processing computer, situated in the base of the tower, connected via over 100 m of electrical and optical fibre cable to the transceiver head mounted on the top of the nacelle. The lidar system was installed and was fully operational within a few hours, thus allowing a demonstration of advance warning of oncoming gusts and providing valuable experience in practical deployment issues.

Figure 3: Evolution of the ZephIR lidar. The left-hand picture 3a shows the lidar head mounted on the nacelle of a Nordex N-90 wind turbine. The central picture 3b shows a prototype ground-based wind profiler at Risø wind energy test site, Høvsøre, Denmark. The right-hand picture 3c shows the ZephIR production model deployed in the field.

The system returned to QinetiQ Malvern having achieved several weeks of successful operation. It was then converted into a ground-based scanning unit for wind profiling (Figure 3b). This was first trialled in December 2003, and has since been used in numerous campaigns in the UK, Europe, and other parts of the world. The experience gained through these trials has built confidence in the robustness and reliability of the core ZephIR design. In late 2004, work started on a production instrument (Figure 3c)**,** designed to perform autonomous wind profiling measurements at heights up to 150 m [9], primarily for site surveys at proposed wind farm sites. The specifications of ZephIR's monostatic, continuous-wave design, and details of several client installations worldwide, are available at:

http://www.naturalpower.com/products-and-services/zephir/zephir-product.html

QinetiQ have granted an exclusive licence to Natural Power (a subsidiary of Fred Olsen) for marketing ZephIR in the renewable-energy industry.

3. Lidar measurement process: Assumptions

3.1 Behaviour of scattering particles

The lidar signals from which wind speeds are derived originate via backscattering of the beam by particles in the atmosphere. The detailed constitution of these particles is generally unknown, and unnecessary for lidar wind speed measurement; they are normally assumed to consist of dust, organic matter (e.g. pollen), soot, or water droplets. They must provide sufficient signal for Doppler analysis and their motion must faithfully follow that of the atmosphere (i.e. they must be "wind tracers"). This latter assumption is very good for small particles where viscous forces are dominant. Larger particles for which this does not apply will rapidly fall to ground.

Raindrops, snowflakes and hail provide a strong contribution to the lidar signal. The identification of such anomalous "non-tracer" scatterers, so that their effects can be removed or further investigated, may be important. Their downward motion could lead to an error in the vertical component of wind speed (usually a parameter of lesser interest), but the important horizontal component will be correct.

A further excellent assumption is that the return signal is dominated by light generated by single-scattering events. While it is possible for light to suffer multiple scattering in dense cloud, it is assumed that any effect on the Doppler spectrum is negligible.

3.2 Uniformity of flow and backscatter

A least-squares fitting to the azimuthal variation of line-of-sight wind speed allows the derivation of wind parameters from conical scan data, and the derivation is straightforward if we assume a uniform flow. This assumption is usually reasonable, but sometimes misleading (see section 5.6). These parameters pertain to a significant volume of atmosphere – the signal originates from a disk whose diameter commonly exceeds 100 m, and whose depth along the beam direction can be over 10 m.

The contribution to the lidar signal from different regions of the lidar probe volume is weighted by the value of the atmospheric backscatter coefficient $\beta(\pi)$ at each point. The value of $\beta(\pi)$ is typically constant to ~10 % throughout the probe volume [8] except in conditions that lead to stable mist layers, or when the lidar beam intersects a low cloud base.

3.3 Beam positional accuracy

Lidar focus calibration is performed in the laboratory, and must be correctly maintained throughout a period of deployment in the field. Obviously errors in the focus setting would result in wind speed measurement at the wrong height. The real-life behaviour of optical beams projected through finite apertures is slightly different from the behaviour of ideal light rays; as discussed in the Appendix, if we specify the "focus range" (meaning the projection distance at which the beam width is a minimum, not the focal point for an ideal ray tracing), we may also need to choose explicitly between two

solutions of the beam width equations. Careful design will eliminate any serious uncertainty about the beam focus: thermal expansion (which could change the length of the transceiver telescope) is compensated, and the position of the focus mechanism can be automatically checked to provide information on any malfunction.

The lidar must be correctly set up, with the vertical and azimuthal orientation adjusted appropriately during installation.

The atmosphere has small-scale refractive-index atmospheric fluctuations, but for typical conditions and short/medium ranges they will have negligible effect on the propagation of the lidar beam [10].

3.4 Optical and electrical interference sources

ZephIR identifies the presence of a wind signal when the power density in the Doppler spectrum exceeds a threshold level. In the absence of any significant source of spurious signal, the only mechanism that can lead to such detection events is the backscatter of Doppler-shifted light from the probe volume into the lidar receiver. Optical interference is highly unlikely – even when the lidar points directly at the sun the spectral power density is insufficient to cause a problem, and interaction between two lidars placed side-by-side (including the possibility of interference from the beam emitted by an adjacent lidar) can be neglected. Careful screening eliminates the risk, for any normal deployment, of spurious spectral features caused by electrical interference.

3.5 Time-of-flight considerations

The round-trip time for light interrogating the atmosphere at a height of 100 m is 0.83 μs. On this timescale the ZephIR scanner moves the focused beam a distance of only 300 μm, and the laser phase drifts by an insignificant amount. The polarisation state of the lidar output is similarly frozen on this timescale. For faster scanning or longer ranges (or, worse, both), the effects of nonzero time-of-flight can be important, so that special optics for “lag angle” correction are needed in pulsed lidars.

4. End-to-end measurement process

4.1 Introduction

The measurement process can be split into a number of steps. This section describes these steps in turn, arriving at an overall end-to-end description of the wind speed measurement process for a CW coherent Doppler lidar wind profiler.

4.2 Transmitter optics

The role of the transmitter is to provide a focused beam at a desired location. This location can be moved around in space with a combination of (i) changing the focus range and (ii) passing the beam through a scanning element such as a rotating prism (wedge). Wind profiling lidars conveniently employ a conical scan with its axis aligned vertically; the cone half-angle θ is commonly of order 30° (i.e. the beam elevation angle is ~60°).

In a monostatic CW system, with the LO and transmit beams properly matched, a Doppler-shifted contribution to the signal is generated via light scattering from any moving part of the atmosphere that the beam illuminates. The contribution from any point is weighted by the square of the beam's intensity at that point [4,11]. Hence it can be shown that focusing of the beam gives rise to a spatial sensitivity along the beam direction that depends on the inverse of beam area; it follows that the sensitivity rises to a peak at the beam waist, and falls symmetrically on either side. There is also a spatial dependence of sensitivity transverse to the beam, but because the beam is very narrow this is of little interest and can be ignored. To a good approximation the axial weighting function for a continuous-wave (CW) monostatic coherent lidar is given by a Lorentzian function [4,6]:

$$F = \frac{\Gamma / \pi}{\Delta^2 + \Gamma^2} \tag{4.1}$$

where Δ is the distance from the focus position along the beam direction, and Γ is the half-width of the weighting function (to -3 dB point, i.e. 50 % of peak sensitivity). Note that F has been normalised such that its integral from $-\infty$ to ∞ gives unity. To another good approximation, Γ is given by:

$$\Gamma = \frac{\lambda R^2}{\pi A^2} \tag{4.2}$$

where λ is the laser wavelength (e.g. the telecommunications wavelength ~1.55×10^{-6} m), R is the distance of the beam focus from the lidar output lens, and A is the beam radius at the output lens (the beam intensity profile is assumed to be an axially-symmetric 2D Gaussian; A is calculated for the point at which the intensity has dropped to $1/e^2$ of its value at the beam centre). For example, if A takes the value 24 mm, then at a focus range R of 100 m Γ has a value of ~8.5 m, or a probe length (to -3 dB points) of ~17 m. Further examples are in Section 5.

4.3 Light scattering by aerosols

Coherent lidar measures the Doppler shift resulting from the component of target velocity along the beam (or line-of-sight) direction. Motion of the target transverse to the beam direction produces no net Doppler shift. Hence, consider a lidar located at the origin (0,0,0) and aimed at a specific location (x,y,z) so that the line-of-sight vector may be notated (x**i** + y**j** + z**k)** or (**x** + **y** + **z**). If the local wind vector at (x,y,z) is similarly notated with components u, v and w, the detected line-of-sight velocity component is:

$$V_{LOS} = \left| (\quad + \quad + \quad) . \left(\frac{\quad + \quad}{\sqrt{x^2 + y^2 + z^2}} \right) \right| \tag{4.3}$$

This is the dot product of the wind vector and the unit vector along the beam direction, and V_{LOS} is a modulus because the standard version of the ZephIR lidar is incapable of distinguishing motion towards/away for any isolated Doppler measurement.

In the backscattering geometry considered here, the scattered light experiences a Doppler shift in frequency given by:

$$\delta\nu = \frac{2V_{LOS}}{c}\nu = \frac{2V_{LOS}}{\lambda} \tag{4.4}$$

where c is the speed of light (2.998 x 10^8 m s^{-1}), and ν and λ are respectively the laser frequency and wavelength.

Since the signal originates from a nonzero probe length containing different scatterers, the overall return exhibits a variety of frequencies. This “spectrum” results from the contributions from different velocities (with relative strengths determined by the weighting function, Eqn 4.1) over all the space occupied by the lidar beam. The contributions add “incoherently”: their phases are randomised, and over the timescale of our spectral estimation the total power in each frequency bin is well approximated by the sum of powers from the individual contributions at the corresponding velocities. This is a key simplification, by no means true of all lidars (pulsed or CW) or all targets. Note that from a single such CW lidar measurement, without further information, it is not possible to identify from what range each component of the spectrum has originated.

For a CW coherent system, the time-averaged optical power P_S backscattered by the aerosols into the receiver is well approximated by:

$$P_S = \pi P_T \beta(\pi)\lambda \,, \tag{4.5}$$

where P_T is the transmitted laser power and $\beta(\pi)$ is the atmospheric backscatter coefficient in (m sr)$^{-1}$. It is notable that Eqn 4.5 contains no dependence on either the focus range or the system aperture size. With a value of 10^{-8} (m sr)$^{-1}$ for $\beta(\pi)$ in clear boundary-layer air, a transmitted power P_T ~ 1 W and λ ~ 1.5 μm, the received power P_S is only ~5 x 10^{-14} W so the need for high sensitivity is plain.

4.4 Receiver optics

In a monostatic system, the backscattered light returns through the transmission optics (the word *transceiver* is commonly used to denote this dual role). Any scanning motion of the beam during the timescale for the radiation to travel to the aerosols and back will result in misalignment (lag angle) of the receiver, but this is insignificant for the range of parameters considered here.

Within the transceiver, optical means are used to isolate the return light, which is passed to the next stages of the detection process.

4.5 Light beating

In a coherent laser radar, the incoming Doppler-shifted radiation is optically mixed with a reference or local oscillator (LO) beam. The mixing of two waves in this manner leads to the well-known “beat” phenomenon in which the resulting amplitude oscillates at the difference frequency. The efficiency of the beating process is optimised when the signal and LO beams overlap perfectly in space (i.e. they occupy identical “modes”). This condition is ensured when both beams propagate in the same single-mode optical fibre and share the same polarisation state.

It is instructive to consider a simple classical description of the light beating process. Superposition of a LO field $E_{LO} \cos \omega_{LO} t$ and a stable signal field $E_S \cos \omega_S t$ results in a fluctuating detector output:

$$i(t) \propto [E_{LO} \cos \omega_{LO} t + E_S \cos \omega_S t]^2 \quad (4.6)$$

Typical detectors cannot follow the very fast frequencies (~10^{14} Hz) of the optical field amplitudes, but they can follow the “slow” (MHz or GHz) fluctuations of the net field intensity – hence the “square-law” form of the current in Eqn 4.6. This is conveniently separated into a “constant” term and a cross term oscillating at the difference frequency:

$$i(t) \propto [{E_{LO}}^2 + {E_S}^2] + 2E_{LO}E_S \cos|(\omega_S - \omega_{LO})|t \quad . \quad (4.7)$$

Since the optical power of the local oscillator beam typically exceeds that of the signal beam by many orders of magnitude, the first term is given by E_{LO}^2 to a very good approximation; quantum fluctuations on this give rise to the shot noise floor of the instrument (section 4.6). For a system without direction sensing (for which there is no frequency shift between the LO and transmitted beams) the measured Doppler shift is given simply by:

$$\delta\nu = 2\pi|(\omega_S - \omega_{LO})| \quad (4.8)$$

from which the value of V_{LOS} is derived via Eqn 4.4.

The coherent detection process involving light beating is also commonly referred to as offset-homodyne or self-heterodyne detection. A rigorous quantum-mechanical theoretical treatment of the detection process is given in [13]. Note that although the detection process is described as coherent, the backscattered radiation itself is incoherent in nature, meaning that its phase is uncorrelated with that of either the transmitted beam or the local oscillator. In practice a signal field originating from atmospheric scattering is the sum of many small contributions from different scatterers, and this sum shows random fluctuations in amplitude and phase (or frequency). The coherent detection process ensures that these fluctuations are reproduced in the

detector output i(t); in the limit of high SNR, analysis of i(t) will give a correct representation of the scattered light's spectrum [12]. The timescale of the random fluctuations is roughly the inverse of the signal bandwidth (see section 4.7).

4.6 Photodetection

The beat signal is detected by directing the optically-mixed beams onto a photodiode, which produces photoelectrons in response to the incident photons (in the square-law sense of Eqn 4.6). The resulting photocurrent is normally passed through further stages of amplification before digitisation. There are generally four contributions to the output of the photodetector module:

1 Dark noise – this is the intrinsic wideband noise floor generated by the detector and amplifier combination in the absence of any incident light.
2 Photon shot noise [14] (sometimes called quantum noise) – the random generation of photoelectrons by the incident LO beam leads to a wideband, spectrally flat (white) Gaussian noise source. The shot noise power spectral density increases in proportion to the optical power of the LO beam.
3 Laser relative intensity noise (RIN) – intensity fluctuations that are in excess of shot noise, caused (e.g.) by relaxation oscillation [15] of the laser output. Such oscillation is typically at relatively low frequency, peaking below 1 MHz, and hence only affects the sensitivity of the lidar at low wind speed. In some systems it is possible to cancel the RIN by use of a dual-channel balanced detector.
4 Beat term resulting from the wind signal – this is the contribution that contains the information on Doppler shifts from which the wind speed is derived. Its power spectral density increases in proportion **both** to the LO power **and** to the signal power (Eqn 4.7).

The requirements for the detector are: high quantum efficiency; sufficient bandwidth to cope with the maximum Doppler frequencies of interest; and a shot noise contribution that significantly exceeds that of dark noise (this third requirement depends on a combination of the detector's intrinsic noise floor and the optical saturation threshold). There are reliable, well-suited, off-the-shelf photodiodes in the telecommunications wavelength band around 1.5 μm.

4.7 Fourier analysis

To extract the Doppler frequency information we perform a spectral analysis of the detector output. This is conveniently done digitally; an example of a typical signal processing procedure is described below. An ADC with a sampling rate of 100 MHz permits spectral analysis up to a maximum frequency of 50 MHz, corresponding with a wind speed V_{LOS} of ~38.8 m s^{-1} (Eqn 4.5, with λ = 1.55 μm). A hardware low-pass filter with a cut-off frequency of 50 MHz, inserted between the detector and ADC, eliminates aliasing problems. Spectra are calculated by digital Fourier transform (DFT) methods; a 512 point DFT gives rise to 256 points in the output spectrum with a bin width of ~200 kHz (corresponding with a line-of-sight velocity range of ~0.15 m s^{-1}). Each DFT represents ~5 μs of data; successive DFTs are then calculated, and the resulting "voltage" spectra are modulus-squared in order to generate a power spectrum. These power spectra are then averaged to find a mean spectrum for the averaging period.

The random fluctuation in the shot noise floor of the spectrum reduces as the square root of the number of averages: the sensitivity increases by this same factor. For 4,000 averages the measurement time amounts to ~20 ms (or a data rate of ~50 Hz). This requires that the processing is capable of 100 % duty cycle, which is achieved in ZephIR with a fast Fourier transform (FFT) block within a field-programmable gate array (FPGA). It has been shown that a standard fast PC with no additional duties to perform can achieve a similar performance. It is possible to accommodate reasonable variations in any of the above parameters (sample rate, DFT size, number of averages) and maintain the 100 % duty cycle.

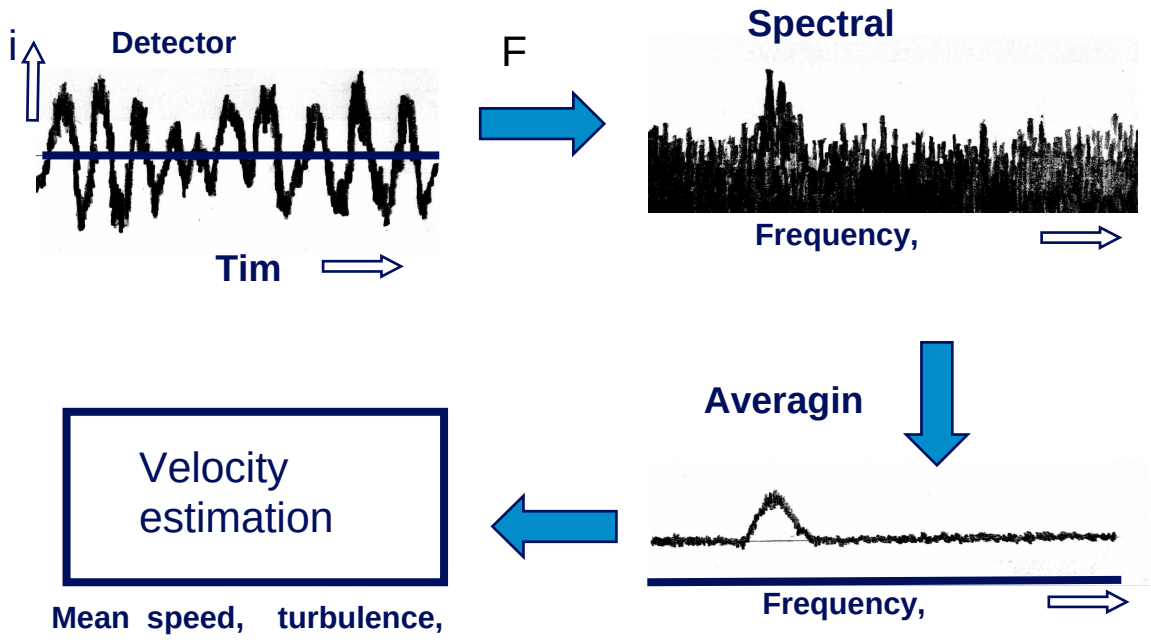

Figure 4: Stages in lidar signal processing: the Fourier transform analysis is carried out by a computer integrated into the lidar system. As an example, 4000 individual spectra might be averaged to achieve high sensitivity and measurable returns even in very clear air. This entire process takes only 20 milliseconds, giving ~50 measurements per second of line-of-sight wind velocity.

The width of the Doppler spectrum is determined by three elements:

1 Instrumental width: this corresponds closely with the 200 kHz bin width mentioned above.
2 Transit-time broadening: during the conical scan, the beam passes through the aerosol particles in a timescale of ~10-15 μs, independent of the lidar focus setting. The corresponding broadening is again of order 200 kHz.
3 Turbulence broadening: the probing of a significant volume results in a range of Doppler shifts from parts of the atmosphere that are moving at different speeds (see section 4.3). This contribution usually dominates the spectral width, except under conditions of very uniform airflow. As turbulence and shear increase, the width increases, and occasionally multiple Doppler peaks may appear.

It is useful at this stage to consider system sensitivity requirements. The signal-to-noise ratio (SNR) for a wind speed measurement by a CW CLR is:

$$SNR = \frac{\eta P_S}{(hc/\lambda)\Delta\nu[1 + D(\nu) + R(\nu)]} \quad (4.9)$$

Here η is an efficiency term incorporating optical losses and photodetector sensitivity, P_S is the input signal power as defined in Eqn 4.5, and hc/λ is the light quantum energy, of order 1.3×10^{-19} J. The signal bandwidth $\Delta\nu$ is determined by the three contributions listed above, and the term inside the square brackets denotes the various noise sources listed in section 4.6. $D(\nu)$ and $R(\nu)$ represent the power spectral density (at frequency ν) from dark noise and RIN respectively in units of the power spectral density of the local oscillator shot noise. Ideally $D(\nu)$ and $R(\nu)$ should both be << 1 over the range of Doppler frequencies of principal interest, so that the shot noise is the dominant noise source.

The SNR as defined here is the power spectral density at the Doppler peak divided by that in the surrounding noise floor. The averaging of many spectra (described in the following sections) does not increase SNR according to this definition, but it ensures that good performance can be obtained even when the SNR is well below unity. For example, an SNR of 0.1 will just exceed a 6σ threshold level for an average of 4000 spectra. From the above it is possible to derive an approximate value of $\beta(\pi)_{min} \sim 10^{-9}$ -10^{-8} (m sr)$^{-1}$ for the minimum detectable backscatter (assuming a transmitted intensity 1 W and a 20 ms measurement time).

4.8 Velocity estimation

So in this example each measurement of line-of-sight wind speed, obtained over a timescale of ~20 ms, generates a Doppler spectrum that shows one or more peaks (of variable width), superimposed on a noise floor that is predominantly white but may have spectral features originating from RIN and dark noise sources. This section outlines the steps necessary to derive an appropriate estimate of the wind speed.

First, the noise floor is “whitened” so that each spectral bin contains the same mean noise level; this is achieved by dividing the power value in each bin of the spectrum by a previously-measured value for the same bin obtained with the shutter closed. A flat threshold is then applied at a pre-determined level above the mean noise. A suitable and conservative choice for the threshold is 6 standard deviations (6σ) above the mean noise level. In the absence of any wind signal (e.g. with the output of the lidar blocked) such a setting will give rise to negligible occurrences in which the noise alone exceeds threshold. It follows that any bin whose level exceeds the threshold is deemed to contain a valid contribution to the wind spectrum. For each 20 ms measurement, the wind spectrum is reconstructed by subtracting the mean noise contribution from the contents of each bin that exceeds threshold, and applying a small recorrection for any distortion resulting from the noise whitening. In order to proceed to the next stage, a single speed value must be derived from the resulting spectrum. A number of options are available, including peak and median values; a common solution is to calculate the mean (or centroid) value $<V_{LOS}>$.

A series of these values of mean line-of-sight wind speed is generated during a conical scan. ZephIR generates 50 values per one-second rotation, and calculates wind parameters from up to 150 (three rotations).

4.9 Least-squares fitting routine

The fitting routine accepts up to 150 pairs of values of $\langle V_{LOS} \rangle$ and azimuth angle ϕ. In conditions of uniform wind flow, these give a rectified cosine wave (Figure 5a):

$$\langle V_{LOS} \rangle = |a \cos(\phi - b) + c| \quad (4.10)$$

The derivation of this function is straightforward and can be found in a number of publications, e.g. [8]. The peaks of the function correspond with the azimuth angle aligned parallel or anti-parallel to the wind direction. The function passes through zero when the azimuth angle is perpendicular to wind bearing since there is no component of velocity along the line of sight. The data are also conveniently displayed on a polar plot (Figure 5), which provides information at a glance on the speed, direction and vertical wind component. A standard least-squares fitting routine provides the estimates of the values of the three floating parameters (a, b and c). The high level of redundancy in the fitting process is advantageous and can be used to identify non-uniform flow. The root mean square deviation of the points from the optimum solution gives an indication of the quality of fit, and this can be related to the value of turbulence intensity.

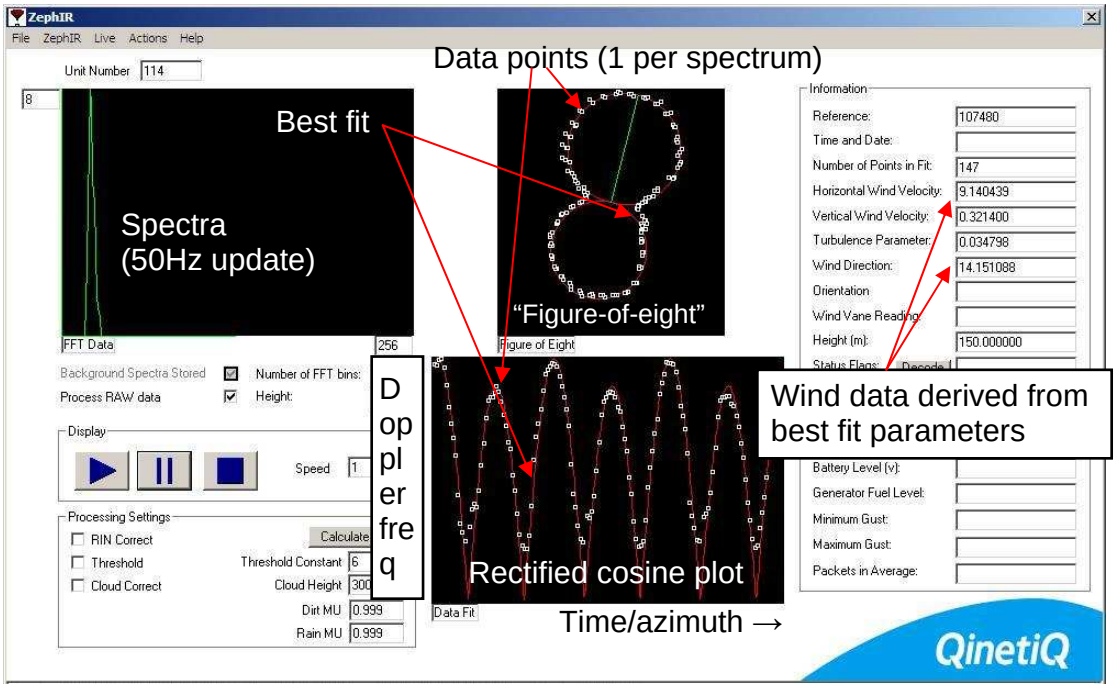

Figure 5: Wind lidar output screen, illustrating many of the features of a wind profile measurement. This example has been obtained at a height 150 m above ground level, one of several heights being probed in sequence. The lower trace shows 147 individual line-of-sight wind speed values (obtained over a total period of 3 seconds) plotted as white squares against azimuth scan angle. The same data, along with the least-squares fit in red, are displayed above this in polar coordinates on the figure-of-eight plot. The wind bearing lies slightly to the E of N. The wind parameters, derived from the fit, appear in the table on the right; the horizontal wind speed at this height is determined to be 9.1 metres per second, or roughly 18 knots. Each of the 147 points in these two traces is derived from an averaged power spectral estimate; one of these 147 PSEs is plotted at upper left.

More work is needed to establish a full understanding of the turbulence information available from lidar signals [16]. Note that much additional information on turbulence is available from the spectral widths and shapes of the individual line-of-sight wind speed measurements, but this information is not currently used to evaluate turbulence parameters. A single number (V_{LOS}) summarises the information from one measurement - this is a serious but convenient simplification; because of data storage constraints, the other information is commonly discarded after the centroiding process.

4.10 Parameter extraction

The wind parameters for each 3-second measurement period are extracted from the best fit as follows (θ is the cone half angle of order 30°):

$$\begin{aligned} &\textit{Horizontal speed (u) } V_H = a/\sin\theta \\ &\textit{Vertical speed (w) } V_V = -c/\cos\theta \\ &\textit{Bearing } B = b, \textit{ or } b \pm 180^\circ \end{aligned} \qquad (4.11)$$

Because of the ambiguity in the sign of the Doppler shift, there are two equally valid best-fit solutions corresponding with values of b separated by 180°. The correct choice is usually easy: we choose the solution that lies closer to a conventional measurement from a met station situated close to ground. Conventionally, a wind profiling lidar incorporates such a station that performs these (and other) measurements and feeds the results to the analysis software.

The 3-second wind parameter values are stored internally for subsequent analysis; they can also undergo further processing for extraction of average values.

4.11 Data averaging

It is a common requirement to calculate 10-minute averaged wind data for compatibility with industry standards. This is most easily achieved by calculation of the arithmetic mean (“scalar average”) of the 3-second values of V_H, V_V and B that have been obtained during the required period. Care is needed to avoid double counting of measurements taken during 3-second periods that straddle the boundary between successive 10-minute periods. A vector average is also possible in which the resultant of the individual measurements is calculated over each 10-minute period. In practice the results from the two methods differ negligibly in reasonably stable conditions.

When the lidar is operating as a wind profiler it is necessary to measure each height in series. Hence, at any given height the wind is not monitored continuously. Instead, a 3-second measurement is followed by a period of order 20 seconds during which the lidar is focused at other heights. Since this sampling is carried out randomly with respect to any behaviour of the wind, this duty cycle of order 15 % has negligible impact on the validity of the resulting 10-minute averaged values.

5. Uncertainty analysis

5.1 Rain / snow / cloud, solid objects

In general the Doppler shift measured by coherent laser radar is very accurate. This is apparent from Eqn 4.5 as long as the laser wavelength remains stable and the signal processing has been correctly performed – both good assumptions in practice. The values of $<V_{LOS}>$ that are derived from the centroids of the spectra can be measured to considerably better than a bin width. A greater source of error arises from uncertainty about what provides the scattering from which the Doppler shift is derived. The scattering is assumed to originate from atmospheric particles moving at the *same* speed as the wind and positioned close to the focus of the lidar beam (section 3.1). An obvious example where this breaks down is when the beam intersects a solid object (e.g. a bird) that is moving at a different speed. However, in such a case the derived value of $<V_{LOS}>$ will stand out as clearly anomalous on the polar plot (Figure 5). The presence of such points will be diluted by the correct values of $<V_{LOS}>$ obtained from uncontaminated parts of the atmosphere, so that only a small bias should result. A further safeguard against these erroneous points is provided by a simple "outlier removal" algorithm. This identifies and removes any points that lie anomalously far from the best-fit solution to Eqn 4.10. The least-squares routine is then rerun on this slightly reduced set of $<V_{LOS}>$, ϕ data pairs (up to 150-N pairs if there are N outliers).

The presence of precipitation within the probe volume leads to a different source of uncertainty. The downward motion of rain and snow inevitably leads to some error in the vertical component of wind speed. However, the presence of rain and snow is normally easily identified from the measurement, so the resulting values of vertical wind can be eliminated from the data. This still leaves a possibility of error in the horizontal component if the drops retain some memory of higher wind speeds as they fall through a strong shear layer. In practice, mast/lidar comparisons performed to date have shown that any such bias is of little significance.

5.2 Cloud effects

Continuous-wave (CW) laser wind profilers do not use time gating (unlike pulsed systems) and rely on focusing the beam in order to measure wind speed at a given height. This technique can lead to problems when the beam impacts a cloud base at higher altitude: under severe conditions the contribution to the Doppler signal from cloud can contaminate that from the aerosols at the desired height. Under normal wind profile conditions (wind increases with height), this will lead to an overestimate of wind speed. The severity depends on a number of factors; the risk increases for:

1 low cloud height
2 high lidar altitude setting
3 low aerosol scattering at desired height
4 high wind shear

A general approach to mitigating this problem needs first to identify the presence of a cloud return and then remove its contribution from the Doppler spectra. Cloud returns have a number of characteristics that potentially allow them to be distinguished from aerosol returns:

1 Velocity usually higher (but there is no guarantee of this)

2 Spectral width usually narrower (again, with no guarantee)
3 Power in Doppler peak has clear dependence on lidar focus. The power is maximised when the lidar beam is focused close to the height of the cloud base.
4 Doppler spectrum is independent of focus range

The latter two characteristics are dependable and can be used as a basis for identification and elimination of spurious cloud returns.

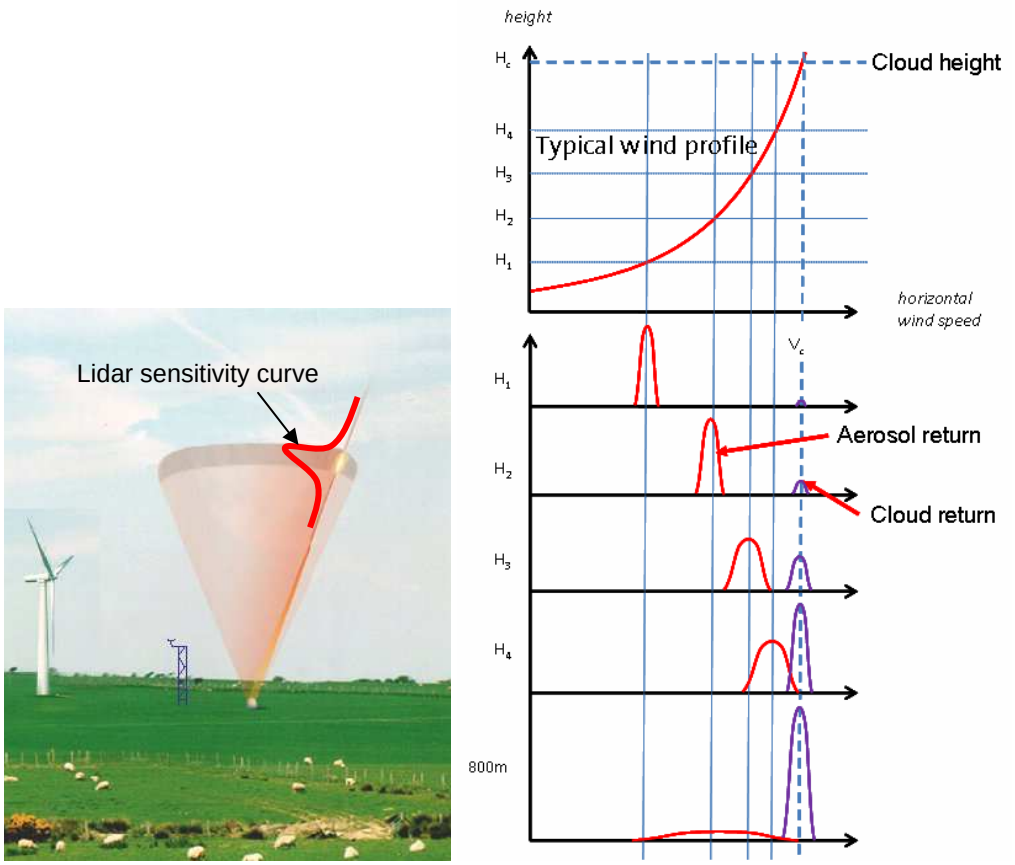

Figure 6: Cloud removal. The left plot shows the lidar conical scan focused at a typical height above ground level. The Lorentzian sensitivity curve is also shown; a spurious return is generated when the far wing of this curve intersects a strongly scattering low cloud layer. The right plot shows the aerosol (red) and cloud (purple) returns as the lidar is focused at various heights – the strength of cloud contamination increases with focus height. The cloud signal is easily identified from the 800 m focus, and these data are then used to eliminate the cloud return at the measurement heights.

The general strategy is outlined in the following steps (and shown in Figure 6):

1. Routinely run the lidar at an additional greater height (typically 800 m – essentially a collimated beam output) and an additional very low height (typically 38 m), immediately before or after the maximum height of interest (say 150 m). Use these extra measurements to decide whether significant contamination from cloud returns is present.
2. For each azimuth angle within the scan at 150 m, identify the cloud spectrum (from the 800 m measurements) obtained at the closest value of azimuth angle.

3. Apply suitable thresholding on the cloud spectrum to identify bins occupied by cloud signal.
4. Subtract a suitably scaled version of the cloud spectrum from the 150 m spectrum. The resulting spectrum will have its cloud contributions reduced below threshold.
5. Run standard thresholding and centroiding routines on corrected spectra and fit to the rectified cosine wave (Eqn 4.10) as usual to obtain wind parameters.

Algorithms based on this approach have been implemented and used to analyse lidar data from systems beside instrumented met masts. Initial results indicated the effectiveness of the cloud removal approach [17], but analysis by Risø/DTU indicated the possibility of up to 5 % bias under certain conditions. A revised version of the cloud correction algorithm (described in a forthcoming QinetiQ report for UpWind) has mitigated the problems, and further testing has been performed at Høvsøre.

5.3 System positioning accuracy

Correct alignment ensures the risks are low, but errors in aligning the lidar during setting up will have an impact on the measurement of wind bearing (if the lidar is rotated from its correct orientation) and vertical wind speed (if the lidar is tilted, so that the axis of its conical scan is not precisely vertical). For a small tilt angle δ, the error in vertical wind speed V_V will vary from $\pm V_H \sin\delta$ (if the tilt is towards or away from the direction of the wind) to zero (if the tilt is perpendicular to the wind). Any bias on V_H is negligible to first order.

5.4 Probe volume effects

As discussed in section 4.2, the lidar samples the motion of air from a nonzero volume of air, centred around the beam waist at the focus. Clearly there is no bias while all the air within the probe volume moves at the same speed. However, there is usually some degree of shear across the sample region. For a linear shear this leads to spectral broadening of the returns, but no overall bias. A strong non-linear shear profile across the probe volume is required to induce any bias of significance; in practice such conditions will be rare (and very rare for measurement heights below ~100 m, where the probe length is small).

5.5 Range resolution

Having aimed our laser in a particular direction, we may wish to estimate how the average backscatter coefficient β, and the Doppler shift, vary as functions of range along the beam. Then extension to 2D and 3D profiles, by varying the aim direction, can be important. There is a large literature on the estimation of atmospheric backscatter profiles and wind speed profiles, mostly dealing with pulsed lidars, and often drawing on results or mathematical approaches known from conventional radar. But there is currently no firm agreement within UpWind about a practical definition of “range resolution”. Some different standard meanings are:

1 Length of “probe volume”. For CW lidar we may choose e.g. the half-height points of the Lorentzian weighting function.
2 “Instantaneous sample depth”. For pulsed lidar this means the thickness of the region from which light is received at any instant. We may choose e.g. the half-

height points of the pulse shape edges.

3 Frehlich et al. in [18] and many other papers often use an “effective range resolution” $\Delta R = \Delta p + \Delta r$ where Δp is the distance that the pulse moves per estimate, and Δr is the spatial extent of the pulse. But this is a definition for notational purposes, not yet a statement of physically achieved “resolution”. Other authors use slightly different combinations of Δp and Δr. These formal treatments typically begin with expressions for the complex time series of the detector output (a convolution of optical pulse and atmospheric backscatter/Doppler response), with more or less detail of the physics of detection / aperturing / turbulence / filtering / sampling / windowing. Then a suitable average gives the output's autocorrelation coefficient, which is sometimes easier to manipulate than the time series or frequency spectrum. Banakh and Smalikho [19] noted that, for given assumptions, an integral showing weighted contributions from each range to the velocity estimate could be expressed in closed form using error functions.

4 We apply some “reasonable” test along the lines “if there are two individual contributions to the detector output, from two individual features or scatterers at ranges L_1 and L_2, and we gradually increase their separation L_2-L_1, then we will call the range resolution the first value of L_2-L_1 at which we can just confidently say that there are two separate contributions”. Confidence is assumed when the crossing point of the two response functions occurs at half (or some other arbitrary fraction) of their peak height. This is like the common approach taken to define optical fringe visibility (the “Rayleigh criterion” is $\theta = 1.22\ \lambda / D$ for the angular separation of two Airy discs, when the maximum of one coincides with the minimum of its neighbour).

5 It is also like the approach defined in the corresponding UpWind sodar document [20], where “spatial resolution” refers to an idealised case where two thin atmospheric layers are to be detected separately: “*Spatial resolution* is the vertical separation Δz_σ of two infinitely thin layers which are resolved in the returned signal. Two peaks in the time series are resolved if the signal drops to at least half power between them. It can be shown that spatial scales smaller than $\Delta z_\sigma = c\tau/2$ are not resolved. This is the spatial resolution due to the pulse length” [20].

6 Because the pulsed-lidar physical detector output is produced by a convolution of the pulse shape and the atmospheric scattering, we can try to invert the process and estimate the β and Doppler profiles by numerical *de*convolution. Deconvolution can be computationally demanding, and badly and nonlinearly affected by noise. It takes advantage of our prior knowledge of some of the functions involved: we may know the transmitted pulse behaviour very accurately, and may also have some knowledge of the likely atmospheric layering. The resulting “resolution” – if that means our ability to distinguish slightly separated layers, or slightly different speeds, or both – is not necessarily limited by pulse durations or Fourier analysis bin widths. The achievable spatial resolution may be much better than $c\tau/2$ at high CNR, but degrade sharply at low CNR as the deconvolution becomes unstable and noise-dominated.

7 Much work at Risø and elsewhere has considered the turbulent fluctuations of air, and modelled their scale sizes and (under the assumption of frozen flow) the resulting spectrum of Doppler lidar output fluctuations. Normally we wish to use the lidar (with its known sensor response) to determine the incompletely known atmospheric properties: the mean “β and Doppler profiles” just mentioned, but also the spectral and statistical parameters of time-varying turbulent flow. But this can be reversed: given that the lidar responds to a probe volume of a fixed but unknown shape (and averages or filters over the

fluctuating scattering within it), and given knowledge or estimates of all *other* parameters, we can use the post-filtering observations, such as spectral behaviour, to find the best-fit shape. Mikkelsen [21] has recently reviewed the probe volume definitions and literature conventions, taking ZephIR and WindCube as examples.

In a practical test, when we base our decision on the behaviour of the processed detector output, the “range resolution” – in any sense where we want to detect returns from different thin slices, and correctly estimate and assign their Doppler shifts – becomes dependent on CNR. Many authors quantify the “confidence” by borrowing results from statistical decision theory (some of them originally developed for radar applications), so that the effects of Gaussian noise are included.

It may be desirable to choose a practical definition that is (a) realisable in a field test or calibration (b) relevant to UpWind interests (c) applicable to all the different systems that we wish to compare (d) confined to high-CNR cases. It may not be easy to arrange a clean, controllable, experimental test of an atmosphere with (a) several layers of varying β but uniform speed (b) several layers of uniform β but varying speed (c) more general cases. Also there are many intermittent and hard-to-control interfering factors in real atmospheres (birds, rain, snow etc.) that will distort the time series and statistics unless they are recognised and removed. But first we recommend discussion on a meaning for “resolution”.

5.6 Flow uniformity

The assumption of uniform flow is inherent in most conical-scan algorithms. Except in situations of strong shear or turbulence it is a reasonable assumption, and the best-fit wind parameters may be used to indicate the average values over the (large) sampled volume. But it is not a necessary assumption, and its implications are important. For instance, the usual signal processing (see sections 4.9 and 4.10), which couples horizontal and vertical components, will *invent* nonexistent winds (a net vertical component, and a horizontal component bias) if the Doppler estimates at different scan angles are affected by local turbulence or the inevitable noises that affect real measurements. And the wind will be misestimated in “complex” terrain, e.g. where the wind remains locally parallel with the ground but the ground is not flat on a larger scale. This is true of ZephIR, WindCube, Galion and other devices that extract wind estimates from angle-diverse data under the assumption of flow uniformity.

Most simply, in the presence of non-uniform flow (section 3.2), a volume-averaged lidar measurement can indicate a wind speed different from that of a point measurement by a mast-mounted cup anemometer. However, under such conditions it can be argued that the lidar will provide a more reliable measure of the mean wind speed, averaged over a volume of air relevant to wind energy applications.

Also, to an important extent the assumption of wind field uniformity can be straightforwardly *checked*, if the conical scan provides measurements at many different azimuth angles. This is a definite advantage over systems that provide only a few angles: three or four suffice to measure the wind, but only if the flow is indeed uniform.

In principle the well-sampled angular data might permit identification and study of data contaminated by the influence of e.g. turbine wakes or topography. In complex

terrain, the flow undergoes stable and unstable non-uniformities, and the figure-of-eight plot (Figure 5) can distort systematically for a given wind direction, reflecting the speeding up and slowing down in certain regions of the scan. The ZephIR lidar provides a lot of information about the flow non-uniformity, with up to 50 points being interrogated around the scan disk.

So work continues on combining ZephIR data with flow-modelling software in order to improve resource assessment in complex terrain. At the price of increased complexity in hardware and processing, we can relax the uniform-flow assumption, and use other scan patterns and/or other algorithms. These may be better suited to estimating 3D wind flows with fewer constraints (although it is unlikely that we would want to estimate real wind fields subject to *no* constraints of continuity or velocity limits).

5.7 Dependence on backscatter level

Under conditions of high backscatter, the spectrum provides an accurate measure of the distribution of line-of-sight velocities within the probe volume (weighted according to Eqn 4.1). A fall in backscattering strength (usually associated with an increase in air clarity) has the same effect as raising the detection threshold (section 4.8), and leads to elimination from the spectrum of weaker components of velocity. The impact of the system noise floor on the detailed spectral shape will also be increased. The centroid values $<V_{LOS}>$ will be unbiased and independent of threshold level when the spectrum is symmetrical. However, for a skewed (asymmetric) spectrum the precise value of $<V_{LOS}>$ will be sensitive to the threshold level. Hence a small difference in measured wind speed is possible between two measurements under conditions that are identical in every way apart from the level of backscatter. However, there is no evidence from the many comparisons so far to suggest that this leads in practice to a measurable discrepancy.

A further possibility to be considered is the effect of saturation (by very strong scattering returns from thick cloud) of the lidar detector, electronics or signal processing. The last of these has the lowest threshold, caused by the finite limits of input voltage to the ADC. In the event that the input signal exceeds these limits, the spectrum will become distorted, featuring higher harmonic components of the true Doppler frequencies. In practice, the range of inputs to the ADC can be tailored to accommodate the highest levels of backscatter that will reasonably be encountered, eliminating the risk of bias.

5.8 Beam obscuration

Lidar can operate successfully even when part of its scan is obscured. This confers great flexibility so that the system can easily be located adjacent to masts, buildings or in forests. Stationary objects pose no major problem other than the loss of wind measurements from the relevant obscured sector of the scan; the fit to Eqn 4.10 will no longer contain data over the full 360 degree range of ϕ. Laboratory experiments on moving belt targets have indicated that accurate ZephIR measurements are obtained even when over half of the scan is obscured. Early studies with slow-scanning 10 µm wind lidars successfully extracted winds from 1/8 or even 1/16 of a rotation [22].

Lidars such as WindCube that use only three or four view angles cannot afford to have even one of them obscured. This is not usually a serious problem in a reasonably fixed environment. If the lidar can be initially aligned so that none of the view angles coincides with an obstruction (mast, tree, tower etc.), then its operation is unaffected.

5.9 Wind direction

ZephIR's two best-fit solutions of Eqn 4.10 give values of wind direction that are 180° apart. Selection between the two options is made with reference to the measurement of wind direction from a ground-based anemometer. This needs to be in disagreement by over 90° with the direction at the chosen height for the incorrect choice to be made. Such a directional shear is conceivable in highly complex terrain, but very unlikely in the reasonably uniform conditions of interest for wind energy applications.

If we make the wrong choice (a wind bearing error of 180°), the value of vertical component of the wind V_V will have the wrong sign. In other words, an updraught will be wrongly identified as a downdraught (of the same absolute speed) and vice versa.

6. Conclusions and recommendations

A detailed examination of the measurement process for coherent monostatic CW lidar reveals that the basic acquisition of line-of-sight Doppler spectra is a well-established method with little scope for gross errors and miscalibration. The subsequent steps required to convert these spectra into a profile of wind speed are more complex, however, and their validity relies on a number of well-established assumptions. Much work has already been performed to test the validity of assumptions outlined in section 3, and to understand the issues discussed in section 5.

Currently, the clearest demonstration of validity must be provided by direct side-by-side comparisons between the lidar system and a fully instrumented met mast of height 100 m or more. Rigorous careful comparisons must be carried out to avoid a number of known problems associated with cup anemometers [23], for example:

1 Shadowing of the cup anemometer by the mast from certain directions.
2 Cup sensitivity to any vertical wind component.
3 Topographic effects leading to non-uniform flow across the area occupied by mast and lidar scan (including turbine wakes).
4 Valid cup anemometer calibration.

The comparisons performed so far are providing mounting confidence in the validity of the lidar technique. Notably, this has been achieved over a rapid timescale, with the first such comparison during the summer of 2004 [9]. There is currently a need for agreement on a unified method to allow meaningful comparison between the performance of different remote sensing systems; the issue of “range resolution” and its meaning as a performance measure for different groups has been noted.

The extraction of turbulence data relevant to the wind industry from lidar signals is an area that will benefit from further research and verification through field comparisons.

A further opportunity exists for computer modelling as a way to check the validity of assumptions. Such an exercise would start with a realistic 3D description of the atmosphere (characterised by wind speed, shear, turbulence and backscatter parameters), and proceed with a mathematical transformation to “lidarise” the data (i.e. calculate what the lidar will see in each given set of simulated conditions). The study could hence be used as a means to investigate any bias under each specific set of conditions.

The first version of this report was written for UpWind in 2006 (see “Remote sensing (UPWIND WP6) six-month progress report”, Risø-I-2527(EN), November 2006). This revision benefits from discussions with Risø scientists and other colleagues at UpWind meetings, the 2008 Risø Ph.D. Summer School, and the 2009 Coherent Laser Radar Conference. The lidar scientists at each company made essential contributions: Mike Harris at Natural Power, Jean-Pierre Cariou and Matthieu Boquet at Léosphère, and Guy Pearson at HALO Photonics. See also M Harris, “Introduction to continuous-wave Doppler lidar” (2010 Risø Ph.D. Summer School), for recent ZephIR developments.

7. References

1 A V Jelalian, “Laser radar systems”, Boston: Artech (1992).
2 J A Zak “Atmospheric boundary layer sensors for application in a wake vortex advisory system”, NASA/CR-2003-212175.
3 M L Chanin, A Gariner, A Hauchecorne and J Portneuve, “A Doppler lidar for measuring winds in the middle atmosphere”, Geophys. Res. Lett. **16** 1273-1276 (1989).
4 C Karlsson, F Olsson, D Letalick and M Harris, “All-fiber multifunction CW 1.55 micron coherent laser radar for range, speed, vibration and wind measurements”, Applied Optics **39** 3716-3726 (2000).
5 G N Pearson, P J Roberts, J R Eacock and M Harris, “Analysis of the performance of a coherent pulsed fiber lidar for aerosol backscatter applications”, Applied Optics **41** 6442-6450 (2002).
6 C M Sonnenschein and F A Horrigan, “Signal-to-noise relationships for coaxial systems that heterodyne backscatter from the atmosphere”, Applied Optics **10** 1600-1604 (1971).
7 M Harris et al, “Single-particle laser Doppler anemometry at 1.55 μm”, Applied Optics **40** 969-973 (2001).
8 V A Banakh, I N Smalikho, F Köpp and C Werner, “Representativeness of wind measurements with a CW Doppler lidar in the atmospheric boundary layer”, Applied Optics **34** 2055-2067 (1993).
9 D A Smith et al, “Wind lidar evaluation at the Danish wind test site in Høvsøre”, Wind Energy **9** 87-93 (2006).
10 S F Clifford and S Wandzura, “ Monostatic heterodyne lidar performance: the effect of the turbulent atmosphere”, Applied Optics **20** 514-516 (1981); L Lading, S Hanson and A Skov Jensen, Applied Optics **23** 2492-2497 (1984).
11 M Harris, G Constant and C Ward, “Continuous-wave bistatic laser Doppler wind sensor”, Applied Optics **40** 1501-1506 (2001).
12 M Harris, G N Pearson, C A Hill and J M Vaughan, “Higher moments of scattered light fields by heterodyne analysis”, Applied Optics **33** 7226-7230 (1994).
13 R Loudon, “The quantum theory of light”, 3rd edition, Oxford University Press (2000).
14 B I Bleaney and B Bleaney, “Electricity and Magnetism”, Oxford University Press, Section 23.4 (1976).
15 A E Siegman, “Lasers”, Mill Valley, California: University Science Books, Section 25.1, (1986).
16 V A Banakh, I N Smalikho, F Köpp and C Werner, “Measurements of turbulent energy dissipation rate with a CW Doppler lidar in the atmospheric boundary layer”, J. Atm. Oceanic Tech. **16** 1044-1061 (1999).
17 D Kindler, A Oldroyd, A MacAskill, D Finch, “An 8 month test campaign of the QinetiQ ZephIR system: preliminary results”, Extended Abstracts of the Intern. Symp. for the Advancement of Boundary Layer Remote Sensing (ISARS 13), 18-20 July 2006, Garmisch-Partenkirchen, Germany. Wiss. Ber. FZKA 7222, 165-167.
18 R Frehlich, “Effects of wind turbulence on coherent Doppler lidar performance”, J. Atm. Oceanic Tech. **14** 54-75 (1997).

19 V A Banakh and I N Smalikho, “Estimation of the turbulent energy dissipation rate from the pulsed Doppler lidar data”, Atmos. Oceanic Opt. **10** 957-965 (1997). They derive a "longitudinal size of the sounding volume" (ct/2)/erf(t/2σ), for a Gaussian pulse shape exp(-t^2/σ^2), with several assumptions about the slowly varying velocity, the autocorrelation velocity estimator, and additive noises. They define spatial resolution as $1/Q_s(R)$ where Q_s is the normalised spatial weighting.
20 S Bradley and S von Hünerbein, “Remote sensing of flow – a description of the measurement at all stages: SODAR”, in “Remote sensing (UPWIND WP6) six-month progress report”, Risø-I-2527(EN), November 2006.
21 T Mikkelsen, “On mean wind and turbulence profile measurements from ground-based wind lidars : limitations in time and space resolution with continuous wave and pulsed lidar systems”, EWEC 2009 (Marseille, France, 16-19 March 2009), online Proceedings paper PO 171.
22 R L Schwiesow, F Köpp and Ch.Werner, "Comparison of CW-Lidar measured wind values obtained by full conical scan, conical sector scan and two-point techniques", J. Atm. Oceanic Tech. **2** 3-14 (1985).
23 L Kristensen, “The perennial cup anemometer”, Wind Energy **2** 59-75 (1999).

8. Appendix: Lidar focus and beam geometry

This appendix describes the influence of beam focus on the performance of lidar wind sensors.

Since the autumn 2009 UpWind meeting we have surveyed the three main coherent wind lidars from European companies (Léosphère / ONERA, Natural Power / QinetiQ, and Sgurr / HALO Photonics) and prepared this summary of their chosen optical beam geometries. Setting and confirming the distance of the focus from the sensor, and the width of the beam at the focus, are important steps in sensor calibration. This distance and this width affect the measurement accuracy and resolution, and can be chosen within certain limits.

When the lidar receives scattered light from many individual aerosols or particulates at a certain range, and the phases of individual contributions are unrelated (so that the net resulting detector output can be represented as a random walk with Gaussian statistics), then the contribution from that range to the mean carrier strength C becomes inversely proportional to the beam area (Section 4.2 above). This “1/area” rule is useful for understanding lidar behaviour, although further examination is needed to define the lidar performance (particularly if we wish to compare continuous and pulsed operating modes). For example, the carrier-to-noise ratio (CNR) involves the bandwidth(s) of the physical processes, not just their statistics.

Below we consider the ZephIR, WindCube and HALO lidars.

Examples of focus calculations

The optical arrangement in **ZephIR** is straightforward: light at approximately 1.57 microns is launched from a single-mode fibre, expands through free-space propagation until it reaches a single focusing lens, then passes through the rotating wedge and continues to propagate outwards through the atmosphere.

We make some simplifying approximations. We assume that the light leaves the fibre as a pure TEM_{00} (fundamental Gaussian) beam; note there are slight differences between ideal ray optics and real beam optics, and here we will use the familiar first-order Gaussian beam propagation equations. The fibre is much narrower than the lens, and the outgoing light reaches the lens with essentially a spherical wavefront whose centre of curvature is on the fibre axis and very near the fibre end. Some of the TEM_{00} is intercepted by the lens edge or other circular aperture, but we assume that this aperturing is negligible: we ignore edge effects, scattering into non-TEM_{00} modes, etc. Attenuation is neglected – this is usually a good approximation for typical UpWind ranges, but it becomes worse at long range and particularly if the beam propagates long distances through low-level atmospheres.

At very large distances, regardless of the fibre/lens distance, the outgoing beam diverges. But for some fibre/lens distances the beam width at first reduces, reaches a local minimum, and then increases again, and the range where the width reaches a minimum is called the “focus range” in this Appendix. In ZephIR this range can be set conveniently anywhere from ~ 10 m to ~ 200 m by translating the fibre. We find that the total range of movement of the fibre (over all useful focus ranges) is less than 1 % of

the fibre-to-lens spacing. Note that the optical strength or “focal length” of the *lens* is fixed, and generally different from this “focus range”.

Figure A1 shows the variations of the focus range (distance of waist from lens), and of the beam waist size at that focus, as the fibre is moved. Figure A2 shows the same information with a logarithmic scale for the fibre movement.

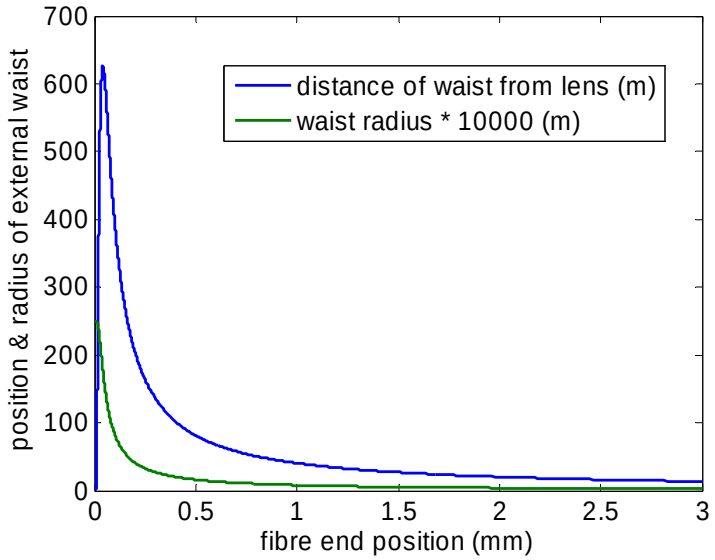

Figure A1: ZephIR nominal design parameters: Variation of focus beam width and focus range.

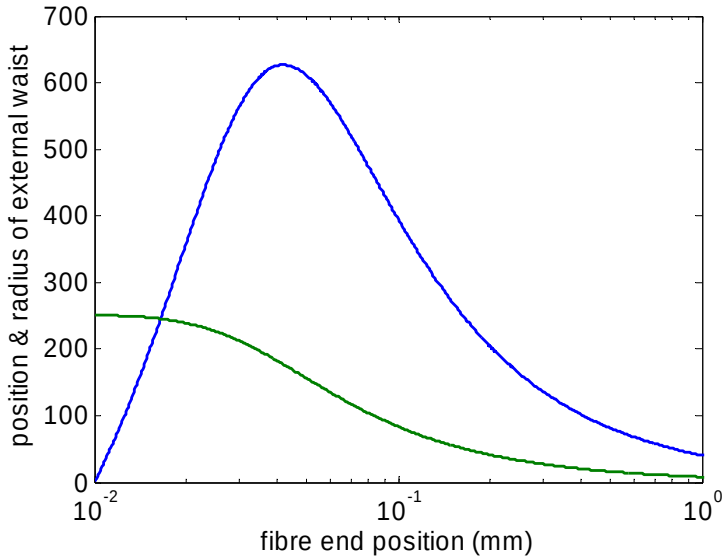

Figure A2: As Figure A1, with logarithmic scaling of fibre end position.

First, we see that a statement of “focus range”, with no further information, may be ambiguous. A quadratic equation relates the beam waist radius w_0 to the beam radius w(z) at distance z:

$$w^2(z) = w_0^2 [1 + (z/b)^2]$$

where $b = \pi w_0^2/\lambda$, and z = distance from the beam waist plane. This equation is symmetric on either side of the waist (although we remember that z is constrained on one side, since the lens intervenes).

To obtain the beam parameters at a certain range r, we substitute $z = r_{waist} - r$ in the beam equations. Note again that w_0 is a local minimum value of the beam’s radius, defined by the e^{-2} intensity criterion, and existing at the waist plane, a distance “focus range” or r_{waist} from the lens; and r_{waist} for an initially collimated beam is not quite equal to the lens focal length f (because of the “slight differences” above).

As we vary the beam width at the lens, we can obtain no solution, one solution, or two solutions for a physically valid beam width at a given focus range r_{waist}. This is illustrated in Figure A3 for a focus range of 100 m (typical for UpWind applications), and wavelength 1.575 μm.

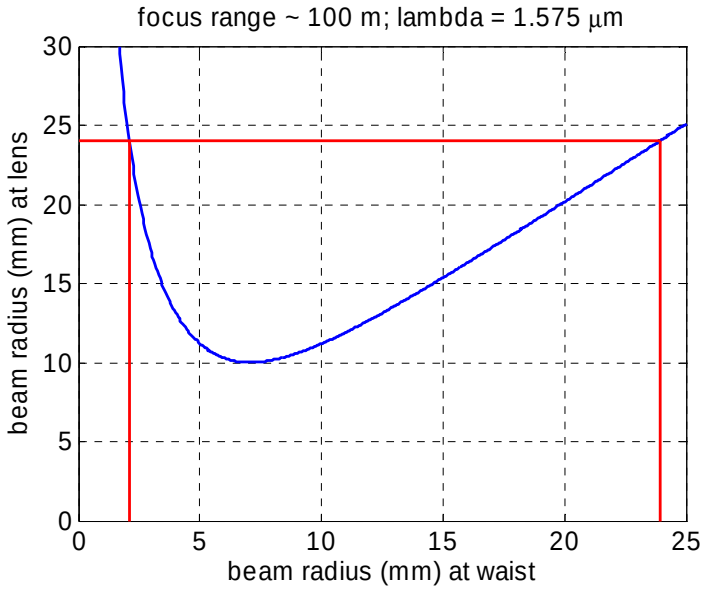

Figure A3: ZephIR nominal design parameters: Variation of beam radius w (at the lens; z = 0) as a function of the beam waist radius at the external focus.

For the nominal ZephIR beam radius at the lens of w(r = 0) = 24 mm, we have drawn a red horizontal line which intersects the blue curve twice, because there are two solutions of the quadratic: a strong focus (w_0 ~ 2.1 mm) to the left of the minimum, and a weak focus (w_0 ~ 23.9 mm, barely smaller than the initial size) to its right.

As w(r = 0) is reduced, the two solutions converge, until we reach a “degenerate” solution with w_0 ~ 7 mm when w(r = 0) is about 10 mm; for w(r = 0) smaller than this there is no solution.

The ZephIR weak focus above – a smaller beam area at long range, but a larger area at short range – is not generally used, because ZephIR is not intended for long ranges.

Secondly, we note that this weaker of two distinct focus geometries can in theory exist for ZephIR at a range beyond 600 m. But it would be very weak, with poor range resolution. As explained in Section 4 above, the lidar sensitivity has a near-Lorentzian profile whose "probe length" is usually taken as its FWHM, equal to the "Rayleigh range" $2b = 2\pi w_0^2/\lambda$. This length becomes reasonably short only when the focus range is less than about 300 m, and in practice ZephIR is specified for ranges up to about 200 m. Only fibre displacements greater than about 100 μm (that is, the right-hand half of Figure A2) are relevant to ordinary ZephIR operation. "Reasonably short" is a loose term, meaning roughly a few tens of metres (comparable with the range resolution of pulsed lidars), or significantly shorter than the focus range itself.

The focusing performance of several ZephIR systems has been checked after extensive deployment by customers in the field, and found to be unchanged from its original calibration: this is a consequence of the design in which the focus position is set via a closed loop system, providing automatic checking for positional discrepancy. Easily conducted tests also confirm that the design goals are met, and that the focusing has no measurable dependence on temperature.

The WLS-7 and WLS-70 **WindCubes** are pulsed lidars, not continuous-wave, but if there are many independent contributing scatterers we have the same Gaussian-statistics result: "1/area" is proportional to the detector carrier power C (per unit range) due to the scatterers at a given range.

We consider here only the effects of variations in this area. The overall CNR calculation for a pulsed lidar has extra complications not examined here in detail. First, the bandwidths necessary for detection can be influenced by lidar pulse shape and duration as well as by scatterer movements. Second, the range resolution is usually obtained by *time*-gating the detector output. This gating can lead to further bandwidth and resolution effects in the signal processing, as well as to the obvious limits in spatial location of the contributing scatterers. So the focus geometry is no longer the sole means of obtaining resolution, but it still influences the CNR, and needs appropriate and explicit choices if CNR is to be preserved at ranges of several hundred or many hundred metres. We might choose, for example, to maximise CNR at a certain range (say 100 m, or the maximum envisaged range), or to maximise the average CNR across all measurement ranges.

If a strong focus at ~ 100 m is kept, but measurements are made at longer range, their CNR falls rapidly as the beam expands away from its focus. The WLS-7 Léosphère model is intended for ranges from 40 m up to 200 m, whereas the WLS-70 is intended for ranges up to at least 2000 m. And indeed, according to the manufacturers [1], the decisions are slightly different for the different models, and different from the ZephIR settings.

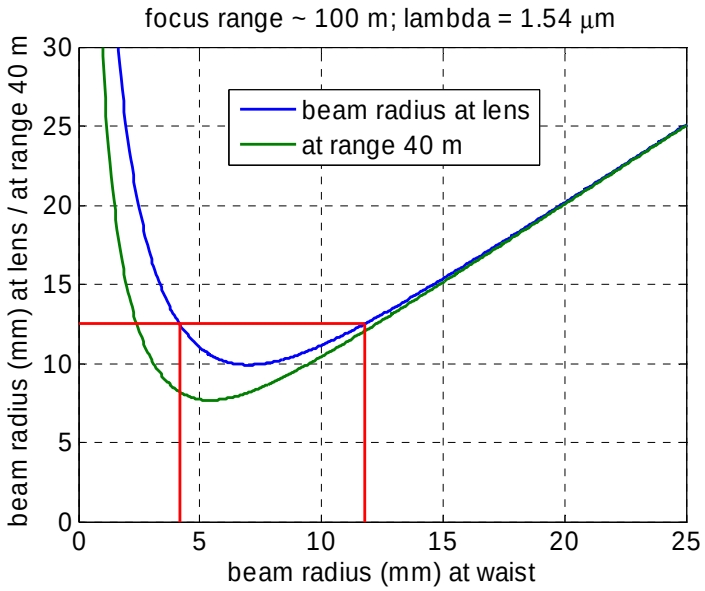

Figure A4: WindCube WLS-7 nominal design parameters: Variation of beam radius w at the lens (r = 0), and at minimum range (r = 40 m), as functions of the beam waist radius at the external focus (r = 100 m from the lens).

In Figure A4 we take nominal values for the WLS-7 of 25 mm beam diameter at the 50 mm diameter lens, with a fixed focus range of ~ 100 m [1]. Again the horizontal red line (at beam radius w(r = 0) = 12.5 mm) intersects each curve twice. The strongly focused beam has a waist radius of 4.2 mm, and a radius at range 40 m (the nominal minimum range) of 8.2 mm. So at 100 m the CNR (per unit range) is about 6 dB above the CNR for this minimum range, because the beam radius is smaller by a factor ~ 2 and the area is smaller by a factor ~ 4.

At 200 m (the nominal maximum range of WLS-7, and by coincidence twice the focus range in this example), the beam radius is again 12.5 mm, and the CNR is about 10 dB lower than at 100 m. This is a deliberate and reasonable compromise, achieving an acceptable CNR throughout the ~ 10 range bins.

The more powerful WLS-70 has nominally 50 mm beam diameter at the 100 mm diameter lens, with a fixed focus range of ~ 400 m [1]. Because of its higher power and longer pulse, the minimum range is ~ 100 m (Figure A5):

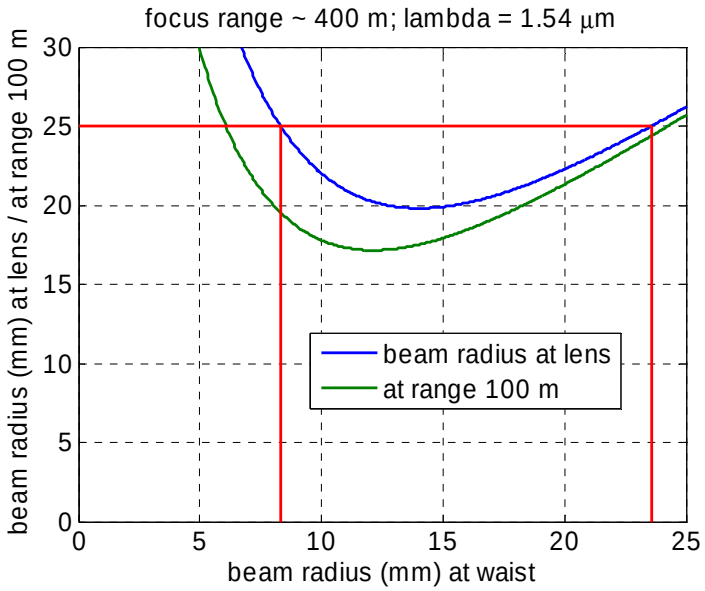

Figure A5: WindCube WLS-70 nominal design parameters: Variation of beam radius w at the lens (r = 0), and at minimum range (r = 100 m), as functions of the beam waist radius at the external focus (r = 400 m from the lens).

Now the red line drawn horizontally (at beam radius w(r = 0) = 25 mm) intersects the blue curve at w_0 = 8.3 mm and 23.6 mm. The strong-focus choice has a radius at range 100 m of 19.5 mm. So at 400 m the CNR (per unit range) is about $(19.5 / 8.3)^2$ ~ 5.5 times, or ~ 7.4 dB, higher than the CNR for the nominal minimum range of ~ 100 m.

For the weak-focus choice, the beam is almost collimated and has a radius of 24.4 mm at 100 m range. So at 400 m the CNR is about $(24.4 / 23.6)^2$ or ~ 0.3 dB higher than the CNR for 100 m.

At a nominal maximum range of 2000 m, the strong and weak focus geometries have CNR penalties of 21.1 dB and 4.8 dB respectively (compared with their CNRs at 400 m).

The relative CNRs for WLS-70 as functions of range, for the two focus choices, are shown in Figure A6:

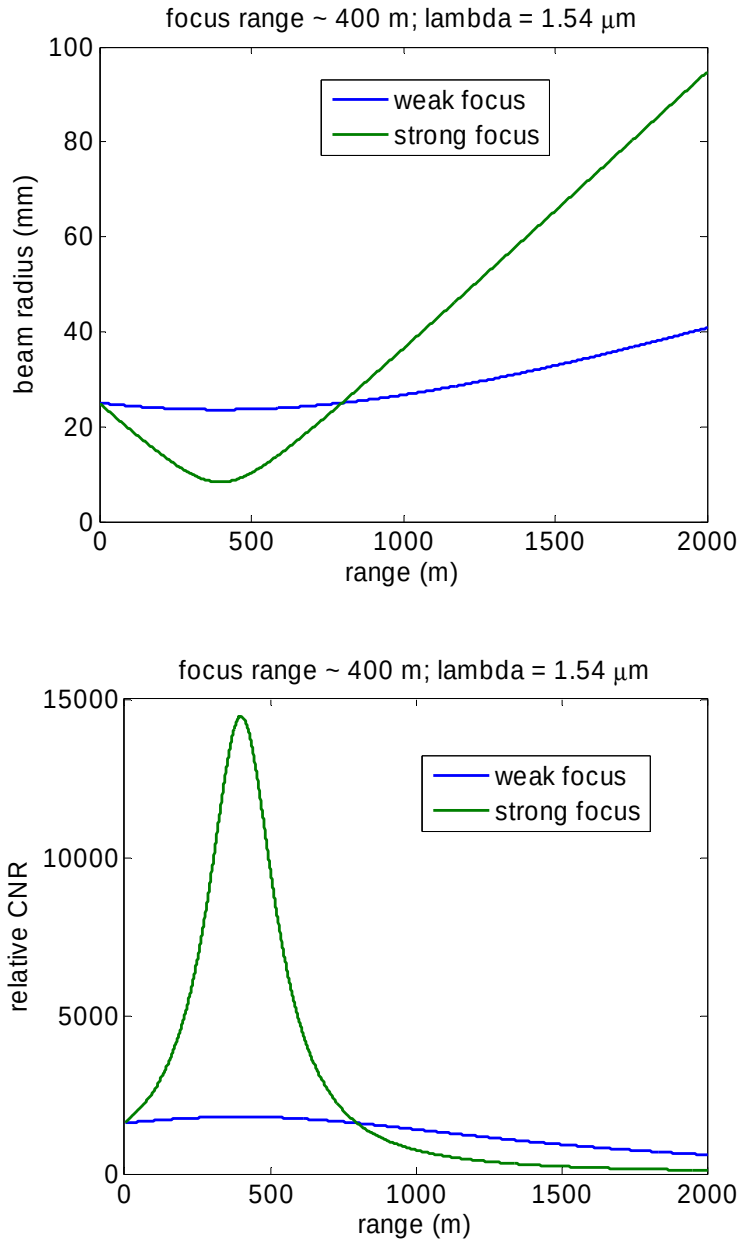

Figure A6: Lidar beam parameters for weak and strong focus at r_{waist} = 400 m, with nominal WindCube WLS-70 parameters of λ= 1.54 μm and w(r = 0) = 25 mm. Top: w(r), Gaussian beam radius as a function of range. Bottom: w^{-2} variation with range for the carrier (C) term in the lidar CNR.

The “standard” and “long range” versions of the pulsed 1.5 µm **Galion** wind lidar have specifications very similar to those of the WLS-7 and WLS-70: “standard” means 40 m to 250 m, and “long” means 80 m to 2000 m. The focus choices are thus very similar [2], and the Galion focus is adjustable by users.

The optical parameters of two HALO Photonics lidars are described by O’Connor et al. [3]. One was “configured for a primary function of observing liquid and ice cloud at all heights up to 10 km”, and its transmitted beam was therefore collimated (= focus at infinity; the beam waist lies at the exit aperture). The second lidar, more relevant to UpWind, was “optimised for boundary-layer studies” and “achieved the required sensitivity with a shorter integration time by having the telescope focus set to approximately 800 m (note that this reduces the instrument sensitivity dramatically above 2 km)…a maximum range of about 2 km”.

The tabulated parameters for this second lidar, as deployed in the 2007 REPARTEE (Regent's Park and Tower Experiment) trials in central London, are: wavelength 1.5 µm, range resolution 30 m, PRF 20 kHz, 20000 pulses averaged per velocity-range profile, lens diameter 8 cm, divergence 33 µrad (full-angle; note that θ in Eqn (5) in [3] is a half-angle divergence), and focus range 801 m. The CNR will behave roughly as:

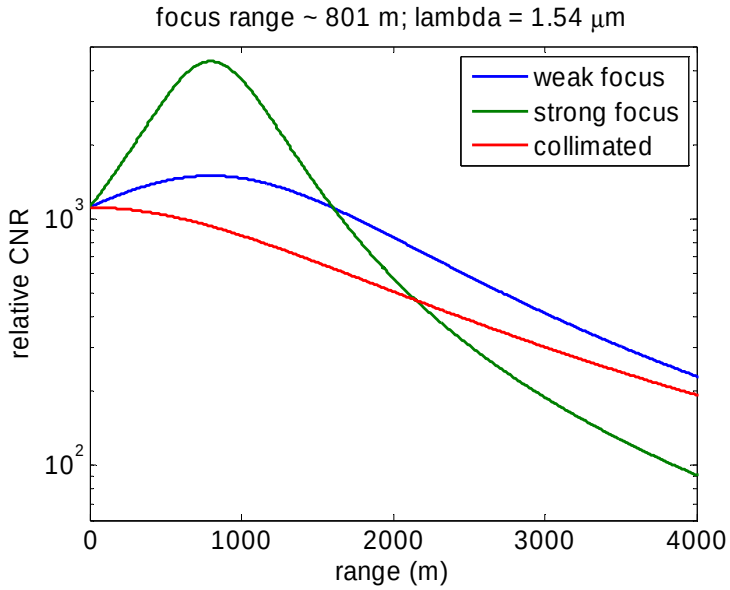

Figure A6: CNR variation with range for a design similar to the HALO lidar fielded at REPARTEE [3]. At the focus range ~ 800 m, the CNR is nearly three times higher for the strong focus than for the weak focus, but for focus ranges beyond 2 km the weak focus or the collimated beam is preferred. Note the logarithmic scale for CNR.

The CNR behaviours for the strong-focus and collimated-beam configurations differ significantly, even for a focus range as long as 800 m. The difference between collimation and weak focus is smaller (~ 2 dB) but may still be considered worthwhile.

None of these curves provide tight range resolution; their -3dB widths are hundreds of metres, i.e. many typical range bins. To repeat: in these pulsed systems

the range-gate processing is the most important influence on resolution. The overall variation indicated by the CNR curves is still important, but it is relatively weak, and in a CW system it would not normally be thought sufficient.

Figure 3 in [3] shows the estimated power spectrum of the fluctuations in vertical velocity, at a height of 90 m ("taken from 3 adjacent range gates at 60-120 m"), for one hour of data (~ 2700 individual velocity measurements).

Conclusion

The accuracy of a wind speed estimate depends on the lidar carrier-to-noise ratio (CNR), and the choice of lidar focus geometry strongly influences the variation of CNR with range. Illustrations of these Lorentzian curves and lidar CNR equations can be found in [4] and many other papers.

This Appendix has described the design choices for today's most common commercial wind lidars, but not how the manufacturers verify that the beam geometry in practice is the same as the design choice. Such procedures are described for ZephIR in the UpWind lidar calibration report; for pulsed systems the checks are potentially more difficult because of the tens-of-metres minimum-range requirement.

The assumptions of Gaussian beams, negligible aperturing etc. should hold rather well for these single-mode fibre-optic lidars, so that (when calibration tests are defined and conducted) we expect good agreement between measured and predicted CNR. Significant discrepancies should cause concern. It is wise to check the assumption of "complex Gaussian" statistics; the beam size and probe volume may change considerably with target range, but for every possible range we wish the detector to see contributions from "many" independently phased scatterers. Thus any calibration target should be carefully chosen and treated.

Typically the lidars are robust and they maintain their initial alignment: once the geometry has been fixed, we usually find that the experimental error in the CNR value, and the contribution of that error to the uncertainty in final wind speed, are much smaller than other errors in the process (notably those induced by turbulence and by the spread of velocities within the sensing volume).

References for Appendix

[1] J-P Cariou, personal communication.

[2] G N Pearson, personal communication.

[3] E J O'Connor, A J Illingworth, I M Brooks, C D Westbrook, R J Hogan, F Davies and B J Brooks, "A method for estimating the turbulent kinetic energy dissipation rate from a vertically-pointing Doppler lidar, and independent evaluation from balloon-borne in-situ measurements". Submitted to J. Atmos. Oceanic Technol.
http://www.met.rdg.ac.uk/clouds/publications.html

[4] M J Kavaya and P J M Suni, "Continuous wave coherent laser radar: calculation of measurement location and volume", Applied Optics **30** 2634-2642 (1991).