

Article

Coherent Focused Lidars for Doppler Sensing of Aerosols and Wind

Chris Hill

Malvern Lidar Consultants, Great Malvern, Worcestershire WR14 1YE, UK; malvernlidar@gmail.com

Received: 16 January 2018; Accepted: 21 February 2018; Published: 16 March 2018

Abstract: Many coherent lidars are used today with aerosol targets for detailed studies of e.g., local wind speed and turbulence. Fibre-optic lidars operating near 1.5 μm dominate the wind energy market, with hundreds now installed worldwide. Here, we review some of the beam/target physics for these lidars and discuss practical problems. In a monostatic Doppler lidar with matched local oscillator and transmit beams, focusing of the beam gives rise to a spatial sensitivity along the beam direction that depends on the inverse of beam area; for Gaussian beams, this sensitivity follows a Lorentzian function. At short range, the associated probe volume can be extremely small and contain very few scatterers; we describe predictions and simulations for few-scatterer and multi-scatterer sensing. We review the single-particle mode (SPM) and volume mode (VM) modelling of Frehlich et al. and some numerical modelling of lidar detector time series and statistics. Interesting behaviour may be observed from a modern coherent lidar used at short ranges (e.g., in a wind tunnel) and/or with weak aerosol seeding. We also review some problems (and solutions) for Doppler-sign-insensitive lidars.

Keywords: coherent Doppler lidar; wind sensing; single-particle

1. Introduction

Near-infrared coherent lidars are familiar in anemometry and turbulence sensing. Their behaviour has been fairly well understood and modelled since the 1960s. Fibre-optic versions are increasingly used in the wind power industry for aerosol targets, and are also becoming popular for solid targets such as vibrating or rotating machinery, structural panels, and turbine blades. Several tutorials and reviews have recently been published, aimed at the growing number of readers—not necessarily optical or laser specialists—who need to understand the main features and limitations of these sensor tools. As they become widely used and extended to different technical areas, it is sometimes necessary to return to the fundamentals and check that older radar/lidar lessons are correctly translated and applied.

A recent review of modern fibre-optic lidars [1] discussed three main points:

- (1) The most common references in the literature of “coherent continuous-wave focused monostatic lidars” are now some 25–40 years old but still worth reading. In particular, the carrier-to-noise analysis of Sonnenschein and Horrigan [2] agrees with alternative treatments based on the popular “antenna theorem” or “back-propagated local oscillator” (BPLO) approach. Their analysis applies to ZephIR and similar modern fibre-optic aerosol lidars.
- (2) Such lidars work over a large range of conditions and spatial scales; for example, the so-called “probe volume” of a variable-focus lidar may easily vary over eight orders of magnitude. There may also be large variations in scattering particle density and average atmospheric backscatter. The assumption that the probe volume contains “many” scatterers can lead to simple mathematics (Gaussian statistics for long random walks where the central limit theorem

holds) and is extremely common but can be faulty (e.g., with very clean air or short measurement range).

- (3) The desire for simple descriptions or “sensor performance metrics” conflicts with the complications of real lidar measurements. For example, the expressions “range resolution” and “bandwidth” have multiple meanings, and it is difficult (and often confusing) to characterise a lidar's performance by a single value. In a well-known sense, the axial resolution of a coherent CW focused monostatic lidar is a Lorentzian function with scale parameter equal to the beam Rayleigh range. In another important sense, the lidar can “resolve” scattering events with much finer range precision.

The present paper concentrates on aerosol scattering and a common commercial application: the “lidar Doppler” estimation of aerosol/lidar relative velocity and thus (by using several or many estimates) of wind flow and wind patterns. We start with a brief review of the beam geometry for a standard coherent lidar (Section 2). Sections 3 and 4 discuss the detector output for direct and heterodyne detection respectively. Section 5 has comments on the large preceding literature and reviews some disagreements about “lidar collection efficiency”, which is one aspect of the dependence of carrier-to-noise (CNR) on beam geometry.

Section 6 discusses a computer simulation of multi-scatterer experiments. Section 7 describes the important practical constraint of sign ambiguity for moving targets and illustrates how, even if I&Q data or other indications of sign are not immediately available, that ambiguity can be removed for typical aerosol targets. For the common conical-scan or sector-scan geometry and its associated VAD (velocity-azimuth display) output, Section 8 discusses examples of measurement bias.

2. Lidar Geometry

Consider the sketches in Figure 1.

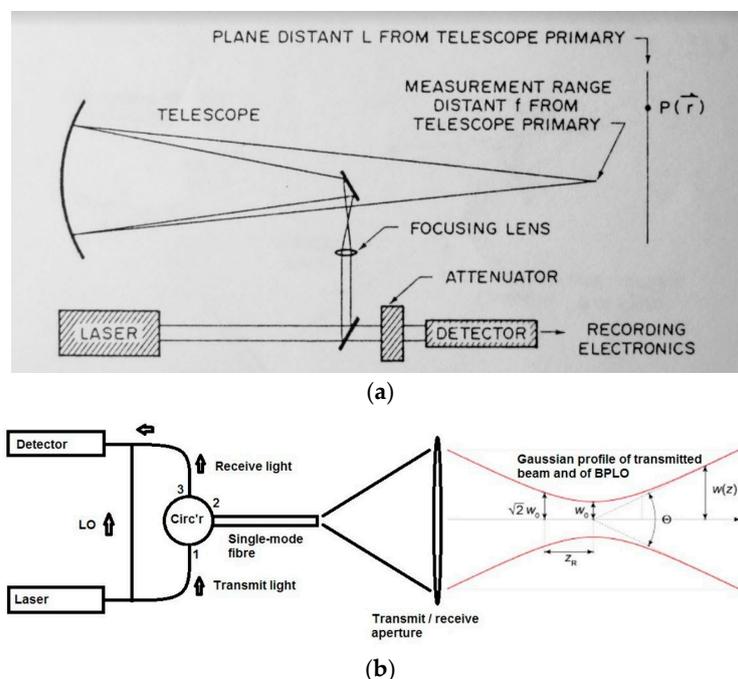


Figure 1. Schematics of laser Doppler systems for remote wind sensing. (a) Reproduced from Lawrence et al. [3] (“A laser velocimeter for remote wind sensing”, Rev. Sci. Instrum. 1972, vol. 43, pp. 512–518) with the permission of AIP Publishing. This is concerned with Doppler measurements of wind-borne scatterers in the atmosphere at relatively short ranges, but is more widely relevant; (b) From Hill [1], concerned with modern fibre-optic lidars. See also below for Figure 2 and the original diagram in Lindelöw [4].

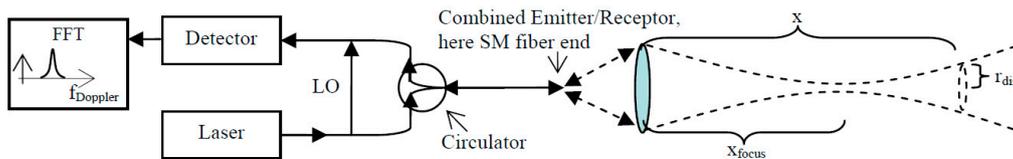


Figure 2. Schematic of focused monostatic CW coherent lidar based on fibre-optics components. Distance to focus is x_{focus} , and distance to target scatterer is x ; these are called f and L in Figure 1a. Reproduced from Lindelöw [4].

The beam is brought to focus at a distance f from the telescope aperture (mirror or lens), and the scatterers are carried across the beam by the wind at a possibly different distance L , with a crosswind velocity component of $V \text{ ms}^{-1}$.

Let the beam be a lowest-order Gaussian beam with a field described by

$$E(x, y, z) = \frac{E_0}{w(z)} \exp\left[-\frac{(x^2 + y^2)}{w(z)^2}\right] \exp\left[jk \frac{(x^2 + y^2)}{2R(z)^2}\right] \tag{1}$$

$$w(z) = w_0 \sqrt{1 + (z/z_R)^2} \tag{2}$$

is the “beam $1/e^2$ intensity radius” at axial distance z from focus

$$R(z) = z + z_R^2/z \tag{3}$$

is the beam radius of phasefront curvature at z

$$z_R = \pi w_0^2/\pi \tag{4}$$

is the beam parameter or Rayleigh range (often notated b or $b/2$)

$$I(x, y, z) = |E(x, y, z)|^2 = \frac{E_0^2}{w(z)^2} \exp\left[-2 \frac{(x^2 + y^2)}{w(z)^2}\right] \tag{5}$$

Some constant factors and the extra on-axis (Gouy) phase shift have been omitted.

According to Equations (1) and (2), the beam extends indefinitely in any XY cross-section plane, but its field strength decays exponentially, with scale parameter $w(z)$. The beam also extends indefinitely along the z axis, but its field strength decays according to $1/w(z)$, with scale parameter z_R . For a uniformly and densely seeded atmosphere, we are usually justified in neglecting the contributions from scatterers that are more than a few beam widths off-axis or more than a few Rayleigh ranges away from the beam focus—see below. For a single scatterer, of course, we are interested in the lidar response for any scatterer position if it is detectable.

For such a Gaussian beam of intensity $I(r) = I_0 \exp(-2(x^2 + y^2)/w^2) = I_0 \exp(-2r^2/w^2)$, illuminating one thin “slice” between ranges z and $z + dz$, some simple properties are

- $P_0 = I_0 (\pi w^2/2)$ is the total power, I_0 is the on-axis intensity, and w is the e^{-2} intensity radius of the beam.
- The power enclosed by a circle of radius h is $P_0 \cdot [1 - \exp(-2h^2/w^2)]$. If h equals $2w$, about 99.97% of the total power is enclosed. This is effectively all the power for most practical purposes, and the mean intensity within the circle of radius h is nearly $P_0/(\pi h^2)$ if $h \geq 2w$.
- The probability distribution for the illumination intensity I is $p(I) = \frac{1}{I} \left(\frac{w^2}{2h^2}\right)$ with normalisation to unit probability within the circle of radius h (i.e., the intensity lies between $I_0 \exp(-2h^2/w^2)$ and I_0).

Similar but z -averaged properties may be derived for a focused Gaussian beam whose width parameter $w(z)$ varies along the beam according to Equation (2). However, the mean intensity and the $p(I)$ then depend on the z integration limits, which may not be symmetric around the focus distance. For example, a “collimated” exit beam has a plane wavefront initially; the beam waist lies at the exit aperture ($z = 0$) and the beam expands as it propagates through positive ranges z ; negative values of z are not relevant because they lie inside or behind the lidar.

Note that our Rayleigh range is defined here as $z_R = \pi w_0^2/\lambda$, so that the beam radius $w(z = z_R)$ is $\sqrt{2}$ times the beam waist radius w_0 . Alternatively, the terms “beam parameter”, “confocal distance”, and “Rayleigh range” are sometimes defined as $2\pi w_0^2/\lambda$ to refer to the distance between the plane where $w(z = -z_R) = \sqrt{2}w_0$ and the corresponding plane on the other side of the waist where $w(z = z_R) = \sqrt{2}w_0$.

3. The Detector Output (Direct Detection)

Consider the average backscattered power reaching the receive aperture: a quantity proportional to the “photon count” due to scattering from all illuminated regions [5].

If there are N scatterers contributing, the instantaneous power is proportional to the magnitude-squared of the sum over N phasors (one from each scatterer). A time or ensemble average power, taken over all relative phases of the individual scatterer reflections (all phases being assumed equally probable), is proportional to the sum of N positive terms, one for each scatterer; each term depends on the scatterer’s position in the Gaussian illuminating beam, and also has an inverse-square dependence of the strength of the spherical scattered wave as it propagates back to the receiver. Cross-terms average to zero, regardless of N ; the averaging to zero results from the random uniform distribution of phases and does not require $N \gg 1$. If we assume for simplicity that the fraction of scattered light received within the aperture (and falling on the detector) is the same for all scatterers at a particular range, that is usually a serious and restrictive assumption. The precise dependence on scatterer position may not be easy to express, if we wish to consider scatterers in the near field—for example, a thin slice so near the lidar aperture that the solid angle subtended by the aperture is not the same for all scatterers in the beam. We can only mention here a large literature on the detailed geometrical “form factors” and “collection efficiency” of direct-detection lidars [6].

This instantaneous power and this average power would be seen in direct-detection mode, where we assume that photoelectrons are generated independently from every small sub-area of the detector, and the total photoelectron current is (in the limit of large count) an accurate, “light-in-a-bucket” record (subject to some detector response bandwidth) of the time-varying optical power that falls on the whole detector surface. For the moment we can consider that a fixed (usually large) fraction of the backscattered laser light that falls within the receive aperture reaches the detector. This backscattered light forms an interference pattern due to its N randomly phased components (a Gaussian speckle pattern, in the limit $N \gg 1$). Using a single-mode optical fibre in the receive channel of a direct-detection lidar has consequences. One advantage is that we know where we are looking (there is a well-defined “receive antenna” pattern—not to be confused with the phase-sensitive heterodyne-lidar antenna pattern below) but a disadvantage is that much light entering the receive aperture is lost because it does not match the fibre mode at the entrance facet; another consequence, often undesirable in direct detection, is that intensity fluctuations are not averaged out.

With this catch-all or “bucket” aperture, the average direct-detection backscattered power is proportional to the sum over the individual N illumination intensities—for any value of N or any seeding density; strictly, this holds without any lower limit, and we assume we never reach the other extreme where the scatterers are so dense that multiple scattering is significant.

In terms of the standard random-walk problem on a complex plane, we have at any one point (or any one sufficiently small detector element) the coherent summation of N randomly directed vectors [1]. If each vector has the same length a , the total “intensity”—the mean value of the square of the resulting distance from the origin—is $\langle I \rangle = Na^2$. The second moment of intensity is $\langle I^2 \rangle = 2N(N - 1)a^4 + Na^4$, which for large N approaches $2(Na^2)^2 = 2\langle I \rangle^2$. The “normalised second

moment" is thus $\langle I^2 \rangle / \langle I \rangle^2 = 2 - (1/N)$. If the vector lengths are drawn from a random distribution [7], we can use the following result:

$$\frac{\langle I^2 \rangle}{\langle I \rangle^2} = 2 \left(1 - \frac{1}{N} \right) + \frac{1}{N} \frac{\langle a^4 \rangle}{\langle a^2 \rangle^2} \quad (6)$$

Many further expressions for the intensity moments of various orders have been studied, often through the indirect mathematical methods of "generating functions". The probability distribution of intensity $P(I)$ tends to be less tractable; few of the results, in terms of integrals of Bessel functions, reduce to useful analytical expressions. (Note that we use $p(I)$ above to describe the spatial variations, and $P(I)$ here to describe the time series at a detector point). For an illuminated collection of N non-fluctuating scatterers, with N "moderately large", Pusey et al. [8] expressed $P(I)$ as the Gaussian-limit exponential $\exp(-I/\langle I \rangle)/I$ plus a series of stated correction terms. Some of the integrals involved in random-walk analysis are difficult to compute with high precision, but this is not of major importance for lidar users.

A many-scatterer Gaussian-statistics limit of the random walk model should not be assumed without thought; also, the definition of "CNR" depends on whether we include times when no scatterers are present. In practice, perhaps with rather detailed "time-frequency" post-processing, we may increase the "CNR" by discarding data from such times and by discarding frequency bands which do not contain the (usually chirping) lidar returns.

Because of the Gaussian beam's circular symmetry, the backscattering from the full 3D scatterer-filled volume is usually treated by an integral over a set of thin discs (with radial scale $w(z)$ and extending from z to $z + dz$) considered to fill the space on one side of the lidar aperture. We note that scattering from a collection of particles all confined to a very thin disk (or spherical shell etc.) may not satisfy the requirement of uniformly distributed phase [7].

4. The Detector Output (Heterodyne Detection)

Now consider the heterodyne operating mode by adding the LO (and BPLO) to the picture, with no other optical change; the backscattered light received at the detector is the same as before—say $P_s(N)$ for N scatterers—but in general, some fraction $(1 - \eta_{\text{het}})$ of this incident light will not be mode-matched with the LO and will not contribute to the heterodyne current, whose power is proportional to $\eta_{\text{het}} P_s(N)$, and the local oscillator power P_{LO} . The overlap of incident light and LO light can be evaluated at the detector surface or (by a considerable extension of the original far-field "antenna theorem") at any other convenient plane. When we evaluate it in the target plane, the (time-averaged or ensemble-averaged) contribution to the heterodyne output power made by any scatterer is proportional to the intensity of the transmitted beam at the scatterer coordinates and also to the BPLO intensity there. Since we have assumed that transmit and BPLO beams are matched, we can describe the spatial variation equally by the product of transmit and BPLO intensities or by the square of the transmit intensity.

In direct detection, each small element of the detector produces its own photocurrent—which can be described by a time series of non-negative real numbers. In heterodyne detection, each small element produces a (usually dominant) shot noise photocurrent, plus a *modulation*—that is, an extra term (which oscillates at the heterodyne frequency, taking both positive and negative values as the interference fringes evolve). The net result of summing the outputs from the various detector elements is thus phase-sensitive: the total photocurrent is never negative, but the coherent (phase-sensitive) sum of the additional modulations—which is what interests us—can be positive or negative.

Once the lidar is shot-noise-limited, any further increase in LO power causes a proportionate increase in shot noise power and does not change the CNR; that is, any multiplying constant of LO power cancels top and bottom in a C/N expression.

The heterodyne current due to several scatterers is the vector sum of the individual currents, one for each scatterer; each term is proportional to the beam intensity and to the BPLO intensity at the scatterer. The mean square of the current (averaged over all relative phases as above) equals the

sum of the individual squares; this is the same result as for the mean square of the length of a 2D random walk.

Consider a single scatterer blown transversely through the beam along a straight line at uniform speed V : its z and x values remain essentially constant, and its y value equals $(t - t_0)V$. As it traverses the beam it produces a detector output current with an intermediate-frequency (IF) component.

$$i(t) = 2R\sqrt{\eta_{\text{het}}P_{\text{LO}}P_s(t)} \cos(\omega_c t + \theta(t)) \quad (7)$$

where R = detector responsivity, ω_c = offset radian frequency, and θ = phase shift. The time variation of $i(t)$ includes variations in some or all of these parameters. Here we may take ω_c to be constant (there is no extra bulk target motion along the z -axis line of sight) and θ to express the time-varying Doppler phase shift. We may also take P_{LO} to be constant (the total local oscillator power at the detector) and P_s to be the backscattered power reaching the detector (not constant). The scalar η_{het} represents, for this single scatterer, the “heterodyne efficiency” with which the LO and backscattered light overlap at the detector. Note again that we write η_{het} proportional to the square of output current, i.e., to output power. In the literature on heterodyne lidars, there are various preferences for splitting the “efficiency” into several terms identified with different parts of the total system [1,9].

Usually, Equation (7) is taken to say that P_{LO} is a fixed quantity without spatial variation, and its value (in units of Watts)—the LO power that reaches the detector—is adjustable at our discretion; whereas, although the transmitted laser power is also a fixed number of Watts at our discretion, the value of $P_s(t)$ depends on other factors such as optical losses, atmospheric attenuation and the position of the scatterer in the beam. In practice, $P_s(t)$ will vary deterministically as the scatterer moves across the Gaussian intensity beam profile, reaching some peak value (see Section 5.4) when the scatterer makes its closest approach to the beam axis. The value of η_{het} for the particular distribution in space of scattered light at the detector (which will usually vary with time), and the particular (usually fixed) LO light distribution in space, is then a further matter of fact or calculation.

That is, in our notation, the spatial dependence in Equation (7), as the scatterer traverses the beam, brings a time dependence to P_s , and possibly to η_{het} and θ , but not to P_{LO} . This heterodyne current component $i(t)$, due to a single scatterer at position $y(t)$, is proportional to the transmitted beam intensity at the scatterer, i.e., the spatial variation is described by $\exp\left[\frac{-2(x^2+y^2)}{w(z)^2}\right]$. Note that the transverse variation of the transmitted beam is the Gaussian in Equation (1); the local oscillator (and hence the BPLO) will be assumed to keep a perfect copy of this Gaussian shape. Equation (14) in the early paper by Sonnenschein and Horrigan [2] describes “the square of the signal current produced by a single scatterer”, and a corrected version is

$$|i_s|^2 \propto \frac{R^4}{\lambda^2 L^4 \left[1 + \left(\frac{\pi R^2}{\lambda L}\right)^2 \left(1 - \frac{L}{f}\right)^2\right]^2} \exp\left\{\frac{-4\left(\frac{\pi R r}{\lambda L}\right)^2}{\left[1 + \left(\frac{\pi R^2}{\lambda L}\right)^2 \left(1 - \frac{L}{f}\right)^2\right]}\right\} \quad (8)$$

L is the range to target, f is the range to beam focus, their R is the transmitted beam radius (our $w(z)$ at the aperture where $z = -f$), and r is off-axis distance ($r^2 = x^2 + y^2$).

In [1], we reviewed some of the differently notated but essentially similar versions of Equation (8) in the literature and illustrated the frequency-chirp behaviour of scatterers that traverse the beam. The product of chirp duration and chirp slope is roughly

$$\text{chirp excursion (Hz)} = \frac{\pi w(z)}{2V} \frac{2V^2}{\lambda R(z)} = \frac{\pi V w(z)}{\lambda R(z)} \quad (9)$$

The chirp changes sign if the wavefront curvature changes sign (i.e., if we consider events on one side of the beam focus and then the other). As we move away from the beam waist $z = 0$ (in

either +ve or –ve direction), the absolute value of $R(z)$ at first decreases, then reaches its minimum of $R(|z| = z_R) = 2z_R$, then increases again. So, on each side of the beam focus, there are in principle two distances z corresponding with any given value of R greater than $2z_R$. These two distances are associated with the same chirp slope $d\theta(t)/dt$ but different beam widths $w(z)$ and thus different envelope durations and different chirp excursions.

5. Previous Literature

5.1. Semiclassical Account of Laser Radar

There are large relevant literatures on coherent lidar, photon correlators, laser Doppler velocimeters (LDVs) and laser transit velocimeters (LTVs), electromagnetic scattering from small particles, interference effects in the presence of more than one scatterer, and so on (see for example [1,7,9], and their references). From an optical-radar viewpoint, the function of the scatterers is to provide, at the receiver, copies of the transmitted waveform [10]; when all the different copies are considered (with their various delays, attenuations, polarisations, frequency shifts etc.), we have a total field at any given detector element whose intensity at any instant is proportional to the mean rate of photoelectron production. That mean rate is vastly increased by the strong steady LO. There are typically very many photoelectrons, but each has a random (Poisson point process) time of origin. There is no one-to-one connection of individual photoelectrons with individual “photons” in this typical semiclassical account of a shot-noise-dominated coherent lidar. It is generally assumed that the photocurrents are so strong that, in line with our discussion above, the “full phase and intensity” information in the optical field is indeed transferred intact to the heterodyne detector output; for example, a very weak FM sideband (due to a faint micro-Doppler vibration) can still be isolated and examined and assigned a conventional SNR that is negligibly degraded by the (intermittent, discrete, Poisson) nature of the photoelectrons.

At power levels several orders of magnitude weaker, when we approach single-photon detection, this picture must change, but we retain it here. We neglect many complexities of scattering theory, vector wave effects, polarisation, and detector physics. But it is worth reviewing the detector output current (above) and referring now to some literature including frequently quoted “fundamentals of coherent lidar” papers.

The strength of the heterodyne current is usually judged by the average modulus-squared $\langle |i(t)|^2 \rangle$. (The average is taken over a time exceeding the longest fluctuation time, and there are usually several types of fluctuations present). This strength is determined by how well the scattered light and the LO overlap: the overlap integral, or antenna efficiency, includes both the magnitudes and the relative phases of the signal and LO terms. The need for transmit and receive antennas to be “matched” was familiar from earlier radar work, and was quickly imported and applied to lidar studies in the 1960s and 1970s. A large literature developed on various sub-topics such as:

- The “best” designs of telescopes and truncating apertures, according to several different metrics of efficiency;
- The benefits of angular selectivity or directionality of heterodyne antennas—and the corresponding requirement to establish and maintain precise alignment in practice;
- The statistics of detector outputs for various types of target: solid, liquid, or gas; few scatterers or many scatterers; concentrated or distributed in range; static or moving;
- The effects of one-way and two-way atmospheric turbulence;
- The differences between monostatic (shared apertures and collinear beams) and bistatic or multistatic lidars.

The literature inevitably swelled as different approaches and notations were developed and published by groups working on laser systems in private companies, universities, and government organisations.

5.2. Local Oscillator (LO) and BPLO

One important sub-topic is the “back-propagated local oscillator” or BPLO, which was presented by Siegman [11] as follows:

Consider the complex LO amplitude distribution falling on the photodevice surface (weighted by the quantum efficiency distribution if necessary). Reverse the direction of propagation of this LO distribution and allow the reversed wavefront to propagate back out through any optical elements that an incident signal wave would traverse. The resulting far-field or Fraunhofer diffraction pattern will be the antenna pattern of the optical heterodyne receiver.

The general principle is that we may reverse the direction of propagation of the wavefront and perform the overlap integral in any convenient place—not just the detector surface, but for example the telescope aperture or the plane(s) of the target(s). Two aspects of this BPLO approach may be especially relevant to modern wind lidars [1]. One is the advantageous choice of overlap plane. We are free to evaluate the overlap at different places: it may be that some, from a practical, computational point of view, are better than others. Zhao et al. [12] compare two expressions for a transmission function that relates a point in the scattering plane to a point in the detector plane. The light successively encounters a primary mirror, a secondary mirror, “a series of optical components, such as steering mirrors and polarizers”, a detector lens, and the detector. The virtual BPLO encounters these in the reverse order, and the corresponding expression is derived by “changing the order of integrals and invoking the reciprocity theorem”. Formally, the expressions are equivalent, but in the BPLO one, “the result of the first several steps of integration is common to all points in the scattering plane and needs only to be calculated once . . . In addition, if the system is well aligned and free from astigmatism, circular symmetric properties of the integrand further simplify the integration to a 1-D calculation . . . Thus the BPLO treatment greatly reduces the amount of computation”.

The amount of computation needed for these integrals, and any “computational advantage” of the BPLO approach, typically decrease when we change to fibre-optic lidars—because they have fewer components and obstructions, and propagation within single-mode fibre needs negligible extra calculation. That is, the non-BPLO approach may be less tedious in fibre lidars than in free-space lidars, although still more tedious than the BPLO approach. We can calculate everything at the detector if we wish; nothing forces us into BPLO calculations.

The second point is a difference between fibre-based lidars and the more familiar free-space optical systems. Single-mode fibres act as spatial filters. They support only one spatial mode (transverse mode), so any light that arrives at a fibre entrance plane in other modes does not propagate any significant distance along the fibre and does not reach the fibre exit—it is lost. Similarly, any virtual BPLO light “arriving” at the fibre from the detector contributes nothing to the antenna pattern (and the overlap integral) unless it matches the fibre mode. The spatial form of the BPLO, once it leaves the fibre and continues through any transmit optics toward the target, is always that of the launched single transverse mode of the fibre. In particular, it is unaffected by apertures or obstructions between the fibre and the detector or by spatial variations in detector response. Nothing we do before the fibre can affect the antenna pattern after the fibre.

To the extent that detector output statistics are affected by the number and nature of the scatterers in the probe volume, this is an important difference between single-mode-fibre lidars and most free-space lidars. If we damage or partly block the fibre-lidar detector, then the probe volume (the physical extent and shape of the region contributing to the spatial overlap), and the carrier statistics, will not change. In free-space lidars there are often apertures, obstructions, and detector imperfections; these are often inaccessible and/or hard to adjust, yet they strongly influence the probe volume and statistics. Moreover, the practical difficulty of measuring detailed beam properties (such as BPLO shape and phase) causes uncertainty about what parts of a target are contributing, and how strongly. The ease of LO and BPLO alignment in fibre sensors, and the enforced limitation to one well-defined mode, reduce this uncertainty.

The situation is symmetrical; we could as easily have said “Nothing we do in free space can affect the illumination pattern at the detector—it is determined by the propagation of the single mode from the fibre end to the detector surface, through any distorting elements that may be present internally”.

Zhao et al.’s caveats still apply: “The LO field at the detector should also be calculated very carefully . . . the field at the waist is usually different from an ideal Gaussian distribution . . . the criterion for ignoring the diffraction effect of the sharp-edged apertures for a Gaussian beam is quite stringent”. The point is that the presence of a truly single-mode filter, somewhere in the optical chain between detector and target, defines the single mode that is relevant (e.g., in diffraction calculations) for both internal and external regions.

We neglect the possibilities of systems that are not reciprocal, e.g., the beams have significant frequency differences, and some components (such as modulators, amplifiers, and regenerative or self-aligning cavities) are frequency-selective or dispersive [3].

In practice, there are also specific, small, but possibly significant departures from ideal theory when fibres are involved:

- The fundamental mode of a single-mode fibre is normally modelled as a free-space Gaussian TEM_{00} in spatial profile, and we make this assumption here; in practice the match is very good but not perfect;
- Fibre-pigtailed collimators (FPCs) can improve the balance between lateral and angular misalignment effects by increasing the effective TEM_{00} radius at the coupling plane [13]. Lens focal length and lens-fibre distance will change accordingly, but otherwise, our efficiency calculations are not affected—the FPC/fibre combination behaves as a single-mode fibre with, at one end, a larger mode area;
- Higher-order terms in the mode field expressions imply that the wavefront in the fibre is slightly curved instead of plane. This curvature is familiar in hollow waveguide physics [14] and leads to measurable asymmetries in waveguide/beam coupling experiments, but we neglect it here.

There is a long history of cross-checks and calibrations of coherent lidars, with hard and distributed targets, and a sometimes baffling range of discrepancies. Experienced scientists, trying hard to account for all terms in the carrier-to-noise equations, still fall short by factors of around 2; indeed, Kavaya [15] noted that CNR theory and experiment often disagreed by nearly 3 dB, and not by a random variable factor ranging from say -3 dB to 3 dB. Note that this shortfall is in CNR (or sensitivity), and not necessarily in accuracy, probe length etc.

5.3. Collection Efficiency for Coherent Lidars

Evidently, then, in modern coherent lidar sensors the nature of the beam overlapping that determines optical efficiency and carrier-to-noise (and their variations as functions of range to target) is essential to lidar calibration and operation.

Some disagreements about this in the wind lidar literature were discussed within the UpWIND project [16–19]. First, Lindelöw [4] reconsidered the range weighting of wind lidars and proposed a redefinition of “probe length” based on the WPP (“wind peak profile”) function; this is relevant to pulsed lidars that are range-gated (whether focused or not). These two issues of focus and range-gating are often present together, but it is possible to separate them and check the predicted discrepancies. For incoherent (randomly phased) returns from many aerosol scatterers, with the algorithms of signal processing and Doppler estimation that are most commonly used, the WPP proposal has not been widely accepted and has not altered the current standard treatments. It is mentioned here because “probe length” or “range resolution” is often presented as a figure of merit (that is, the shorter the better), and thus enters discussions of whether one lidar or another is “better”, or whether pulsed or CW operation is “better”. Readers should check which definition is being used.

A second topic is the “fibre lidar collection efficiency” relevant to focused monostatic coherent lidars, both pulsed and CW. Most published models use a Lorentzian for this efficiency or sensitivity

function, whether or not the scaling with “1/area” (which is a Lorentzian for a Gaussian beam, with the conventions above) is recognised explicitly. Lindelöw [4] and then Lindelöw with Risø/Leosphere colleagues [17] preferred a function that is spatially narrower (tighter) than the Lorentzian. The discrepancy becomes larger at longer focus ranges, amounting to ~25% difference at typical large-turbine sensing heights of 100–150 m.

Lindelöw’s model [4] “takes into account small receptor apertures and co-propagation of the local oscillator and the received backscatter in a single mode fiber . . . A focused lidar will transmit a narrow beam of light with a waist at distance x_{focus} , typically at 20–200 m. A sketch of a fiber optic based focused monostatic coherent lidar is presented in Figure 31. The lidar has a combined emitter/receptor in the form of a fiber end positioned in front of the focusing lens”. This “Sketch of system and principle of a focused monostatic cw coherent lidar based on fiber optic components” is our Figure 2 below, similar to Figure 1b above.

If we consider this as an imaging system, then an object or source at range x (in this case a thin illuminated disk of air) creates an image on the other side of the lens. In Lindelöw’s treatment, the first stage of light propagation (from fibre through lens to this target object) and the last stage (just before recoupling to fibre) are described by standard Gaussian beam equations, but there is an intermediate stage (backscattering by the target, then formation of an image) under different assumptions and using thin-lens equations.

This mixing of Gaussian beams and rays has consequences illustrated in Figure 3.

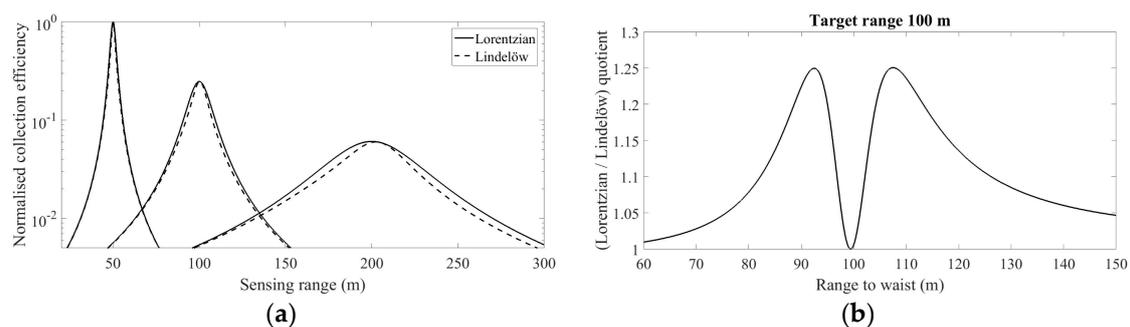


Figure 3. Two models for heterodyne lidar collection efficiency. (a) Normalised curves of range dependence for three different fixed focus ranges (50 m, 100 m, 200 m). Lorentzian (solid lines): efficiency scales as the reciprocal of beam area. Lindelöw (dashed lines): modified Gaussian optics/ray optics model proposed by Lindelöw [4]; (b) The quotient (Lorentzian efficiency)/(Lindelöw efficiency) for a fixed target range of 100 m and a varying focus range. Wavelength 1.575 μm , equivalent fibre core beam radius 4.5 μm , and lens focal length 200 mm (roughly representative of QinetiQ/ZephIR lidar).

The two “collection efficiency” functions are normalised so that their peak heights are equal for short focus range (where the differences between ray optics and Gaussian beam optics are negligible). This single normalisation factor is then applied for the other focus ranges. Any one curve in Figure 3a shows the efficiency for the given choice of focus (e.g., 100 m) as the target range, and only the target range, is varied. The Lorentzian curves peak at exactly the selected focus ranges (by construction, because our “focus range” means the distance where beam area is smallest); the curves for Lindelöw’s model peak at slightly more distant ranges.

Note that the focus distance is assumed to be changed by varying the fibre-to-lens distance as mentioned above [20]. This means a slight but not always negligible dependence of beam size at the lens on x_{focus} . The exact form of the curves depends on this and on the choice of normalisation, but a comparison of the two approaches is not affected (because they use the same beam physics, and give identical results, for propagation to the target).

In Figure 3a, the efficiencies are plotted for three choices of x_{focus} (50 m, 100 m, 200 m). The quotient of the two efficiencies, for a fixed target range (100 m) and a varying focus range, is shown in Figure 3b.

Note that these are two different feasible experiments: the first considers how the fixed-focus lidar responds to scatterers placed at different ranges, and the second varies the fibre-to-lens distance to explore the response to scatterers placed at a fixed range.

Lindelöw [4] notes the disagreement between his approach and the early theory of Sonnenschein and Horrigan [2], which for our current geometry is equivalent to the BPLO approach. This “S&H” paper predates 1.5 μm fibre lidars but applies to them without modification because we assume untruncated pure TEM₀₀ beams throughout. He notes that beams may overfill or underfill the detector, and that the older literature does not consider the modern layouts where received and LO beams co-propagate in optical fibre. But the transmit/BPLO approach remains valid, if the BPLO is correctly calculated (including any variations or truncations or damage of the “photodevice surface”, and any spatial filtering in the fibre). Results may be expressed as a diffraction integral in the detector plane, the target plane or some other convenient plane. The standard approach is correct—or, at least, correctly developed within its stated approximations—and Lindelöw’s, because of inconsistent shifting between ray and beam optics, is not. But the differences are small, of the same order as current errors in calibration, and so not easy to verify in practice—often less than 1 dB (in efficiency or power ratios) or a few % (in the range for peak efficiency). Lindelöw’s -3 dB widths (FWHM) are 20–25% narrower than the conventional predictions (see also [21] and its Figure 2.9). No experimental checks are offered here, and few relevant ones are in the References; it would be good to see more.

Brewer et al. [22], before detailing a full diffraction-integral approach, describe a geometrical approach similar to Lindelöw’s. They are more concerned with bistatic direct-detection lidars for imaging distant targets, but they are still treating single-mode fibres (which, although not carrying LO or BPLO beams, still reject any received light which is not mode-matched or, in geometric optics language, does not fall within the fibre numerical aperture). Some changes or extra details are needed before we apply their results to a monostatic focused coherent lidar: transmit and receive fibre/telescope optics are matched (in their worked example, the transmit aperture diameter must be taken as 10 cm), the far-field assumption is dropped, and care is needed when solving quadratic Gaussian-beam equations.

This last point is already familiar from Hill and Harris [20]. Brewer et al.’s worked example uses a nearly collimated transmitted beam that has a specified width (spot size) at a distant target. In their “far-field” solution a slightly diverging transmitted beam appears to have a virtual beam waist some 67 m behind the lidar. They ignore the slightly converging “near-field” solution with a real beam waist a similar distance in front of the lidar. The choice makes little difference to an ideal free-space imager, but more difference to a lidar used for near-field wind measurements.

Frehlich and Kavaya [23] develop useful general expressions and, after translation between various technical notations, show agreement with several other published accounts including Sonnenschein and Horrigan [2] where there are several misprints and, as occasionally happens, two cancelling factor-of-2 errors; Michael Kavaya has maintained and distributed a list of these and other errata.

Sonnenschein and Horrigan proceed to sum over a large collection $N \gg 1$ of scatterers (representing a random diffuse target) and neither they nor Frehlich and Kavaya consider the single-particle or few-particle cases in detail.

Lawrence et al. [3] show good agreement between their CO₂ lidar and a conventional cup anemometer, but do not analyse the few-particle case.

5.4. Single-Particle Mode and Volume Mode

This interesting distinction in lidar processing should perhaps be better known. It gives an opportunity to recall the many friendships and the long productive career of Dr Rod Frehlich. Indeed, large parts of the present paper and of its References are comments on and restatements of his determined, thorough exploration of radar and lidar physics.

A NASA report by Kavaya et al. [24] includes Frehlich’s comparison of single-particle mode (SPM) and volume mode (VM) measurements. The contributions from scattering particles can be defined in different ways. Essentially, VM adopts the theory and CNR definitions followed here and in most of the literature, while SPM measures a particle’s peak contribution—that is, if a particle is detected in the probe volume, then we record the photocurrent contribution from the instant of closest approach to beam axis.

For decades, optical scientists have grumbled about the uncertainties in the shape(s) of laser beam(s), the resulting distortions in fringe patterns, and the positions and trajectories of scatterers. “Non-diametral traverses of particles across the scattering volume have also not been considered in the broadening formulae. There is no convenient way to account for these effects . . . ” [25]; see also Schulz-Dubois [26]. Our point here is that these two definitions give different results for range dependence, as shown in Figure 4. Errors may arise in calibration and interpretation unless a consistent choice is made.

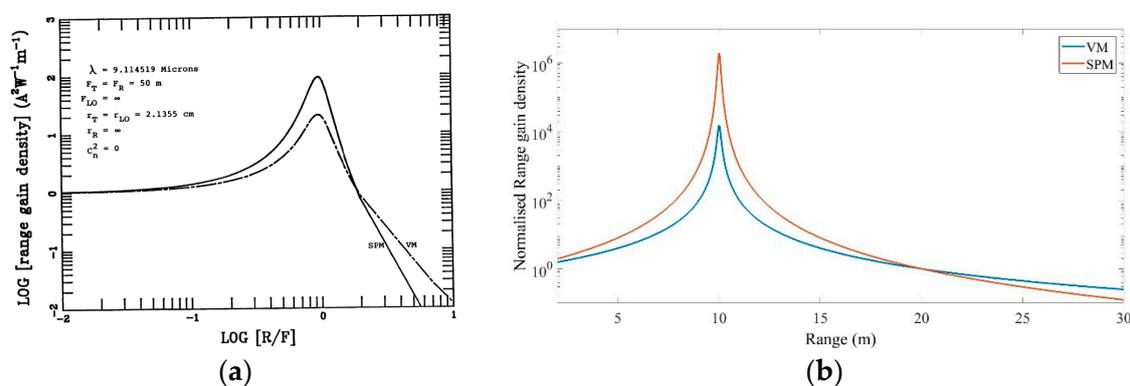


Figure 4. (a) “Normalized range gain density” from Kavaya et al. [24], for volume mode and single-particle mode ($\lambda = 9.11 \mu\text{m}$, beam radius to $1/e^2$ intensity points at the lens = 30.2 mm, focus distance = 50 m); (b) Same quantities calculated for a ZephIR-type lidar ($1.57 \mu\text{m}$, 25 mm, 10 m).

6. Simulation of Multiple Scatterers in a Lidar Beam

The simulated measurement volume is assumed here to have rectangular (usually square) cross section, with room for circles around the beam axis whose radii are at least $2w(z)$ for all relevant z . The decrease in average intensity as we move away from the focus along the z axis follows, as in Equations (1)–(5), a Lorentzian (rather than the steep Gaussian function for x and y), but for most practical purposes a choice of $(-3z_R, 3z_R)$ for the z -axis limits will include “almost all” the relevant scattering. Indeed, these choices are overkill in the sense that many of the simulated scatterers are illuminated so feebly that their contribution to the results is negligible; a conservative choice increases confidence in the simulations.

Thus, we choose a focused lidar beam geometry (a certain transceiver aperture and focused Gaussian beam parameters) and then simulate the movement of identical scattering particles which enter a measurement “box” at random times and at random positions on the x - z “face” of the box, the box being sufficiently large that any particles arriving outside it would not be “significantly” illuminated. Each scatterer has a random Poisson-process arrival time and crosses the beam at uniform speed, with a constant x and constant z each drawn from a uniform random distribution.

The scattering probe volume is shaped by our focused monostatic lidar [24], so the illumination is non-uniform, and the wavefronts encountered away from focus are non-planar. At any instant, $N_{\text{tot}}(t)$ scatterers are within the box, and the time or ensemble average of this number is proportional to the rate of the Poisson process and to the size of the box (the assumed measurement volume). For precise experimental checks, the measurement volume should be carefully specified. For example, in Figure 1b, the physically relevant volume lies to the right of the transmitting lens. Regions to the

left are inside the lidar itself, and presumably not filled by scatterers; thus the “relevant” volume may not be symmetric around the beam focus.

At any instant, the number of scatterers “contributing”, in the sense that they are significantly illuminated, is at most N_{tot} and may (if we size the box conservatively) be much smaller. During some time intervals, there may be no scatterer in the box or none that contributes significantly. The number of scatterers illuminated in a fixed measurement volume will generally change with time (as would happen in real life, because of random or wind-blown scatterer movement). The assumed measurement volume itself may be altered during a simulation run (notably because of a varying beam focus), but of course, consistency checks are advisable—runs for different measurement volumes may take very different amounts of computation, but their average results should be similar so long as the box and N_{tot} are sufficiently large.

There are many possible experimental scenarios, and many fluctuation processes and associated timescales. It is important to keep in mind what is meant by an “average”—whether over time, probability distribution, or ensemble of experiments.

For a collection of scatterers distributed (on average uniformly) across a slice of atmosphere at a certain range z , and for randomly distributed phase terms, the mean squared quantity $\langle |i(t)|^2 \rangle$ contributed by that slice is proportional to the scatterer number density and inversely proportional to the illuminated area of the slice—so it peaks at the beam focus, where the area is least (see also Section 4 of [1] for this “1/area” relation). If we fix the number density and vary the beam area, the number of scatterers involved scales with beam area; but, because the average contribution per scatterer to the detector *power* scales as the intensity squared (i.e., as the inverse square of beam area), the net result is a power scaling with 1/area. We also see that the number of scatterers in a thin slice is proportional to beam area A , and the beam intensity at each scatterer is proportional to $1/A$. So, the fractional backscattered power per slice is a constant, independent of beam size. So long as we continue to use a time-averaged definition of power, this conclusion applies when we vary the beam size for arbitrarily small seeding densities.

The heterodyne current due to several scatterers is the vector sum of the individual currents, one for each scatterer as in Equation (7); for matched transmit and BPLO beams, each term can be considered to vary as the beam intensity at the scatterer. The mean square of the current (averaged over all relative phases as above) equals the sum of the individual squares; this is the same result as for the mean square of the length of a 2D random walk.

We can illustrate some of the familiar results in the literature [1]. We assume a fixed transmit power, a fixed LO power, and the usual optical arrangement described above (focused monostatic coherent lidar, matched transmit and BPLO). First, Figure 5 shows the chirping behaviour of scatterers that cross the beam away from the focal plane.

This modelling approach was discussed in [1], and similar time-frequency plots are well known. Renard et al. [27] discuss airborne lidar options, and the general features of chirps are clear in their exemplary figure (reproduced here as Figure 6).

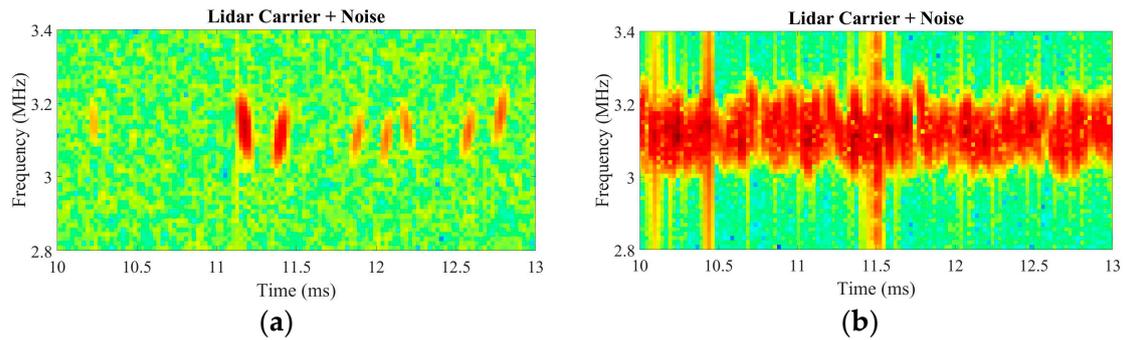


Figure 5. (a) Spectrogram of simulated lidar data for several scatterers crossing the beam. At any instant, during this short (~3 ms) simulation, the number of scatterers significantly illuminated within the probe volume is either 0 or 1; (b) Spectrogram of simulated lidar data (same duration ~3 ms) for a larger density of scatterers crossing the beam. At any instant, it is likely that several scatterers ($N > 1$) are significantly illuminated. The scatterers are given small random Doppler shifts to simulate slightly oblique paths across the beam. As usual with spectrograms there are smearing and rippling effects, and compromises are needed to choose filters and block lengths.

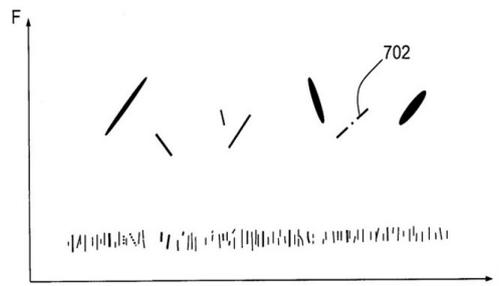


Figure 6. Example of time-frequency plot in Renard et al. [27]. See also Baral-Baron [28].

Second, we model an experiment where the beam waist position and radius are varied (for constant beam width at the lens); the target is modelled as randomly generated scatterers (of equal reflectance) within a box that encloses the waist as described above. The scatterer density (m^{-3}) is constant on average. We see in Figure 7 the variations in probe volume, CNR, and normalised second moment of intensity.

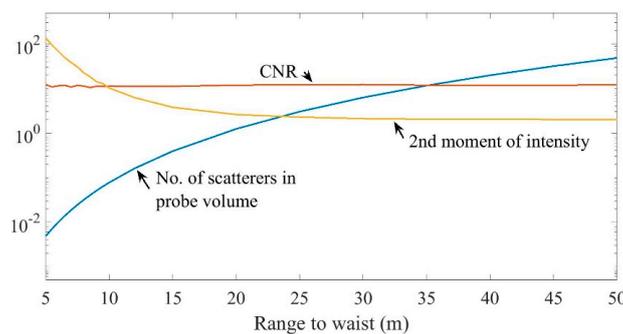


Figure 7. Results from one simulation of a lidar with uniform aerosol seeding density, fixed aperture, and varying focus range. Average CNR is almost independent of focus range, despite the variation of 4 orders of magnitude in probe volume. Normalised 2nd moment of intensity $\langle I^2 \rangle / \langle I \rangle^2$ is large at short ranges but closely approaches 2 (characteristic of complex Gaussian statistics) if there are “many” randomly arranged scatterers (say a few tens or more) in the probe volume. Figure 1.4 in Banakh and Smalikhov [9] is similar, illustrating their analytical expressions for “echo power” and the number of efficiently scattering particles, with specific model assumptions for a CO_2 atmospheric lidar ($\lambda = 10.6 \mu m$).

Here, the simulation assumes constant measurement time (120 ms) at each focus range. At the shortest ranges, there are so few scatterers contributing that the results for CNR and 2nd moment are noisy. (The “No. of scatterers in probe volume” plotted is a theoretical average number given by $(2\lambda^3/\pi^2) \times (\text{focus distance}/\text{beam waist radius})^4 \times (\text{seeding density})$, so it is not noisy). As the range increases to tens of metres, the number of scatterers contributing rises quickly, because of this fourth-power dependence of probe volume on range [24]. The CNR is almost constant with range as we expect, and the 2nd moment of intensity falls to very nearly 2 when the probe volume contains tens of scatterers. This computer run uses $\lambda \sim 1.57 \mu\text{m}$ and $w_0 \sim 15 \text{ mm}$, with a seeding density of $500,000 \text{ m}^{-3}$, and with the scatterer reflectance and shot noise level set to give a CNR of about 10 in a bandwidth of 10 MHz. The 2nd moment is more than 100 at the shortest range of 5 m but within 1% of 2.00 for ranges above 45 m. Of course, more stable results at the shorter ranges can be obtained if (for example) we simulate a different experiment where the total number of contributing scatterers (rather than the measurement time) is fixed at each range.

Such a simulation is easily run for hundreds or thousands of scatterers on a PC. This was harder 40 years ago, for example when Mayo [29] modelled the “triple stochastic” physics of turbulent flow, particle arrivals, and photodetections.

Two more papers emphasising the practical side rather than simulations, but relevant to understanding the various models, are the following:

Jarzemski and Srivastava [30] discuss some interference effects for the case of two illuminated particles and treat some experimental time series at length. But they restrict their explanations to particles near the beam focus and use a corresponding simplification of the Sonnenschein and Horrigan equation. They neglect the information in the envelope of the two-scatterer lidar output, saying that “it does not contain phase properties of backscatter”. For the linked amplitude and phase modulations in such FM interference experiments, see Hill et al. [31] and its references.

Harris et al. [32] draw fresh attention to the possibilities for single-particle lidar anemometry and to the need to consider different fluctuation statistics because the probe volume varies dramatically; they do not consider time series analysis, chirp behaviour, or spectral moment estimation in detail.

This section has only sketched a very wide topic of relevance to short-range and/or thinly seeded sensing (e.g., wind tunnel lidar and high-altitude anemometry); there are whole industries based on particle/suspension analysis through optical scattering.

7. Direction Sensing for VAD Lidars with Sign Ambiguity

7.1. Deciding the Sign

We mention two related issues in coherent lidar Doppler sensing. One is the requirement to attach signs (+ or −) as well as magnitudes to these estimates. The other is the estimation of the overall target motion vector after the estimation of several Doppler shifts. Apart from brief comments here and in Section 7.2 below, we cannot review the huge literature on estimating scatterer speed through the Doppler effect. For the moment we assume that the detector output yields a well-defined spectral peak and thus some reasonable mean-Doppler estimate.

Suppose for illustration that we are trying to estimate the magnitude and direction of the wind and using a coherent lidar and a “conical scan” or “velocity-azimuth display” (VAD) mode. With a system as described above, a change of sign (+/−) of the velocity component does not affect the detector output: that is, one lidar measurement (one brief “snapshot” at a single viewing angle) does not tell us whether the wind component along the line of sight is towards the lidar or away. Moreover, for a uniformly flowing wind, although a collection of several or many measurements at different viewing angles can tell us the wind speed (magnitude of wind vector), it still does not resolve this 180-degree ambiguity. For example, the ZephIR lidar collects approximately 50 separate estimates of wind velocity component per second, spaced evenly around its 360-degree conical scan, and the

resulting “figure of 8” pattern allows an accurate estimate of wind speed. However, the same pattern would appear for a uniform wind flowing in the opposite (180 degrees different) sense (Figure 8).

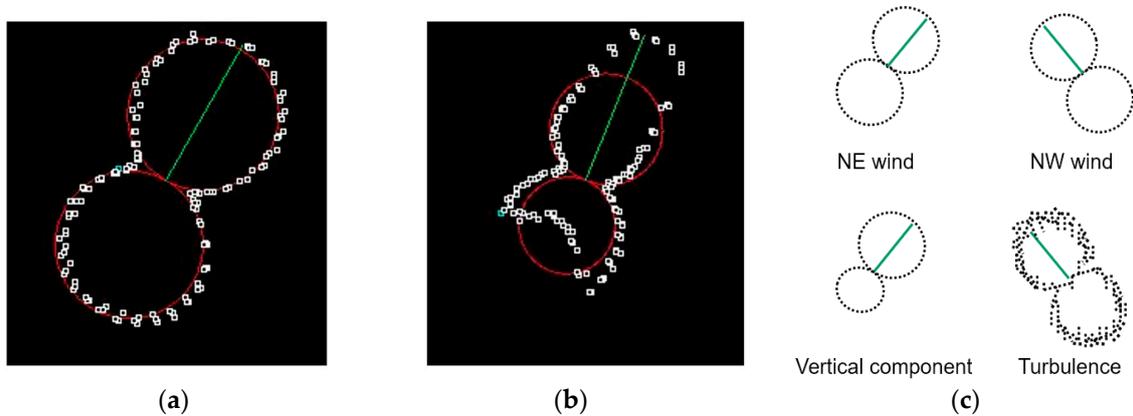


Figure 8. The appearance of Doppler estimates in a polar plot for a conical-scan lidar without direction (+/−) discrimination. (a) “Figure-of-8” for a steady wind; (b) Perturbations soon after the passage overhead of an A320 jet. Data from trials with ZephIR at Birmingham airport; (c) Simplified guide to interpretation of “figures-of-8” (with thanks to Dr M Harris). The results from Hill et al. [33] here and in Figures 9a and 10 below are reproduced by permission of QinetiQ Limited.

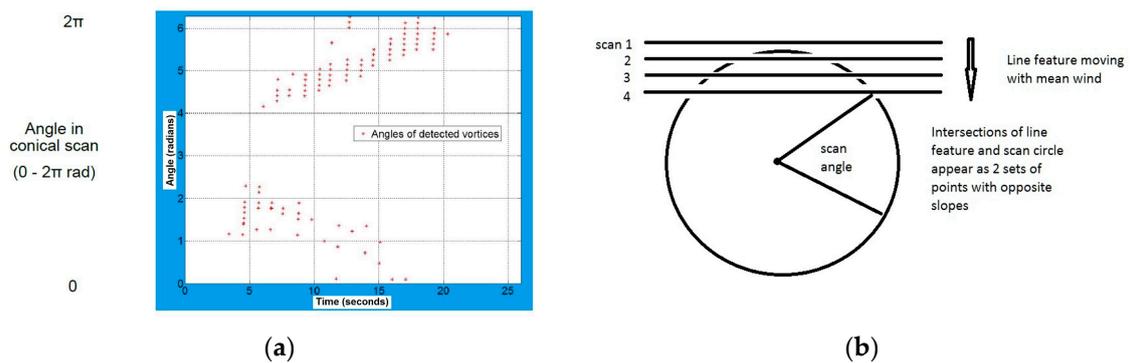


Figure 9. A passing aircraft leaves a pair of trailing wake vortices which form (briefly and approximately) a stable “line feature” in the atmosphere. This line drifts with the mean wind across the scan circle defined by the measurement points of a conical-scan lidar, and the angles where the lidar detects a significant disturbance are recorded and plotted. (a) Example of data from trials with ZephIR at Birmingham airport, adapted from Hill et al. [33]; (b) Sketch of the principle. This plan view is simplified; in practice the lidar processing and scanning may take account of vertical as well as horizontal drifts and may have complicated algorithms for estimating the vortex and wind parameters.

Wind speeds are routinely estimated with fewer and/or less widely spread viewing angles if the CNR is sufficient and the wind is known (or assumed) to be steady so that we are confident of a sinusoidal variation of speed with scan angle or equivalently of the typical figure-of-8 polar plot of Figure 8. Schwiesow et al. [34] reported success with as little as 1/16 of a full scan (that is, less than half a radian), and Leosphere’s Windcube lidars typically use 4 or 5 separate, fixed beam directions [35]. If we assume a steady wind and thus a figure-of-8 of known symmetric shape, far fewer than the 50 estimates per scan suffice for a good estimate of its size and orientation. But the same problem arises in all such cases: we need some extra information in order to determine the sign of the wind.

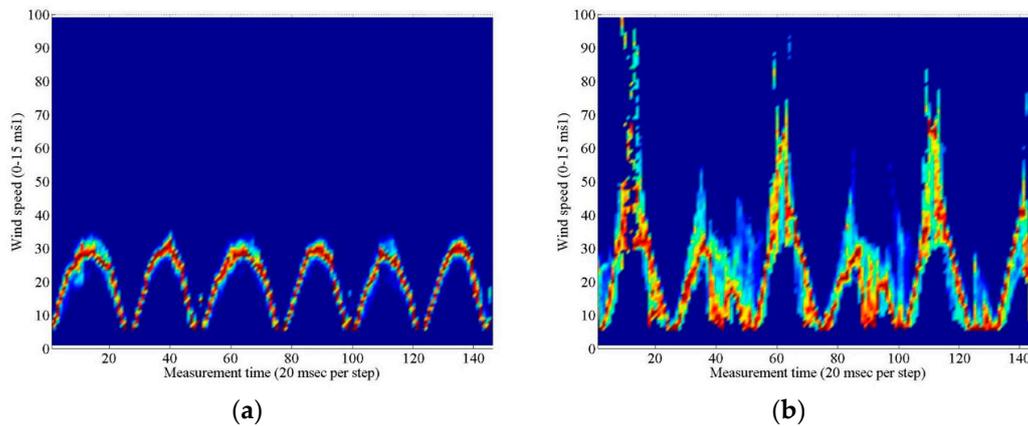


Figure 10. Spectrogram plots of ZephIR lidar data from airport trials. (a) Rectified cosine wave for a fairly steady airflow on the southern approach to Birmingham airport [33]. A polar plot would produce a “figure of 8” similar to Figure 8a; (b) Perturbed cosine wave after the passage of a Boeing 777. The vortex lines provide mainly positive Doppler contributions during one half-scan (vortex flow towards the lidar, faster than the mean wind) and mainly negative ones during the other. The scan angles at which the perturbation is judged “significant” according to some metric can be identified and added to a plot such as the one in Figure 9a.

Of course, a separate non-lidar instrument may be added to measure, at one point or more, the wind direction. In “ordinary” atmospheric conditions, it is sufficient to measure with a conventional anemometer (vane/cup/ultrasonic) near the surface, and trust that the wind direction does not veer sharply between measurements and between adjacent measurement heights; in this way, by assuming a reasonable continuity, we can assign the wind vector to the correct quadrant at successively greater heights [36]. But what if such a non-lidar instrument is unavailable or inappropriate?

There are several ways forward. An artificial frequency shift can be introduced, typically by an acousto-optic modulator, so that the zero-Doppler region appears shifted from zero frequency in the detector output. Or the in-phase and quadrature (I&Q) components of the complex carrier can be separately detected, so that their relative phase can be tracked. Both approaches tend to involve significant extra hardware and cost.

A third way, cruder but feasible, relies on the fact that real wind flows are not wholly “uniform”; they always have fluctuations, for example of density, velocity, and backscatter. Nonetheless we usually assume that they remain “uniform” or well-behaved over some extent and duration: that is, we adopt a “frozen flow hypothesis” whereby the spatial patterns containing these fluctuations are being transported with some overall mean wind motion. The patterns need only be sufficiently distinct, and evolve sufficiently slowly, for us to track them with adequate confidence: it is a matter of “sufficient” CNR and of timescales. If we look at the same small volume of atmosphere at intervals of hours or many minutes, we may see no significant correlations of the departures from uniformity; on a timescale of seconds or tenths of a second, we may have considerable success.

For example, Figure 9 shows unambiguous direction-sensing for a ZephIR lidar without frequency offset or I&Q processing [33]. In this case, the atmosphere is perturbed by aircraft wake vortices that typically last for many seconds and are (approximately) transported across the sky with the local mean horizontal wind component, while descending slowly towards the ground. During a series of one-second conical scans the lidar detects the perturbations caused by the vortex pair. For a typical single scan, there are several detections (marked by the small red symbols in Figure 9a); as time evolves, we see that these detections cluster in two groups that diverge. For one group the scan angles increase, and for the other they decrease, because the line of the vortex pair—which is approximately a straight line, considered on this scale of a few hundred metres—is entering the lidar’s conical-scan pattern. For a line feature leaving the scan pattern, the two sets converge, and this distinction tells us the

direction of the wind. For an ideal thin line, the two sets of detections coalesce when the line becomes tangent to the scan circle; in practice the region of lidar detection sensitivity is 10–30 m across, and so is the pair vortex structure from a typical medium aircraft. So, on this type of plot, the detections are clustered only loosely, and we expect scatter of several metres or more (or the corresponding spreads in angle/time), but after a few consecutive seconds of scanning, we are in no doubt of the overall tendency (divergence in angle of the two groups, or convergence?) and thus in no doubt of the +/– sign decision.

A conventional X-Y time-frequency (rather than polar) plot shows a rectified cosine wave when the air is relatively unperturbed. When an aircraft passes overhead and the trailing vortices are sampled by the lidar scan, there are considerable additions to (and subtractions from) the cosine function (Figure 10).

The ZephIR example above, an early demonstration with eyesafe fibre-optic lidar, is deliberately simple: the airflow is strongly perturbed, and we also have prior expectations about the general form of the pair of counter-rotating vortices. But the principle [37–41] applies to the natural fluctuations of velocity, backscatter etc. in “ordinary” air, with no helpful jet aircraft passing, and hence was of interest for Malvern and Risø work with ZephIR-type lidars for wind applications. Evidently the “metrics” for correlation-tracking of this sort (typically based on the first few spectral moments viz. carrier strength, mean Doppler, and Doppler spread) will be adjusted to suit if we believe that particular geometrical features may be present; for example we may expect to see these linear vortices, or 2-D loops, or sheets of separated air flow, or 3-D volumes of turbulent air; but even a simple correlation plot without special assumptions may strongly indicate the likely sense of the overall wind motion. Many demonstrations are now in the literature with larger measurement sets, more detailed correlation and wavelet algorithms, and heftier processors [42,43].

7.2. Some Issues in Frequency Estimation

The estimation of a Doppler shift in wind lidar is usually treated as a “mean frequency” problem and attacked by forming power spectral estimates or autocorrelation functions. For light winds, there are obvious technical issues.

First, in the absence of a frequency-shifting scheme such as the modulator mentioned above, the informative carrier lies close to DC, where interferences such as local oscillator noise and laser relative intensity noise (RIN) are usually strongest. The influences of “intensity” noise and “phase” noise on frequency estimation depend on the estimator used, and we need to know the complex (intensity and phase) variations associated with laser instability; an estimate of the power spectral density attributed to RIN is insufficient. Evidently a two-point or “instantaneous frequency” algorithm, as it relies on the phases alone, should be unaffected by AM when the CNR is high, but other algorithms use more or less of the AM or envelope information—for good reasons—and will be correspondingly affected.

Second, how is a wind speed defined when the measurement volume—usually, as we have seen, a long thin volume extending through several or many metres—contains some scatterers approaching us and some scatterers receding? The coherent addition of complex components with different Doppler signs (clockwise and anticlockwise phasor rotations) does not necessarily give the same result as the incoherent addition of the sign-insensitive outputs.

This is an issue not just of “ambiguity” but of distinctly different results in the detector output. If we suppose for simplicity that there are only two scatterers present in the beam, one moving towards the lidar and one away, and that their individual contributions to the detector output have equal magnitudes and opposite signs, the net result in a frequency-shifted (direction-sensitive) coherent lidar is the sum of two oppositely rotating phasors.

Also, the common use of a blocking or highpass filter (to remove both carrier and dominant noise near DC) will in any case cause a bias in subsequent Doppler estimation. For example, the lowest few

hundred kHz of ZephIR data (in practice, two or three frequency bins of the spectrogram display) are not used in Figure 10.

8. Errors and Bias in Doppler Lidars for Steady Winds

Finally, we discuss two examples of how (according to the manufacturer’s publications) the extraction of wind velocity from the lidar detector output must be imperfect. Any error or bias is likely to be very small, but lidar users—including those involved in calibration and verification studies—are advised to check if and how any corrections are applied.

There is a practical issue in conical-scan or sector-scan lidars where lidar measurements are accumulated (usually in the form of averaged power spectral estimates or autocorrelations) during an appreciable interval and hence over an appreciable range of angles. For example, the standard ZephIR choice of 50 measurements per one-second scan means that each measurement represents a “wedge” or segment of angular extent $360/50 = 7.2$ degrees of scan. This, in turn, may mean a small but not entirely negligible error in estimates of the figure-of-8 size (wind speed), orientation (wind direction), and lobe symmetry (vertical wind component).

Suppose the lidar forms a spectral estimate at a given scan angle θ (not to be confused with the phase angle above) by accumulating short-term estimates over the range $\theta - \delta$ to $\theta + \delta$ and then finds a mean frequency (first spectral moment) from the accumulated estimate.

Denote the wind heading by θ_{true} , so that an ideal non-direction-sensing lidar would detect a maximum line-of-sight component at θ_{true} and $\theta_{\text{true}} + \pi$ radians, and a minimum line-of-sight component at $\theta_{\text{true}} + \pi/2$ and $\theta_{\text{true}} + 3\pi/2$ radians. Short-term estimates are incoherently accumulated at a steady rate (per radian of scan). The power per estimate will fluctuate around a steady mean (given assumptions of uniform operating conditions, backscatter etc.). But, neglecting the fluctuations, we may assume that each incremental scan segment ($\theta, \theta + d\theta$) contributes a spectral increment of the same total power centred around a Doppler frequency proportional to $\text{abs}(\cos(\theta))$.

The measured wind speed can thus be taken as proportional to the average of $\text{abs}(\cos(\theta))$ over the interval $(\theta - \delta, \theta + \delta)$. Evidently, the wind speed may be underestimated near $\theta = \theta_{\text{true}}$ and $\theta_{\text{true}} + \pi$ and overestimated near $\theta_{\text{true}} + \pi/2$ and $\theta_{\text{true}} + 3\pi/2$. The direction-sensing lidar behaves slightly differently because of the signed function $\cos(\theta)$; the mean frequency is then proportional to $\cos(\theta) \cdot \sin(\delta/\delta)$.

First, we show results (as functions of θ) for 10 measurements per scan, that is $\delta = \pi/10$. It is assumed that scan segments are contiguous (without overlapping). In Figure 11a the scan origin $\theta = 0$ (the midpoint angle for the first segment) coincides with the true wind direction. In Figure 11b we assume the maximum possible “offset” so that the true wind heading coincides with the edge of a segment, that is, the angle where one segment finishes and the next one begins.

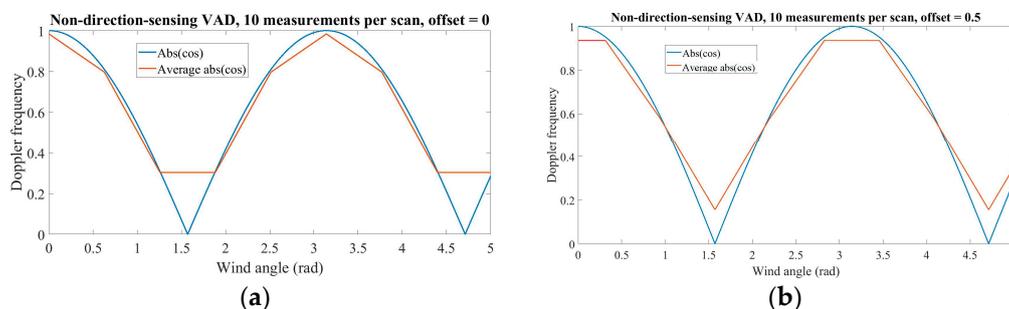


Figure 11. Rectified cosine function and its time-windowed average, representing the Doppler-estimating behaviour of a non-direction-sensing lidar; compare this with the typical ZephIR display in Figure 10a. For each scan rotation, the average of the cosine is obtained for each of 10 contiguous segments (arcs) of $\pi/5$ radians. (a) The true wind heading coincides with the midpoint angle of a segment (i.e., zero offset); (b) The true wind angle coincides with the edge of a segment (i.e., offset = 0.5 segment).

The “errors” shown here might be unacceptable. Figure 12 shows similar results for the ordinary ZephIR case of 50 measurements per scan, that is $\delta = \pi/50$ or $2\delta = 7.2$ degrees. Now the bias would almost certainly be negligible in practice. It is still noticeable very close to $\pi/2$ and $3\pi/2$ but, usually, these regions are not used—any measurements within them are discarded, because (as just mentioned) we do not trust the detector noise floor in the lowest spectral bin(s)—and the effect of this slight data loss on curve-fitting algorithms is unimportant.

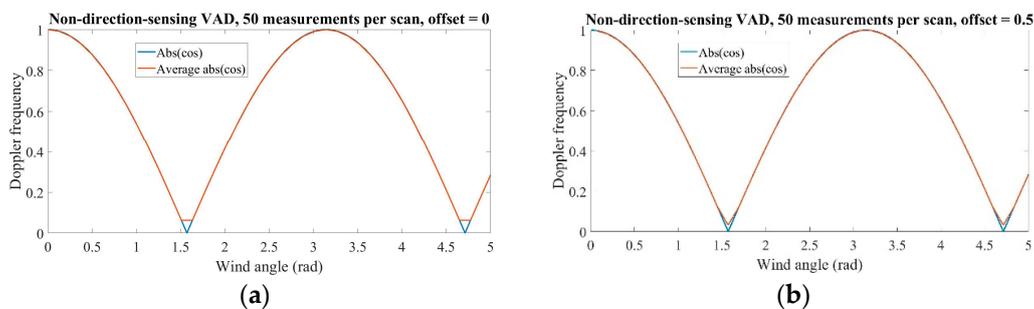


Figure 12. As for Figure 11, with 50 contiguous segments per scan rotation. The bias is now negligible. (a) The true wind heading coincides with the midpoint angle of a segment (i.e., zero offset); (b) The true wind angle coincides with the edge of a segment (i.e., offset = 0.5 segment).

Non-scanning lidars (such as standard WindCube) do not have this problem to the same extent, as their field of view is tighter—essentially, the beam angular width of at most a few milliradians rather than the wedge of a few degrees. But they have another issue if the speed and direction are estimated directly (as described in [19]) from a small number, typically 4, of fixed-angle measurements. When a horizontal wind speed is obtained as $V_h = \sqrt{u^2 + v^2}$, where its orthogonal components u and v are derived from lidar line-of-sight speed estimates, any small zero-mean random noises on these estimates will create small zero-mean random noises on u and v and (in general) non-zero-mean random noise on V_h . For zero-mean Gaussian noises on the initial estimates, this bias of V_h will be positive.

For each type of lidar, given our assumptions about uniform horizontal winds and our known beam parameters, a compensation function could reduce the bias of the individual measurements (before curve-fitting an equation for the wind vector). Or the biased measurements could be fitted to a modified equation that included terms to describe the bias.

9. Conclusions

This paper is a companion to [1] and other recent overviews of fibre-optic lidars; a common theme is the reexamination of basic equations and principles in radar or lidar. In [1], the emphasis was on the hundreds of fibre-optic lidars installed today for wind and especially wind turbine farm measurements; again we stressed what happens when the number of scatterers (components) changes through 0, 1, 2, . . . up to the “many” of near-Gaussian statistics. The present paper has considered the simulation of few-scatterer and multi-scatterer lidar experiments, and some problems (and solutions) for Doppler-sign-insensitive lidars.

We reviewed the BPLO description of the integrals for lidar CNR and efficiency. In a monostatic Doppler lidar with matched LO and transmit beams, “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” [20]. The sensitivity function Q used by Banakh, Smalikho et al. [44] is an exact Lorentzian, and (once we make the necessary adjustments in notation) it is the same Lorentzian used by Sonnenschein and Horrigan [2], Frehlich and Kavaya [23], Qi Hu et al. [45], and others. We contrasted this with Lindelöw’s description of

“collection efficiency”, which departs from the “1/area” rule. A fair and separate question, though, is whether the usual calibration experiments are sufficiently precise to tell these two descriptions apart.

We reviewed the single-particle mode (SPM) and volume mode (VM) modelling of Frehlich et al. and some numerical modelling of lidar detector time series and statistics.

We have tried to highlight and explain behaviour that may be observed from a modern coherent lidar used at short ranges (e.g., in a wind tunnel) and/or with weak aerosol seeding. It is worth noting that non-Gaussian statistics (and opportunities for improved processing) may arise at focus distances of 10–20 m, even when there are 10^6 – 10^7 scatterers per cubic metre.

Finally, we reviewed some issues in direction sensing, sign ambiguity, and limited angular resolution for standard Doppler-sensing coherent lidars.

It has been interesting to watch these fibre-optic lidars grow into accurate reliable calibrated tools of the wind energy trade, with accompanying expectations and responsibilities.

Acknowledgments: I thank the two referees for helpful suggestions. For a better understanding of the sensing needs of the wind energy industry, I am indebted to many colleagues in the UpWIND consortium led by the Danish Technical University (DTU). I have been helped by technical discussions with Jean-Pierre Cariou (Leosphere), Mike Courtney, Mike Harris (ZephIR), Michael Kavaya, Petter Lindelöw, Jakob Mann, Guy Pearson (HALO Photonics), Anders Tegtmeier Pedersen, Kevin Ridley, and Michael Vaughan. These colleagues and their organisations are not responsible for any errors and omissions here.

Conflicts of Interest: The author declares no conflict of interest.

Glossary

BPLO	Back propagated local oscillator. A fictitious aid in analysing heterodyne lidars.
CNR	Carrier-to-noise ratio. Strictly this should mean a ratio C:N but it often appears as a quotient C/N or a decibel measure. In lidars with nonlinear processing steps such as frequency estimation, the CNR or carrier-to-noise (at the detector, before processing) is often distinguished from SNR or signal-to-noise (after processing); these two noises “N” may have quite different statistics.
CW	Continuous wave.
FM	Frequency modulation.
I&Q	In-phase and quadrature (the two parts of a complex quantity, usually the detector output current, giving a phase angle $\tan^{-1}(Q/I)$).
LO	Local oscillator.
SPM	Single-particle mode: lidar scattering events are individually registered and treated as follows: <i>if</i> a particle is detected in the beam, <i>then</i> the peak resulting detector output counts towards an overall metric. See VM.
TEM ₀₀	Transverse electromagnetic mode of lowest order; the first in a nominally complete orthogonal set of “Gaussian beams” into which a propagating monochromatic lidar beam can be decomposed.
VAD	Velocity-azimuth display. Usually refers to a conical-scan lidar.
VM	Volume mode: the entire time series of a scattering event counts towards an overall metric. This is normal procedure in an atmospheric lidar illuminating a volume densely seeded with scatterers; individual scattering events are not usually isolated and examined even if that is technically feasible. But other metrics exist (see SPM) and may show different dependences on target range.
Windcube	A brand of pulsed lidars operating near 1.5 μm ; based on research at ONERA, Palaiseau, and later developed and marketed by Leosphere.
ZephIR	A brand of CW lidars operating near 1.5 μm ; based on research at QinetiQ, Malvern, and later developed and marketed by ZephIR Lidar.

References and Notes

- Hill, C.A. Modern fibre-optic coherent lidars for remote sensing. *Proc. SPIE* **2015**, *9649*. [[CrossRef](#)]
- Sonnenschein, C.M.; Horrigan, F.A. Signal-to-noise relationships for coaxial systems that heterodyne backscatter from the atmosphere. *Appl. Opt.* **1971**, *10*, 1600–1604. [[CrossRef](#)] [[PubMed](#)]
- Lawrence, T.R.; Wilson, D.J.; Craven, C.E.; Jones, I.P.; Huffaker, R.M.; Thomson, J.A.L. A laser velocimeter for remote wind sensing. *Rev. Sci. Instrum.* **1972**, *43*, 512–518. In their lidar (our Figure 1a) the received light must make at least one extra reflection after the beamsplitter. In a common UK terminology this is an autodyne lidar, because the backscattered radiation “is collected by the telescope and allowed to reenter the laser cavity where it is amplified”. The overlap of the reemerging radiation with the (BP)LO may or may not be stable and easily described—this also depends on the active cavity mode structure and the radiation frequencies. Most fibre-optic lidars, including the one in our Figure 1b, are not intended as “autodyne” in this sense. The word has various meanings.
- Lindelöw, P. Fiber Based Coherent Lidars for Remote Wind Sensing. Ph.D. Thesis, Technical University of Denmark, Lyngby, Denmark, 2008.
- Henderson, S.W. Review of fundamental characteristics of coherent and direct detection Doppler receivers and implications to wind lidar system design. In Proceedings of the 17th Coherent Laser Radar Conference, Barcelona, Spain, 17–20 June 2013.
- Hao, C.H.; Guo, P.; Chen, H.; Zhang, Y.C.; Chen, S.Y. Determination of geometrical form factor in coaxial lidar system. *Proc. SPIE* **2013**, *8905*. [[CrossRef](#)]
- Jakeman, E.; Ridley, K.D. *Modeling Fluctuations in Scattered Waves*; CRC Press: Boca Raton, FL, USA, 2006. See chapter 4’s review of random-walk models and partly developed speckle.
- Pusey, P.N.; Schaefer, D.W.; Koppel, D.E. Single-interval statistics of light scattered by identical independent scatterers. *J. Phys. A* **1974**, *7*, 530–540. [[CrossRef](#)]
- Banakh, V.; Smalikhov, I. *Coherent Doppler Wind Lidars in a Turbulent Atmosphere*; Artech House: Norwood, MA, USA, 2013.
- Rye, B.J. Spectral correlation of atmospheric lidar returns with range-dependent backscatter. *J. Opt. Soc. Am.* **1990**, *7*, 2199–2207. Here Barry Rye thanks the late Ken Hulme for emphasising the superposition of scattered “copies”.
- Siegman, A.E. The antenna properties of optical heterodyne receivers. *Appl. Opt.* **1966**, *5*, 1588–1594. Siegman refers readers to his 1964 IEEE-MTT conference paper and to an unpublished memo by R E Brooks of the TRW company.
- Zhao, Y.; Post, M.J.; Hardesty, R.M. Receiving efficiency of monostatic pulsed coherent lidars. 1: Theory. *Appl. Opt.* **1990**, *29*, 4111–4119. [[CrossRef](#)] [[PubMed](#)]
- Wallner, O.; Winzer, P.J.; Leeb, W.R. Alignment tolerances for plane-wave to single-mode fiber coupling and their mitigation by use of pigtailed collimators. *Appl. Opt.* **2002**, *41*, 637–643. [[CrossRef](#)] [[PubMed](#)]
- Tacke, M. The influence of losses of hollow dielectric waveguides on the mode shape. *IEEE J. Quantum Electron.* **1982**, *18*, 2022–2026. [[CrossRef](#)]
- Kavaya, M.J. Coherent lidar: Factors of two among friends. In *NOAA Working Group on Space-Based Lidar Winds*; NASA/Langley Research Center: Hampton, VA, USA, 2002.
- Hill, C.A. Calibration and optical efficiency of fibre-based coherent lidars. In *UpWIND Report*; Technical University of Denmark: Roskilde, Denmark, 2010.
- Lindelöw, P.; Courtney, M.; Parmentier, R.; Cariou, J.P. Wind shear proportional errors in the horizontal wind speed sensed by focused, range gated lidars. *IOP Conf. Ser. Earth Environ. Sci.* **2008**, *1*. [[CrossRef](#)]
- Cariou, J.P. *Pulsed Coherent Lidars for Remote Wind Sensing*; Remote Sensing Summer School at Risø-DTU: Roskilde, Denmark, 2010.
- Cariou, J.P.; Boquet, M. LEOSPHERE pulsed lidar principles. In *Technical Report for UpWIND Work Package*; Technical University of Denmark: Roskilde, Denmark, 2010.
- Hill, C.A.; Harris, M. Lidar measurement report. In *Technical Report for UpWIND Work Package 6*; Technical University of Denmark: Roskilde, Denmark, 2010.
- Beuth, T. Analyse und Optimierung von Fokussierten LiDAR-Systemen für Windkraftanlagen (Analysis and Optimisation of Focused Lidar Systems for Wind Turbines). Ph.D. Thesis, Karlsruhe Institute of Technology, Karlsruhe, Germany, 2016.

22. Brewer, C.D.; Duncan, B.D.; Barnard, K.J.; Watson, E.A. Coupling efficiencies for general-target-illumination lidar systems incorporating single-mode optical fiber receivers. *J. Opt. Soc. Am.* **1998**, *15*, 736–747. [CrossRef]
23. Frehlich, R.G.; Kavaya, M.J. Coherent laser radar performance for general atmospheric refractive turbulence. *Appl. Opt.* **1991**, *30*, 5325–5352. [CrossRef] [PubMed]
24. Kavaya, M.J.; Henderson, S.W.; Frehlich, R.G. Theory of CW lidar aerosol backscatter measurements and development of a 2.1- μm solid-state pulsed laser radar for aerosol backscatter profiling. This is Contractor Report CR-4347; NASA: Washington, DC, USA, 1991. The “probe volume” for a reasonably tightly focused Gaussian beam is estimated as $\lambda^3(F/D_0)^4$ where F is focal distance and D_0 is beam diameter at the lens. There are various acceptable versions of this in the literature (with different notations and factors of about π), such as $(8\lambda^3/\pi^2)(f/w_{\text{lens}})^4$, but our $8\lambda^3R_{\text{lens}}^4/(\pi^2w_0^4)$ in [1] seems misprinted.
25. Durst, F.; Melling, A.; Whitelaw, J.H. *Principles and Practice of Laser-Doppler Anemometry*, 2nd ed.; Academic Press: London, UK, 1981.
26. Schulz-Dubois, E.O. Photon correlation techniques in fluid mechanics. In Proceedings of the 5th International Conference, Kiel-Damp, Germany, 23–26 May 1982.
27. A Renard et al., French patent 2 948 459 (24 July 2009), US 2011/0181863 (23 July 2010); X Lacondemine et al., French patent FR 1202729 (12 October 2012), US 8,976,342 B2 (date of patent 10 March 2015). If we know the crosswind V and the chirp duration, we know the particle’s distance from focus (though not whether it lies on the near or far side). If we do not know V , we need both duration and slope (measured, of course, at “sufficient” CNR). These papers correctly explain the beam/particle physics but there are minor problems: First, the radius of curvature is not shown decreasing (to a minimum at z_R) and then increasing again at longer ranges from the waist; second, there is a language or translation difficulty in the US version at: “In other words, at the Rayleigh distance z_R , the radius of curvature of the wavefronts is large and, as they move away from this point 205, their radius of curvature decreases”. The French original is *En d’autres termes, au niveau de distance de Rayleigh z_R , les fronts d’onde sont tres incurvés, puis en s’éloignant de ce point 205, l’incurvation diminue*. But the curvature of the wavefronts, not their radius of curvature, is “large”—a local maximum—at z_R . I thank Xavier Lacondemine and Jean-Pierre Cariou for confirming that the English text is wrong here [1]. The French text is correct but *en s’éloignant de ce point 205* can be misinterpreted.
28. Baral-Baron, G. *Traitements Avancés Pour L’augmentation de la Disponibilité et de L’intégrité de la Mesure de Vitesse 3D par LiDAR, dans le Domaine Aéronautique (Advanced Methods for Increasing the Availability and Integrity of Lidar Airborne 3D Velocimetry)*. Ph.D. Thesis, Supélec, Gif-sur-Yvette, France, 2014.
29. Mayo, W.T., Jr. Modeling laser velocimeter signals as triply stochastic Poisson processes. In Proceedings of the Minnesota Symposium on Laser Anemometry, Minneapolis, MN, USA, 22–24 October 1975.
30. Jarzembski, M.A.; Srivastava, V. Interference of backscatter from two droplets in a focused continuous-wave CO_2 Doppler lidar beam. *Appl. Opt.* **1999**, *38*, 3387–3393. [CrossRef] [PubMed]
31. Hill, C.A.; Harris, M.; Ridley, K.D.; Jakeman, E.; Lutzmann, P. Lidar frequency modulation vibrometry in the presence of speckle. *Appl. Opt.* **2003**, *42*, 1091–1100. [CrossRef] [PubMed]
32. Harris, M.; Pearson, G.N.; Ridley, K.D.; Karlsson, C.J.; Olsson, F.Å.; Letalick, D. Single-particle laser Doppler anemometry at 1.55 μm . *Appl. Opt.* **2001**, *40*, 969–973. [CrossRef] [PubMed]
33. Hill, C.A.; Bennett, J.; Smith, D. Airport trials with the Aviation ZephIR lidar. In Proceedings of the 15th Coherent Laser Radar Conference, Toulouse, France, 22–26 June 2009.
34. Schwiesow, R.L.; Köpp, P.; Werner, C. Comparison of cw-lidar-measured wind values by full conical scan, conical sector scan and two-point techniques. *J. Atmos. Ocean. Tech.* **1985**, *2*, 3–14. [CrossRef]
35. Leosphere. Available online: <http://www.leosphere.com> (accessed on 21 February 2018).
36. The assumption of “reasonable continuity” should be examined:

The teeing area is protected rear and right by a thick stand of tall pine trees, so that, even in a stiff wind, you are often hitting from a pocket of comparative calm. The problem is that the trees end about halfway to the green, which is the point where any wind that happens to be around begins to affect the flight of the ball. Perhaps because this is the lowest section of the course, the further and larger problem is that the wind here swirls a lot. You can glance up at the tops of those big old pine trees to your right and note that they are blowing one way, then flick your eyes over to the green and watch the flag blowing in the exactly opposite direction. Picking the correct club and shot at this hole on anything but a dead calm day is, therefore, always something of a guessing game. (J Nicklaus and K Gorman, *My Story*, Random House, 1997).

The typical working ranges (50–250 m) and Lorentzian-halfwidth range resolutions of the current ZephIR lidar are well suited to modern par 3 holes and hard-to-judge approach shots. It is common for the wind direction on some areas of the author's home course to be quite different from that of the low grim clouds overhead; for example, a prevailing breeze from the west may swirl over the Malvern Hills, forming a roll or vortex structure that creates a westwards motion near ground level. The potentially large and rapid wind variations throughout such hilly or complex terrain are of interest to turbine farm companies as well as golf broadcasters.

37. Eloranta, E.W.; King, J.M.; Weinman, J.A. The determination of wind speeds in the boundary layer by monostatic lidar. *J. Appl. Meteorol.* **1975**, *14*, 1485–1489. [[CrossRef](#)]
38. Kolev, I.N.; Parvanov, O.P.; Kaprielov, B.K. Lidar determination of winds by aerosol inhomogeneities: Motion velocity in the planetary boundary layer. *Appl. Opt.* **1988**, *27*, 2524–2531. [[CrossRef](#)] [[PubMed](#)]
39. Zuev, V.E.; Matvienko, G.G.; Samokhvalov, I.V. Laser sensing of wind velocity by a correlation method. *Izv. Atmos. Ocean. Phys.* **1977**, *12*, 772–776.
40. Zuev, V.E.; Samokhvalov, I.V.; Matvienko, G.G.; Kolev, I.N.; Parvanov, O.P. Laser Sounding of Instantaneous and Mean Speed of Wind Using Correlation Method. Available online: <https://ntrs.nasa.gov/search.jsp?R=19870000871> (accessed on 15 January 2018).
41. Samokhvalov, I.V. *Correlation Method of Laser Sounding Measurements of Wind Speed*; Nauka: Novosibirsk, Russia, 1985. (In Russian)
42. Dérian, P.; Mauzey, C.F.; Mayor, S.D. Wavelet-based optical flow for two-component wind field estimation from single aerosol lidar data. *J. Atmos. Ocean. Technol.* **2015**, *33*, 1751–1778. [[CrossRef](#)]
43. Hamada, M.; Dérian, P.; Mauzey, C.F.; Mayor, S.D. Optimization of the cross-correlation algorithm for two-component wind field estimation from single aerosol lidar data and comparison with Doppler lidar. *J. Atmos. Ocean. Technol.* **2016**, *33*, 81–101. [[CrossRef](#)]
44. Smalikho, I.N. On measurement of the dissipation rate of the turbulent energy with a CW Doppler lidar. *J. Atmos. Ocean. Opt.* **1995**, *8*, 788–793. See also: Mikkelsen, T. 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. In Proceedings of the European Wind Energy Conference & Exhibition 2009, Marseille, France, 16–19 March 2009. This literature defines a probe scale $\Delta z = \lambda(f/w_{\text{lens}})^2$, roughly $\sim\pi$ times longer than the conventional Gaussian beam parameter or Rayleigh range.
45. Hu, Q.; Rodrigo, P.J.; Iversen, T.F.; Pedersen, C. Investigation of spherical aberration effects on coherent lidar performance. *Opt. Express* **2013**, *21*, 25670–25676. [[CrossRef](#)] [[PubMed](#)]



© 2018 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).